KU LEUVEN

ARENBERG DOCTORAL SCHOOL
Faculty of Engineering Science

Mining the ECG: Algorithms and Applications

Carolina Varon

Dissertation presented in partial fulfillment of the requirements for the degree of Doctor in Engineering

April 2015

Mining the ECG: Algorithms and Applications

Carolina VARON

Supervisory Committee:
Prof. dr. ir. P. Wollants, chair
Prof. dr. ir. S. Van Huffel, promotor
Prof. dr. ir. J.A.K. Suykens, promotor
Prof. dr. ir. B. Puers
Prof. dr. L. Lagae
Prof. dr. B. Buyse
Prof. dr. HP. Brunner-La Rocca
(Maastricht University)
Prof. dr. D. Marinazzo
(University of Gent)
Prof. dr. R. Aarts
(Technische Universiteit Eindhoven)

Dissertation presented in partial fulfillment of the requirements for the degree of Doctor in Engineering

April 2015

© 2015 KU Leuven – Faculty of Engineering Science
Carolina Varon, Kasteelpark Arenberg 10 box 2446, B-3001 Leuven (Belgium)

Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotocopie, microfilm, elektronisch of op welke andere wijze ook zonder voorafgaande schriftelijke toestemming van de uitgever.

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm or any other means without written permission from the publisher.

Para mi abuelita Blanca

Acknowledgments

And there I was, looking at the Milky Way, when I remembered that dreams brought me here, why I wanted to study so hard, why I wanted to finish my PhD, and why I want to learn and learn until I explode like a type Ia supernova, as a friend once said. During the last four years I learned that dreams can be transformed, and that I need my dreams to keep going.

First, I want to thank my promotor Prof. Sabine Van Huffel, who gave me the opportunity to start my PhD in Biomed. Sabine, I have learned so much during my time here, and I have so many memories and stories to tell, which means that I will always remember you and thank you for this opportunity. Our talks during our trips in your car, the nice dinners we had in conferences, meetings, and social events always encouraged me to keep going. Even though I had difficult times during my PhD, you were always there trying to help me as much as you could, and you always believed in me. Thank you very much for that, thank you for the support, especially during these last months that felt like eternity. I am sorry for the pranks that Steven and I pulled on you, but I have to say that I laughed very hard and I hope you did the same.

I also want to thank my Co-promotor Prof. Johan Suykens for believing in me since my master thesis. During my PhD I had many meetings with you, where you were always patient to guide me and give me the best advice to direct my PhD. You were always there to listen to my ideas, and steer them in the right direction. I learned so much from you, and I hope we can still work together in the future.

Most of my research work was possible thanks to the close collaboration with Dr. Katrien Jansen, Prof. Lieven Lagae, Prof. Dries Testelmans, Prof. Bertien Buyse, and Pascal Borzée. You were always there with very interesting research questions and challenges that motivated me to learn new methodologies and especially new things in the fields of epilepsy and sleep, that I never imagined I was going to learn. For me, the visits to the hospital, and the contact with

you really showed me that everything I have learned can be used somewhere. More specifically, that my effort and sacrifice could hopefully help someone in the future.

During my PhD I was involved in the Recap project, which gave me several things to think about. I want to thank all the partners in the consortium, especially Stuart, Nicolas, Jan, Michelle, Keji, and Prof. Hans-Peter Brunner-La Rocca, with whom I worked very close to get the best we could from this project. We managed to do it, and people are talking about it now! I also want to thank Siegfried Jaecques for the time we spent traveling to our Recap meetings, for the nice advice you gave me for my future career and for life in general. I will never forget our small accident and your explanation about the non-existence of zombies. Thanks also to the different partners in the NXT_SLEEP consortium, and to Anouk Van de Vel for your interest in my research. I hope we can collaborate in the future.

An important part of my PhD was the participation in several conferences, where I had the chance to meet incredibly nice people like Juan Bolea, Marianna Meo, Antonio Fasano, Valeria Villani, and Alessandro. I had so much fun during the social events, and it was always nice to know that you were going to be there and we were going to support each other during our talks. Special thanks to Luca Faes and Prof. Daniele Marinazzo for the nice collaboration we started not so long ago. I am sure we will come up with nice things in the future.

Thanks to all my colleagues in Biomed, Katrien, Ivan, Kris, Yipeng, Maarten, Rosy, Aileen, Bogdan, Kirsten, Wang, Wout, Anca, Tim, Diana, Amir, Vladimir, Miguel, Adrian, Nicolas, Bharath, Lieven, Devy, Milica, Laure, Vanya, Griet, Ninah, Rob, Bori, Nico, Otto, Steven, Maria Isabel, Thomas, and Alex. The party nights at Sabine's place were always really nice, and it has been a pleasure to work in such a diverse group. Gracias a Mauro y a Richi por los descansos, las conversas y por darme la oportunidad de compartir ideas y frustraciones. Muchas gracias a Carlosijn, por los sabios consejos, por el apoyo en momentos muy complicados, por los chistes que nunca llegaban a mala hora y los que todavía me hacen reír. Gracias Carlosijn por explicarme con paciencia y espero que podamos colaborar mucho más en el futuro.

Thanks to Hanski, Jan Deca, Jinbo, Mathieu, Daniel Rincon, Ivonna, Dusan, Wouter and Karolien for the dinners, parties, and social interaction that kept me away from becoming one with the computer. To Tim and Thomas thanks for the game nights, and Tim it was really nice to share the office with you. Bori and Adi it is so nice to share time with you two, but please do not forget to restart with the food competition. Thanks to Maria and Jorg for the fun, the matches, the song competition thing, and the nice atmosphere. Muchas

gracias a mi Patty Chu, Camiloco, Paulo, Vane, Hugatie y Nataly por estar siempre pendientes, por los mensajes que me llenaban de fuerza cuando más lo necesitaba. Los quiero mucho y es increíblemente bonito saber que cuento con ustedes siempre.

Nayiaaaa thanks for your patience and for not giving up despite the fact that in the last months I haven't been the best friend. I really appreciate your friendship and I am thankful to you for listening and supporting me here in Belgium. Dery de mi corazón, que sería de mí sin sus visitas. Después de tantos años me siento tan feliz de poder contar con usted y poder compartir tantas cosas bonitas junticas. Gracias por estar ahí mi Dery!

A doña Yolanda le quiero agradecer por todo el cariño que siempre me ha dado desde que nos conocimos. Usted ha sido como una mamá para mí. Muchas gracias por tantos cuidados, tantos mensajes bonitos, y por abrirme las puertas de su casa y su familia. Gracias a Manio por estar siempre pendiente, por los consejos y por tu amistad. Te agradezco siempre todas las cosas bonitas y el apoyo que siempre me has dado desde Colombia.

Ahora voy a agradecerle al Santi, al amigo, al enemigo, al compañero de casa, al cómplice, al compinche, etc., no sé cómo agradecerte por tantas cosas que has hecho por mí. Porque sé que a pesar de todas las bobadas siempre estas ahí y siempre estas pendiente de mí. Muchas gracias Santi por escucharme y por todo lo que hemos pasado juntos.

Ook een super bedankt aan Mia, Francois, Claudia, Koen en natuurlijk Quinten. Dank u voor jullie steun en voor de toffe weekends die we samen doorbrachten. Jullie zijn mijn familie hier in België, en dat betekent echt zeer veel voor mij. Steeds als ik bij jullie ben, voel ik me echt thuis.

Gracias a mi familia en Colombia, especialmente a mi tío Fernando por creer siempre en mí. Porque siempre ha estado pendiente de mí, por ser esa persona tan especial que me brinda seguridad desde que yo era pequeñita. A mi tati no sé cómo agradecerte tanto. Hemos pasado por tanto y acá estamos, fuertes y felices. A lilo, mi persona favorita, cada vez me siento más orgullosa de ti. Estoy tan feliz de compartir esto con ustedes, y de saber que siempre estamos los tres junticos en todo.

Y para rematar con broche de oro: special and the biggest thanks in the history of mankind go to Steven. Mi lindo you cannot imagine how happy and lucky I feel to share my life with you. Thanks for supporting me with every single idea, plan and experiment that I have. Thanks for always believing in me and for understanding what is going on in my head. Thanks for understanding that for some nights and weekends I could not spent time with you, in particular during the last few months. Thanks for your care and love, for our jokes,

systems, plans, trips, failures, dreams, and so on. I know that no matter how difficult and hard a day has been, I can always go back home and rest next to you.

Carolina
Leuven, April 2015

Abstract

This research focuses on the development of algorithms to extract diagnostic information from the ECG signal, which can be used to improve automatic detection systems and home monitoring solutions. In the first part of this work, a generically applicable algorithm for model selection in kernel principal component analysis is presented, which was inspired by the derivation of respiratory information from the ECG signal. This method not only solves a problem in biomedical signal processing, but more importantly offers a solution to a long-standing problem in the field of machine learning. Next, a methodology to quantify the level of contamination in a segment of ECG is proposed. This level is used to detect artifacts, and to improve the performance of different classifiers, by removing these artifacts from the training set. Furthermore, an evaluation of three different methodologies to compute the ECG-derived respiratory signal is performed. It is shown that for long-term signals with transients and non-stationarities, the R-peak amplitude offers the best solution taking into account computational complexity, and resemblance between the estimated and real respiratory signals. The next step of this work covers the quantification of the cardiorespiratory interactions by means of phase rectified signal averaging, information dynamics and subspace projections. These methodologies provide complementary information to the typical heart rate variability analysis, for the assessment of the autonomic control.

All these algorithms are applied in two main fields: sleep and epilepsy. In particular, the effect of sleep apnea on the ECG and respiration is analyzed, and based on this effect two new features are proposed and used to discriminate between normal activity and apnea episodes. These two features quantify the changes in the morphology of the ECG and assess the cardiorespiratory interactions during apnea events. A similar set of features is extracted and investigated for use in epileptic seizure detection. The ECG and the ECG-derived respiratory (EDR) signals are used to improve early detection of seizures and are promising for the development of accurate and robust closed-loop systems in epilepsy. In this respect, features based on the morphology of

the ECG, assessed by means of principal component analysis are proposed. In addition, in collaboration with pediatric neurologists Prof. Dr. K. Jansen and L. Lagae, ways to early detect seizures in childhood epilepsy were investigated, achieving accuracies of 93% and 85% for the detection of partial and generalized seizures, respectively. Features extracted from the ECG, such as those used in heart rate variability (HRV) analysis, together with the analysis of cardiorespiratory interactions reveal important autonomic dysfunctions in children suffering from West syndrome and absence epilepsy. These more fundamental findings may play an important role in understanding epilepsy, and in a long-term can change clinical practice and the associated treatment procedures.

Samenvatting

Dit onderzoek spitst zich toe op de ontwikkeling van algoritmes om diagnostische informatie uit ECG signalen te extraheren, en deze te gebruiken voor de verbetering van automatische detectiesystemen en thuismonitoring systemen. In het eerste deel van dit werk wordt een algemeen toepasbaar algoritme voor modelselectie in kernel principiële componentenanalyse beschreven, dat zijn oorsprong kent in de afleiding van ademhalingsinformatie uit het ECG signaal. Deze methode lost niet enkel een probleem op in biomedische signaalverwerking, maar zo mogelijk nog belangrijker, ook een langlopend en gekend probleem in het domein van machinaal leren. Vervolgens wordt een methodologie voorgesteld om de graad van ruis in een segment van een ECG signaal te kwantificeren. Dit wordt dan gebruikt om storingen in het signaal te detecteren en te verwijderen uit de trainingsgroep, en zo de prestatie van verschillende classificatiemethodes te verbeteren. Daarnaast wordt een evaluatie van drie verschillende methodologieën om de ademhaling uit het ECG signaal af te leiden uitgevoerd. Hierbij wordt aangetoond dat voor lange, niet-stationaire en transiënte signalen de R-top amplitude de beste oplossing geeft als men rekenkundige complexiteit en verschil tussen geschatte en effectieve ademhaling in rekening brengt. De volgende stap in dit onderzoek behandelt de kwantificatie van de interactie tussen hartslag en ademhaling door middel van fasegelijkrichting, informatiedynamica en deelruimte projecties. Deze methodologieën bevatten bijkomende informatie bovenop de typische analyse van hartslagvariabiliteit voor de beoordeling van autonome controle.

Al de bovenstaande algoritmes zijn toegepast in twee domeinen: slaap en epilepsie. In het bijzonder wordt het effect van slaapapneu op het ECG signaal en de ademhaling geanalyseerd. Op basis van dit effect, worden twee nieuwe functies voorgesteld en gebruikt om een onderscheid te maken tussen normale activiteit en apneu. Deze twee functies kwantificeren de veranderingen in de morfologie van het ECG signaal en de cardiorespiratoire interacties tijdens apneu. Een gelijkaardige set van functies wordt gedefinieerd en onderzocht voor gebruik bij detectie van epileptische aanvallen. Het ECG

signaal en de ademhaling afgeleid van het ECG signaal worden gebruikt om de vroegtijdige detectie van aanvallen te verbeteren, en dit is veelbelovend voor de ontwikkeling van accurate en robuuste gesloten lus systemen in epilepsie. In deze context worden functies voorgesteld die gebaseerd zijn op de morfologie van ECG signalen die op hun beurt door PCA afgeleid zijn. Daarnaast zijn, in samenwerking met pediatrische neurologen Prof. Dr. K. Jansen en L. Lagae, manieren om vroegtijdig aanvallen te detecteren bij kinderen onderzocht, waarbij een nauwkeurigheid van 93% en 85% werd bereikt bij de detectie van partiële en gegeneraliseerde aanvallen respectievelijk. Functies afgeleid van het ECG signaal, zoals degene die gebruikt worden in de analyse van hartslag variabiliteit, samen met de analyse van de cardiorespiratoire interacties, onthullen belangrijke autonome stoornissen bij kinderen die aan het West syndroom en afwezigheid epilepsie lijden. Deze fundamentele resultaten kunnen een belangrijke rol spelen in het beter begrijpen van epilepsie, en kan op langere termijn aanleiding geven tot een wijziging in de klinische praktijk en in de bijhorende behandelingen.

Nomenclature

Variables and Symbols

$x, s, \ldots$	Vectors
$\{x_i\}_{i=1}^{N}$, $x_i \in \mathbb{R}^d$	Set of N points
$A, B, \ldots$	Matrices
A_{ij}	ij-th entry of the matrix A
A^T	Transpose of the matrix A
$K(x_i, x_k)$	Kernel function evaluated on vectors x_i, and x_k
1_N	N-dimensional vector of ones
$\mathcal{F}$	Feature space
$\varphi(x)$	Feature map applied to x
$f(\cdot)$	Function
$\min_x f(x)$	Minimization over x. Minimal $f(x)$ is returned
$\operatorname{argmin}_x f(x)$	Minimization over x. Optimal x is returned
$\alpha_i^{(l)}$	i-th entry of the l-th eigenvector
$\lvert a \rvert$	Absolute value of a
σ^2	Bandwidth of an RBF kernel
ω	Contamination level of the ECG
μ_R	Mean of the RR interval time series

Abreviations

AC	Acceleration capacity
ACF	Autocorrelation function
AE	Absence epilepsy
AED	Antiepileptic drug
ANS	Autonomic nervous system
AP	Anchor point
APC	Atrial premature complex
AUC	Area under the receiver operating characteristic curve
AV node	Atrioventricular node
BPRSA	Bivariate phase rectified signal averaging
BLF	Baseline fit
CEN	Central apnea
CO	Control group
CRPS	Cardiorespiratory phase synchronization
CT	Computerized tomography
DC	Deceleration capacity
DBP	Diastolic blood pressure
ECG	Electrocardiogram
ECOC	Error-correcting output code
EDR	ECG-derived respiration
EEG	Electroencephalogram
EMG	Electromyography
FFT	Fast Fourier transform
FLE	Frontal lobe epilepsy
FN	False negative
FP	False positives
FS	Fixed-size method
FS2	Modified fixed-size method
GAB	Generalized absence
GTN	Generalized tonic
HADS	Hospital Anxiety Depression Self-Assessment
HB	Heart beat
HF	Heart failure
HIP HPA	Hypopnea
HR	Heart rate
HRV	Heart rate variability
ICA	Independent component analysis
ICF	Incomplete Cholesky factorization
ILAE	International league against epilepsy
IS	Infantile spasms

KKT	Karush-Kuhn-Tucker
KPA	Kernel parallel analysis
KPCA	Kernel principal component analysis
KPSS	Kwiatkowski-Phillips-Schmidt-Shin test
KSC	Kernel spectral clustering
LA	Left atrium
LDA	Linear discriminant analysis
LS-SVM	Least-squares support vector machine
LV	Left ventricle
LVSD	Left ventricular systolic dysfunction
MDD	Model selection based on distance distribution
MDL	Minimum description length
MIX	Mixed apnea
MRI	Magnetic resonance imaging
MSE	Mean squared error
NYHA	New York Heart Association
OSA	Obstructive sleep apnea
OSH	Obstructive hypopnea
PCA	Principal component analysis
PPV	Positive predictive value
PRSA	Phase rectified signal averaging
PSD	Power spectral density
PSG	Polysomnography
QoL	Quality of life
RA	Right atrium
RBF	Radial basis function
Recap	Regional care portals
REM	Rapid eye movement
RMS	Root mean square
RR	RR intervals
RSA	Respiratory sinus arrhythmia
RV	Right ventricle
SA node	Sinoatrial node
SBP	Systolic blood pressure
SDNN	Standard deviation of normal to normal beats
SNR	Signal-to-noise ratio
SUDEP	Sudden unexplained dead in epilepsy
SV	Support vector
SVM	Support vector machine
TDS	Time delay stability
TLE	Temporal lobe epilepsy
TN	True negative
TP	True positive
VNS	Vagus nerve stimulation
VPC	Ventricular premature complex

Contents

Acknowledgments iii

Abstract vii

Samenvatting ix

Nomenclature xi

Contents xv

List of Figures xxiii

List of Tables xxix

1 Introduction 1

1.1 Research motivations 1

1.2 Research objectives 3

1.3 Chapter overview and personal contributions 7

1.3.1 Part I. Algorithms 7

1.3.2 Part II. Applications 9

1.4 The electrocardiogram 11

1.4.1 Physiological aspects 11

1.4.2 Characteristic waves of the electrocardiogram 13
1.4.3 Abnormal heart beats - Ectopic beats 13
1.4.4 Mechanical effect of respiration on the ECG 14
1.5 The autonomic nervous system 15
1.5.1 Heart rate control 16
1.5.2 Respiratory control 17
1.6 Cardiorespiratory interactions 17
1.6.1 Respiratory sinus arrhythmia 18
1.6.2 Cardiorespiratory phase synchronization 18
1.6.3 Time-delay stability 19
1.7 Ambulatory monitoring 19

I Algorithms **21**

2 Machine Learning Techniques **23**

2.1 Supervised learning techniques for classification 24
2.1.1 Linear discriminant analysis 24
2.1.2 Support vector machines 25
2.1.3 Least-squares support vector machines 27
2.2 Unsupervised learning techniques 28
2.2.1 Principal component analysis 29
2.2.2 Kernel principal component analysis 29
2.2.3 Kernel spectral clustering 32
2.3 Large scale problems 37
2.3.1 Incomplete Cholesky factorization 37
2.3.2 Fixed-size LS-SVM 39
2.4 Conclusions 41

3 Model Selection for KPCA **43**

3.1 Noise and dimensionality 44

3.2 MDD: model selection based on distance distributions 47

3.2.1 Computation of the distance matrix 48

3.2.2 Maximization of the information content 49

3.3 Large scale KPCA 54

3.4 Experimental Results 55

3.4.1 Toy examples 57

3.4.2 Denoising of digits 62

3.5 Conclusions 67

4 ECG pre-processing and Analysis **69**

4.1 ECG Artifacts 70

4.1.1 Sources of noise and artifacts 70

4.1.2 Noise level estimation 71

4.1.3 Artifact detection algorithms 78

4.2 QRS detection and the construction of the tachogram 83

4.2.1 Flattening the ECG 84

4.2.2 The Pan-Tompkins algorithm 84

4.2.3 Construction of the tachogram 86

4.2.4 Test on the MIT/BIH arrhythmia Physionet data 87

4.3 Heart rate variability 90

4.3.1 Time domain analysis 90

4.3.2 Frequency domain analysis 92

4.4 ECG-Derived Respiration 93

4.4.1 R-peak amplitude 93

4.4.2 Principal component analysis 94

4.4.3 kernel principal component analysis 94
4.4.4 Evaluation of the EDR algorithms 95
4.5 Conclusions 99

5 Quantifying cardiorespiratory interactions 101
5.1 Phase rectified signal averaging 102
5.1.1 Univariate PRSA 102
5.1.2 Bivariate PRSA 105
5.1.3 Quantification and interpretation 105
5.2 Information dynamics 108
5.2.1 Entropy decomposition 109
5.2.2 Quantification of entropy measures 111
5.2.3 Statistical significance 112
5.2.4 Interpretation on the cardiorespiratory interactions . . . 112
5.2.5 Why use information dynamics? 113
5.3 Subspace projections 113
5.3.1 Formulation 114
5.3.2 Subspace definitions 114
5.3.3 Quantification and interpretation 118
5.4 Conclusions 120

II Applications 121

6 Cardiorespiratory interactions in epilepsy 123
6.1 Introduction to epilepsy 124
6.1.1 Definitions 124
6.1.2 Diagnosis and treatment 125

6.1.3 Seizure types 125
6.2 Cardiorespiratory interactions in West syndrome 128
6.2.1 Data 129
6.2.2 Methods 130
6.2.3 Results in West syndrome 131
6.2.4 Discussion 136
6.3 Cardiorespiratory Variability in TLE and Absence Epilepsy . . 137
6.3.1 Data 139
6.3.2 Data processing 139
6.3.3 Statistical analysis 140
6.3.4 Results in TLE and AE 140
6.3.5 Discussion 143
6.4 Cardiorespiratory interactions towards seizure onset 145
6.4.1 Data 146
6.4.2 Pre-processing 147
6.4.3 Dynamic computation of cardiorespiratory interactions . 147
6.4.4 Results and discussion 147
6.5 Conclusions 152

7 Detection of epileptic seizures from single-lead ECG 153
7.1 Introduction to seizure detection 154
7.2 Datasets 155
7.3 Methodology 156
7.3.1 Preprocessing and R-peak detection 156
7.3.2 Continuous quantification of ECG morphology changes 157
7.3.3 Continuous measurement of cardiorespiratory interactions 157
7.3.4 Seizure detection 159

7.4 Results and discussion 161
7.4.1 Detection using morphology information 161
7.4.2 Detection using (B)PRSA 166
7.4.3 Unifying seizure detection algorithm 171
7.5 Conclusions 171

8 Sleep apnea detection from single-lead ECG 173
8.1 Introduction to sleep apnea detection 174
8.2 Methods 175
8.2.1 Datasets 175
8.2.2 ECG pre-processing 177
8.2.3 ECG contamination level 177
8.2.4 ECG-derived respiration 177
8.2.5 Feature based on the QRS morphology 178
8.2.6 Feature derived from subspace projections 179
8.2.7 Feature selection 183
8.2.8 Classification 183
8.3 Results 185
8.3.1 Comparison between respiratory signals: real and EDR 185
8.3.2 Feature selection 185
8.3.3 Apnea classification 188
8.4 Discussion 190
8.5 Conclusions 193

9 Conclusions and Future directions 195
9.1 Conclusions 195
9.1.1 Algorithms 196
9.1.2 Applications 198

9.2 Future directions 201
9.2.1 Algorithms 201
9.2.2 Applications 202

A Telehealth in cardiology **205**
A.1 Introduction 205
A.2 Methods 207
A.3 Results 211
A.4 Discussion 217
A.5 Conclusions 218

B Graphical user interfaces **221**

Bibliography **225**

Curriculum vitae **247**

List of publications **249**

List of Figures

1.1 Simplified diagram for a monitoring system based solely on the ECG 4

1.2 Schematic chapter overview 8

1.3 Schematic representation of the human heart 12

1.4 Electrode placement for the standard 12-lead ECG recordings . 12

1.5 Characteristic waves and segments of a normal ECG 13

1.6 Example of an isolated ventricular ectopic beat 14

1.7 Example of a simultaneous ECG and respiratory effort measured in the abdomen using a respiratory belt 15

2.1 Dataset and cluster identifiers obtained using kernel spectral clustering 36

2.2 Clustering results for different values of the kernel parameter σ^2 36

2.3 Results using the regular and the modified fixed-size algorithm . 41

3.1 Noise estimation in 2 and 50 dimensions after permuting the input data 45

3.2 Quadratic Renyi entropy of a structured dataset contained in different dimensions 46

3.3 Density functions of distances for a dataset of uniformly distributed points in spaces of different dimensions 47

3.4 Toy datasets with different signal to noise ratio 52

3.5 Distance distributions and eigenvalue spectra of datasets with different SNR 52

3.6 Real and estimated distance distributions and eigenvalue spectra 53

3.7 Model selection using MDD 54

3.8 Implementation of large scale KPCA with the proposed model selection algorithm MDD. 56

3.9 Model selection criteria for different values of σ^2 57

3.10 Denoised datasets after model selection 59

3.11 Eigenvalue spectra for the square example using kPA and MDD 60

3.12 Mean computation time on toy examples 61

3.13 Model parameters and MSE for different ratios of training and validation points 62

3.14 Grid search for the model parameters 64

3.15 Denoising results on the UCI dataset 65

3.16 Eigenvalue distributions and information for different noise estimation approaches 66

3.17 MSE for different signal to noise ratios of the UCI and USPS datasets 67

4.1 ECG segments contaminated by artifacts 72

4.2 ECG segments and their corresponding power spectra 74

4.3 Autocorrelation functions of four different ECG segments of one minute 75

4.4 Autocorrelation functions and degrees of a long term ECG signal segmented on a minute-by-minute basis 76

4.5 ECG segments of 10s and eigenvalue distributions) of the data and noise 77

4.6 One minute ECG segmented into epochs of 5 seconds 81

4.7 Performance of Algorithm 4 applied to one ECG signal and for different window lengths 81

4.8 Weights and information content for all the 9132 ECG segments 82

4.9 Computation of upper ECG envelope using the secant method 85

4.10 Computation of the flattened ECG signal 85

4.11 Computation of the EDR signals using the R-peak amplitude and PCA 95

4.12 Computation of the mean magnitude squared coherence 97

4.13 Boxplots of correlation and mean magnitude squared coherence between the reconstructed and the original respiratory signals . 98

5.1 Phase rectified signal averaging (PRSA) 103

5.2 Methodology to perform the bivariate version of PRSA 106

5.3 Quantification of the PRSA curves 107

5.4 Entropy decomposition using information dynamics 110

5.5 Decoupling of heart rate and respiration using orthogonal subspace projections 115

5.6 Decoupling of the RR interval time series from the respiratory signal using orthogonal subspace projections 119

6.1 Comparison between a normal EEG activity and a chaotic EEG behavior with high-amplitude waves characteristic in West syndrome. 129

6.2 Time domain HRV parameters with significant differences ($p < 0.05$). 132

6.3 Frequency domain HRV parameters 132

6.4 PRSA and BPRSA curves in West syndrome 134

6.5 Parameters derived from the (B)PRSA curves 135

6.6 Power ratios of the heart rate components in West syndrome . 135

6.7 EEG of a typical absence seizure 138

6.8 Parameters derived from the RR interval time series in temporal lobe and absence epilepsy 141

6.9 Mean and standard deviation of the respiratory rate computed from the EDR signals 142

6.10 Boxplots of the entropy terms for different types of epilepsy . . 144

6.11 Cardiorespiratory signals during an absence seizure 148

6.12 Information dynamics analysis around a generalized seizure . . 149

6.13 Self entropy and self conditional entropy around the seizure onset, and during seizure-free activity 150

6.14 Cross entropy and transfer entropy around absence seizures and generalized tonic/tonic-clonic seizures, compared with values for seizure-free regions 151

7.1 Pattern of heart rate changes caused by a temporal lobe seizure 155

7.2 Procedure to construct the *QRS*-matrix, and derive its non-zero eigenvalues 158

7.3 Computation of the (B)PRSA curves 160

7.4 Signals derived from the ECG after R-peak detection 162

7.5 Threshold for λ_5 163

7.6 Labels for each segment of 5 seconds after classification and clustering 164

7.7 General procedure to detect epileptic seizures 165

7.8 (B)PRSA curves 167

7.9 Features extracted from the (B)PRSA curves of a focal seizure and a generalized seizure 169

8.1 Eigenvalues of the covariance matrix of $\mathbf{X}^T$ 179

8.2 Heart rate components derived using orthogonal subspace projections 182

8.3 Selection of the threshold for the contamination level ω 184

8.4 Features extracted using orthogonal subspace projections . . . 186

8.5 Examples of the heart rate components derived using R_{amp} . . 187

8.6 Features selected as the most discriminative ones 189

8.7 Real and estimated apnea-hypopnea indices 190

9.1 Diagram of an ECG monitoring system that allows to extract cardiac and respiratory information 196

A.1 Heart failure patients on the telehealth program. 210

A.2 Amount of data transmitted to the hospitals for each telehealth system 212

A.3 Physiological parameters at day 1 and age of the patients in both groups under investigation 213

A.4 Slopes for each physiological parameter and for each telehealth system 213

A.5 Scores of the self-care behaviour questionnaire 215

A.6 Scores of items 1 and 6 216

B.1 Graphical user interface for the computation of the contamination level of segments of the ECG. 222

B.2 Graphical user interface for the manual inspection of R-peaks. . 223

B.3 Graphical user interface for the computation of the EDR and HRV parameters. 224

List of Tables

3.1 Comparison of different model selection criteria and two large-scale techniques, on toy examples 58

3.2 Comparison of different model selection criteria and two large-scale techniques on the UCI dataset 63

3.3 Size of the training set, model parameters, and MSE on denoised test data and original USPS patterns 64

4.1 Comparison of the performance of the proposed algorithms . . 80

4.2 Performance of the QRS detection algorithm on the Physionet MIT/BIH arrhythmia dataset 88

4.3 Indices of heart rate variability 91

6.1 Characteristics of West patients 130

7.1 Set of features derived from the ECG, RR interval time series, and respiratory signal R corresponding to any of the EDR signals 161

7.2 Seizure detection results using morphology changes in the ECG. 166

7.3 Seizure detection results using (B)PRSA. 170

8.1 Characteristics of patients in the Leuven dataset 176

8.2 Set of features in apnea detection 181

8.3 Performances of the tested classifier 191

A.1 Main characteristics of the telehealth systems 209

Chapter 1

Introduction

1.1 Research motivations

The electrocardiogram (ECG) is a very well-known diagnostic tool and it is among the most preferred tests in clinical practice [152]. A reason for this is that it is a relatively easy-to-record signal, and it contains an enormous amount of relevant information that can be used for the assessment of cardiac health. In addition, the costs associated with such a test are relatively low when compared to other medical tools like those involving imaging. In fact, the ECG is probably one of the few medical tools available to people in several developing areas. For all these reasons, a lot of attention has been given to the development of ambulatory systems based on ECG. This is important not only to supply remote and/or developing areas with basic health care, but also to diagnose and monitor different diseases in a home environment.

Even though several studies in the literature have been dedicated to the analysis of the ECG signal, there are still many challenges that need to be tackled before fully relying on an ECG monitoring system [51]. Progress on these challenges has become possible thanks to the gradual improvement of the instrumentation and computational resources. For example, nowadays it is possible to use portable ECG monitors to collect data with a relatively high accuracy, but many improvements can still be made on the algorithms to process and analyze such data. One of these improvements is related to the way portable ECG monitors handle data in the presence of artifacts caused by movement or degradation of the electrodes. In this respect, this research work aims to provide a scoring system that allows to label different segments

of the signal depending on their contamination level. These scores can be used to identify segments of interest, in other words, segments that will need visual inspection. By doing so, it is possible to save a considerable amount of time spent to visual screening of long-term ECG recordings by doctors.

Apart from proposing a methodology to deal with ECG artifacts, this research also aims at maximizing the amount of information that can be derived from the ECG signal. That is to say, whereas heart rate information can be derived from the time differences between consecutive R-peaks, much more information is contained in the morphology of the characteristic waves of the ECG. Hence, this research proposes a new feature that captures changes in the morphology of the QRS complexes and that can be implemented in an online monitoring system. Changes in this morphology are well known to be caused by the mechanical interaction between the respiration and the ECG [143], and they have been used to derive information about the respiratory pattern [54, 119, 217]. Therefore, it is important to improve and evaluate these methodologies for their application in ambulatory systems. This poses new challenges since the presence of transients and non stationarities in the signals, have until today not been sufficiently tackled by any of those methodologies. If an accurate respiratory signal can be estimated under these circumstances, it is possible to analyze not only cardiac, but also respiratory activity from only one signal. This represents a huge reduction in the amount of sensors that are needed for the correct assessment of the cardiorespiratory system.

The ultimate goal in ambulatory ECG analysis, to which this research mainly contributes, is the development of a system with the following properties:

- few sensors attached to the body
- low cost
- accurate information about the cardiorespiratory system.

The development of such a system can be advantageous for the diagnosis, monitoring and follow-up of many disorders such as epilepsy and sleep apnea. Epilepsy for example, is a neurological disorder that affects around 50 million people worldwide, and according to the World Health Organization, it accounts for 0.5% of the global burden of disease. Around 80% of these 50 million people are found in developing regions, where diagnostic tools are very limited, and where no treatment or monitoring strategies are available. In addition, the occurrence of seizures is unprovoked and can be very sporadic in some cases, thus a continuous monitoring is necessary to provide an adequate diagnosis. In these cases, the standard video-EEG monitoring is used, but it requires that the patient wears many sensors attached to the scalp, which can be very obtrusive

and interfere with their regular activities. For this reason, a huge improvement can be achieved if a system based on fewer and less obtrusive sensors can be used. In this particular application, a system based on ECG can be useful due to the serious effect that epilepsy has on the cardiac and respiratory control mechanisms [33, 59, 135, 149].

The second field where ambulatory ECG systems can be used is the detection of sleep apnea. Sleep apnea is the most common sleep-related breathing disorder, characterized by disruptions of normal breathing patterns during sleep, and it is recognized as a major risk factor for mortality due to its effect on the cardiovascular system [147]. The prevalence of sleep apnea is around 4% in men and 2% in women, and it is known that at least 75% of cases of severe sleep-related breathing disorders remain undiagnosed [221]. Currently, these disorders are diagnosed using polysomnography (PSG), which is a sleep test that monitors different physiological signals such as heart rate, respiration, EEG, muscle tone and eye movement. This type of test needs to be performed in a sleep laboratory, under the supervision of an expert who attaches several electrodes to the patient. Although PSG is an important diagnostic tool for sleep medicine, it is an uncomfortable and costly procedure, especially when multiple nights of observation are required. Therefore, it is clear that the use of a less obtrusive and more comfortable system will also improve the diagnosis and therapy of sleep-related breathing disorders.

This research will study the applicability of both novel and existing algorithms for the analysis of ECG in these two aforementioned applications. Information about morphology, cardiac and respiratory activity will be derived, and changes in their patterns caused by epilepsy and sleep apnea episodes will be investigated.

1.2 Research objectives

The principal goals of this research are twofold. On the one hand, it aims to develop algorithms for the extraction of informative features from the ECG that can be used for the quantification of cardiac and respiratory activities. On the other hand, it evaluates the application and interpretation of those informative features in sleep and epilepsy research. To achieve these two goals, the whole track can be divided into three main blocks as indicated in Figure 1.1. The first block deals with the assessment of the contamination level in the ECG signal. This is crucial, since the diagnostic power of the ECG can be compromised by different contaminants such as electrode motion and power line interference. The second part of the track comprises the development of novel features

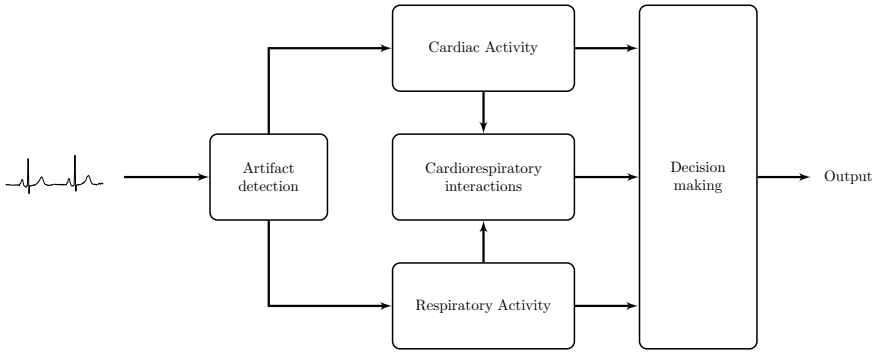

Figure 1.1: *Simplified diagram for a monitoring system based solely on ECG.*

that are easy to interpret and easy to compute solely from the ECG signal. Not only cardiac information is extracted here but also information related to respiration. In the last block, the algorithms and features derived before are used to study cardiorespiratory interactions in applications, in particular in epilepsy and sleep. In general, different challenges (Chl.) can be identified, both in the development of algorithms and in their applicability, as listed below:

Algorithms

Chl 1. **Development of a methodology to handle large-scale applications.** It is well-known that the continuous, long-term monitoring of ECG signals involves the computation of different parameters at several points in time. However, limitations in computational resources demand the development of algorithms that allow selecting those points in time that are most representative for the problem under investigation. In other words, those points that best describe the underlying behavior of the parameters must be selected and used to build machine learning models, such as classifiers to discriminate normal from abnormal ECG activity. This challenge will be discussed in Chapter 2.

Chl 2. **Development of a model selection approach for kernel principal component analysis.** Different algorithms have been proposed to extract respiratory information from the ECG signal. One of these algorithms is based on kernel principal component analysis (KPCA) [217], which takes into account nonlinear interactions between respiration and the morphology of the ECG. However, this methodology requires the optimization of model parameters in order to produce results that are comparable, or that can even improve the results provided by other

algorithms. This is a very well-known challenge in machine learning, and its solution can impact not only the accurate computation of the ECG-derived respiratory (EDR) signal, but also many other application fields. In this context, Chapter 3 tackles this challenge and aims to propose a model selection methodology for KPCA.

Chl 3. **Development of a scoring system that allows to differentiate contaminated ECG epochs from "clean" segments**. A level of contamination can be assigned to each ECG segment, and it can be used to determine the confidence of the results obtained for each particular segment. For instance, for online detection systems this can be useful to determine the reliability of the ECG signal at each moment in time. For support systems that are based on offline detection of medically relevant events, this scoring can be used to label the segments that require further visual interpretation. Such a scoring system will be discussed in Chapter 4.

Chl 4. **Quantification of morphological changes of the ECG signal.** Currently, many studies on autonomic function use the ECG signal to compute the tachogram, which is a time series that contains information about instantaneous heart rate. Although this time series is an essential tool to investigate autonomic control, a large amount of information contained within the ECG itself is often ignored. This is the case for morphological changes of the characteristic ECG waves such as the QRS complexes, which also seem to be affected by many disorders. For this reason, one of the challenges of this research is to capture this morphological information during continuous monitoring. In Chapter 4, a novel feature will be proposed, which will quantify the morphological changes of the QRS complexes.

Chl 5. **Evaluation of different EDR algorithms on real and continuous datasets.** This evaluation is typically done using nearly stationary segments. However, in order to determine if these algorithms can be used in real life applications, transients and different dynamics will be considered and investigated in Chapter 4.

Chl 6. **Quantification of cardiorespiratory interactions using only the ECG signal.** As can be seen in the previous items, the respiratory activity can be estimated from the ECG signal. Therefore, different ways to quantify the interactions between this estimated signal and the heart rate will be discussed in Chapter 5. Furthermore, different EDR algorithms will be evaluated for this purpose.

Applications

Chl 7. **Quantification of cardiorespiratory interactions in epilepsy.** Many epilepsy studies have focused on the analysis of the heart rate during epileptic seizures. However, little is known about the effect that epilepsy has on respiration, and on the interactions between respiration and heart rate. Hence, in order to take into account these effects, and to envision a seizure detection system, the algorithms developed in this research will be analyzed and evaluated against the goal standard in epilepsy research, namely video-EEG monitoring. By doing so, the applicability of the proposed algorithms on epilepsy research will be assessed. As a result, not only detection systems will be considered, but also unknown effects of epilepsy will be unveiled. This can of course, in a long term, improve clinical practice, and influence the development of new treatment and diagnosis strategies in epilepsy. This challenge will be covered in Chapter 6.

Chl 8. **Development of seizure detection algorithms based on single-lead ECG.** Current state-of-the-art seizure detection algorithms, based on ECG are only used for the detection of partial seizures originating from the temporal lobe. The reason for that relies on the fact that generalized seizures are typically more difficult to detect from the heart rate signal. In Chapter 7, this challenge is tackled by including respiratory information, hence, an algorithm that is capable of detecting partial and generalized seizures will be proposed.

Chl 9. **Development of an algorithm for sleep apnea detection from single-lead ECG.** In line with the last two approaches, the proposed methodologies will also be evaluated for the detection of sleep apnea in Chapter 8. This is a typical problem, to which algorithms based on single-lead ECG are applied. In addition, state-of-the-art apnea detection algorithms often use EDR signals. In this research, information about the morphology of the ECG and the interactions between respiration and heart rate will also be included, in order to improve the detection of apneic events. The proposed features and signals will be investigated in order to quantify their changes induced by apnea and hypopnea events. With this in mind, a novel algorithm for the detection of sleep apnea from single-lead ECG signals will be proposed.

1.3 Chapter overview and personal contributions

This book is divided in two main parts. Part I deals with the different algorithms proposed in this research for the analysis of the ECG signal, and part II describes the application and implementation of these algorithms on two main problems in medical diagnosis: epilepsy and sleep. Each part is subdivided into different chapters that will be outlined below. A schematic overview of these chapters is presented in Figure 1.2, where the color of the chapters indicates the level of the personal contributions. In other words, the darker the color, the larger the contribution of this research to the content of each chapter.

1.3.1 Part I. Algorithms

Chapter 2

This chapter describes the different machine learning techniques used in this research. It focuses on the different problem formulations for supervised and unsupervised methods. Supervised learning techniques include classification using linear discriminant analysis, support vector machines (SVM), and least-squares support vector machines (LS-SVM). The unsupervised methods described in this chapter are principal component analysis (PCA), kernel PCA, and kernel spectral clustering. In addition, this chapter tackles Chl 1., and it presents an original contribution of this work, namely a modified fixed-size algorithm to handle large-scale datasets, published in [204].

Chapter 3

The model selection approach for kernel principal component analysis corresponding to Chl 2. is presented in this chapter, and is one of the main contributions of this research. It is based on noise level estimation using pair-wise distance distributions, and it proves to outperform the state-of-the-art model selection algorithms for denoising. This original contribution is published in [204].

Chapter 4

This chapter discusses the different steps that need to be taken to perform ECG pre-processing and analysis in general. It describes some ECG-based

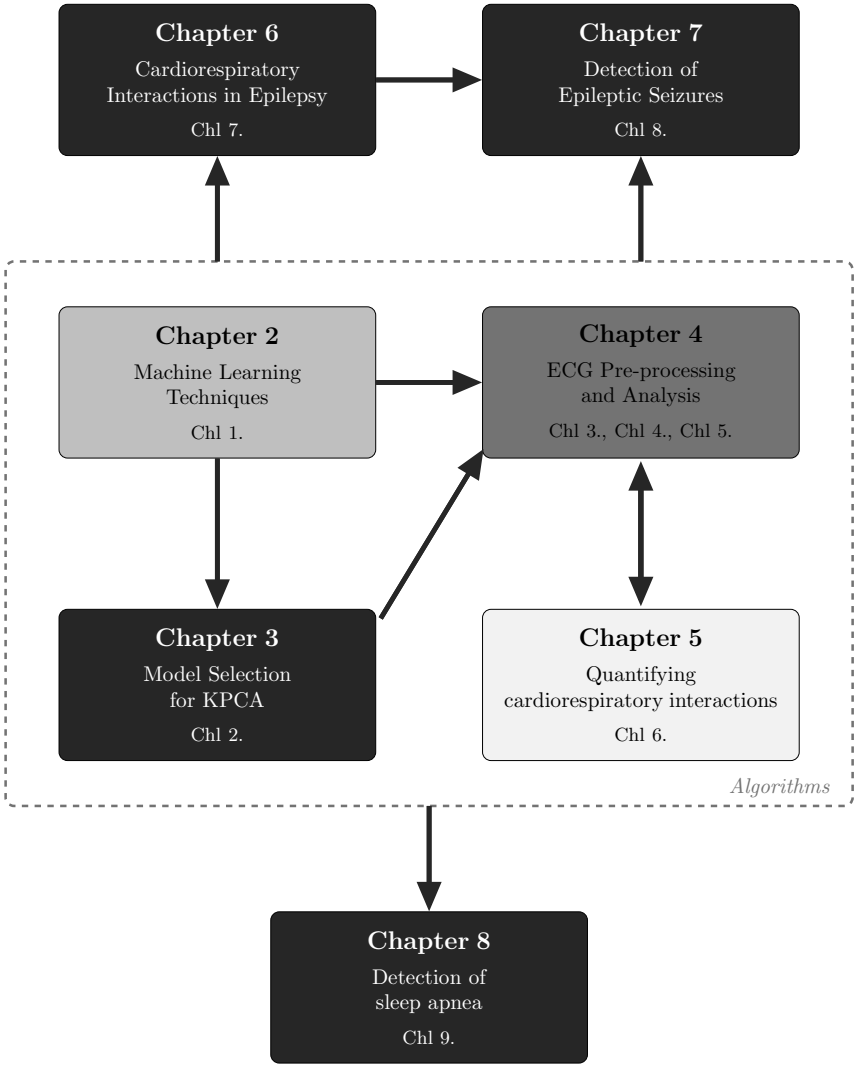

Figure 1.2: *Schematic chapter overview. The darker the color, the larger the personal contribution to each particular chapter. Chapters 2 to 5 make up the algorithmic part of this work, and they are all applied to the different diagnostic problems in epilepsy and sleep described in chapters 6 to 8.*

methodologies that use some of the techniques presented in Chapters 2 and 3. In particular, it discusses two novel approaches to detect artifacts from the ECG signal (Chl 3.). One approach is based on the work presented in [213], where the cosine similarity between the autocorrelation function of the ECG segments was used to discriminate between "clean" and "contaminated" ECG segments. The other approach is based on the noise level estimation approach discussed in Chapter 3. These two techniques are evaluated on a dataset collected in the University Hospital Leuven, UZ Leuven. Furthermore, a modified algorithm for the detection of the QRS complexes from the ECG is proposed and evaluated on a public dataset. Part of this work is published in [206].

The last part of this chapter covers the derivation of respiratory information from the ECG signal, or the so-called ECG-derived respiration (EDR). It illustrates the methodology based on principal component analysis (PCA) to derive morphological information from the QRS complexes (Chl 4.). Furthermore, it presents the results of the comparison between three different EDR algorithms (Chl 5.), namely the R-peak amplitude, PCA, and kernel PCA, applied on a public dataset. Part of this work is published in [217].

Chapter 5

This chapter presents three different methodologies that are used to quantify the interactions between heart rate and respiration using only the ECG signal (Chl 6.). These methodologies are phase rectified signal averaging, information dynamics, and subspace projections. The latter is published in[41].

1.3.2 Part II. Applications

Chapter 6

This chapter presents a large contribution of this research. It describes how the cardiorespiratory interactions are affected in children suffering from different types of epilepsy (Chl 7.). In particular, West syndrome, absence epilepsy, and temporal lobe epilepsy are discussed. Moreover, it shows how the algorithms discussed in the first part of this book can be applied on ECG signals recorded from epileptic patients, and which information can be extracted and related to the pathophysiology of epilepsy. The work presented in this chapter is published in [211, 101, 100, 207, 210, 208].

Chapter 7

This chapter builds up on the findings discussed in Chapter 6, and it proposes two different algorithms to detect epileptic seizures from single-lead ECG signals (Chl 8.). The first algorithm is based on PCA, and it quantifies changes in the morphology of the ECG caused by epileptic seizures. The second algorithm on the other hand, quantifies the behavior of the cardiorespiratory interactions before and during an epileptic seizure, by means of phase rectified signal averaging (PRSA). These algorithms achieve positive predictive values (PPV) larger than 80% for both partial and generalized seizures. This work is published in [205, 209, 102].

Chapter 8

A different application of the algorithms presented in the first part of this book is sleep apnea detection. In this chapter, a methodology to automatically detect sleep apnea is presented (Chl 9.), and it combines many of the algorithms discussed before. These algorithms correspond to the scoring approach based on the autocorrelation function; the characterization of morphological changes by means of PCA; the derivation of the respiratory signal by means of the R-peak amplitude, PCA, and KPCA; and the quantification of cardiorespiratory interactions using subspace projections. This proposed methodology achieves accuracies of about 84% for two different datasets. This work is presented in [206, 213, 214].

The remainder of this chapter will introduce some important concepts that are needed to tackle the different challenges of this research. Section 1.4 will discuss some physiological aspects of the electrocardiographic signal, its characteristic waves, abnormalities, and the mechanical interaction between respiration and its morphology. Next, Section 1.5 will cover the autonomic nervous system, in particular the mechanisms responsible for cardiac and respiratory regulation. After that, three different forms of cardiorespiratory interactions will be described in Section 1.6, and finally some practicalities that need to be considered in ambulatory monitoring will be presented in Section1.7.

1.4 The electrocardiogram

1.4.1 Physiological aspects

The heart consists of two main pumps illustrated in Figure 1.3, that are responsible for providing enough pressure to pump oxygenated blood through the entire body. Each pump consists of two chambers: one atrium and one ventricle. The right atrium receives the deoxygenated blood from the body and it pumps it to the ventricle, which then sends it to the lungs to be oxygenated. This blood is then sent back to the heart through the pulmonary veins, which arrive at the left atrium. The latter is then in charge of moving the blood to the left ventricle, which is then responsible to pump the blood through the body. As can be seen, the heart goes through two main phases that make up one *cardiac cycle*. One is the filling phase or so-called *diastole*, and one is when the heart contracts to pump the blood to either the lungs or to the entire body, called *systole*. In a normal heart, this cardiac cycle starts with the electrical stimulation of the sinoatrial (SA) node (see Figure 1.3). This electrical impulse is then conducted to the atrioventricular (AV) node, and later (about 100ms) it is transmitted to the ventricles. After that, the heart takes a refractory period of at least 200ms to go back to its resting electrical potential, and then it is ready to start the cycle again [84]. The rate at which this cycle is repeated (i.e. heart rate) is determined by the SA, which is seen as the pacemaker of the heart. Furthermore, the SA node consists of cells that are capable of self-exitation, hence, it has its intrinsic firing rate. However, this rate is constantly modified by the central nervous system, and the way this is done will be discussed in section 1.5.

The electrical impulse generated in the SA can travel through the heart due to the electrical properties of the cardiac cells, and this propagation can be measured as potential differences on the surface of the skin. This way of measuring the cardiac activity is known as the *electrocardiogram* or ECG, and its characteristic waves will be described in the next subsection. First, the position of the skin electrodes with respect to the propagation of the cardiac impulse, or so-called *cardiac vector*, determines the type of information that can be retrieved. The reason for this relies on the fact that the magnitude and direction of the propagation of the electrical impulse, i.e. the cardiac vector, changes and rotates over time. Therefore, the information retrieved from one pair of electrodes (i.e. lead) is related to the vector changes with respect to one particular axis. In clinical practice, this vector is projected onto many different axes, by means of the standard 12-lead ECG configuration. This is done in order to completely characterize the electrical activity of the heart. The 12-lead configuration consists of 10 electrodes, four positioned in the limbs and

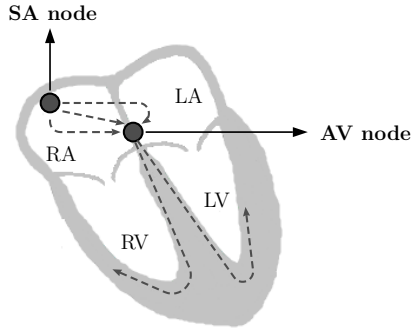

Figure 1.3: *Schematic representation of the human heart. RA, right atrium; RV, right ventricle; LA, left atrium; LV, left ventricle. The propagation of the electrical stimuli is indicated by the dashed line going from the sinoatrial (SA) node to the atrioventricular (AV) node, and finally to the ventricular tissues.*

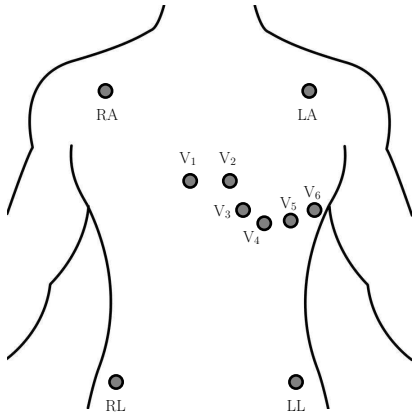

Figure 1.4: *Electrode placement for the standard 12-lead ECG recordings.*

six on the chest as indicated in Figure 1.4. From these electrodes, 12-leads are obtained: 6 are measured between V1-6 and the Wilson's central terminal, lead I = $LA - RA$, lead II = $LL - RA$, lead III = $LL - LA$, and three augmented leads (aVR, aVL, and aVF). In this research, only lead II was used. The reason for this is that this lead together with lead V5, are normally the most informative leads for medical diagnosis [51]. This is why lead II is typically used in ambulatory systems.

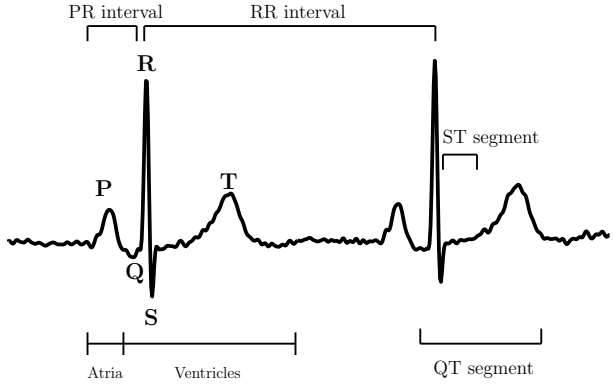

Figure 1.5: *Characteristic waves and segments of a normal electrocardiogram. Two heart cycles are indicated.*

1.4.2 Characteristic waves of the electrocardiogram

Figure 1.5 shows a segment of a normal ECG, which visualizes the combined effect of repolarization and depolarization of the different regions of the heart. As mentioned before, one cardiac cycle, or heart beat, starts when the electrical impulse is fired in the SA node. Next, this impulse propagates through the atria causing them to contract, and depolarize. This effect can be observed from the low-amplitude *P-wave*. Its amplitude can range between 0.1mV to 0.2mV [162]. Once the propagation reaches the AV node, it is delayed for less than 100ms, which results in an iso-electric segment, called the *PQ or PR segment*. After that, the impulse propagates to the ventricles causing contraction, which can be observed from the dominant *QRS complex*. The depolarization of the ventricles might last for up to 350ms, which causes an iso-electric segment after the QRS, called the *ST segment*. Finally, the ventricles relax and this can be seen as the repolarization wave, or *T-wave*. As can be seen, the ECG contains information about depolarization and repolarization of the different chambers of the heart. This makes the ECG a very important tool to detect abnormalities in the propagation pattern, such as arrhythmias. Arrhythmias can be caused by an irregular stimulation of the SA node, or by stimulation originated from other regions of the heart.

1.4.3 Abnormal heart beats - Ectopic beats

As mentioned before, the SA node consists of cells that have self-excitatory properties. Nevertheless, this is not the only region of the heart with the

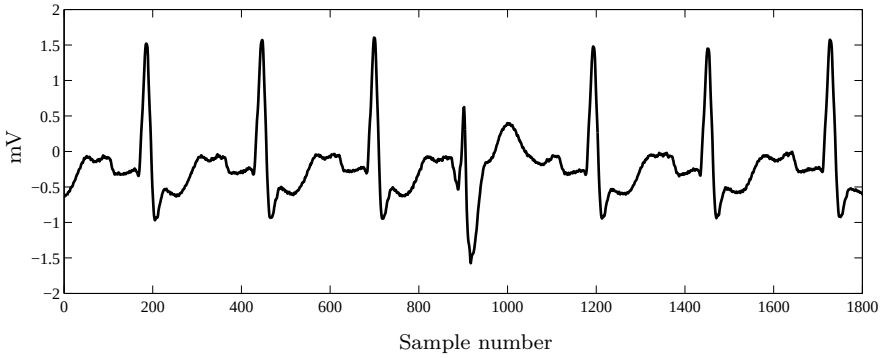

Figure 1.6: *Example of an isolated ventricular ectopic beat (4th beat).*

ability of generating a depolarization wave. Examples of these regions are the atrial and ventricular tissues [84]. Under abnormal conditions, these regions can take over the role of the SA node as a pacemaker, and cause *ectopic depolarization* that results in an *ectopic beat*. If the ectopic beat is originated in the atria/ventricles it is called atrial/ventricular premature complex (APC/VPC) or ectopic beat. An example of a ventricular premature beat, represented in Figure 1.6, clearly shows that the typical propagation pattern was disturbed. These ectopic beats should be ignored when analyzing the firing frequency of the SA node.

1.4.4 Mechanical effect of respiration on the ECG

It is well known that respiration influences the ECG in different ways. However, at this point only the mechanical influences will be described, while other effects will be described later in this chapter. First, the respiratory movements cause changes in the position of the electrodes with respect to the orientation of the heart vector [150]. These changes are translated as changes in amplitude in the different characteristic waves in the ECG.

A second mechanical influence of respiration is related to the electrical impedance of the thorax. In [85] it was found that changes in the thoracic impedance are closely correlated with the changes in the volume of air contained in the lungs. As a consequence, the amplitude of the ECG is also altered.

It is clear that the morphology of the ECG is strongly affected by the breathing cycle. Therefore, algorithms, as outlined in [54, 119, 217], have been proposed to extract the respiratory information from the ECG. In addition, these two

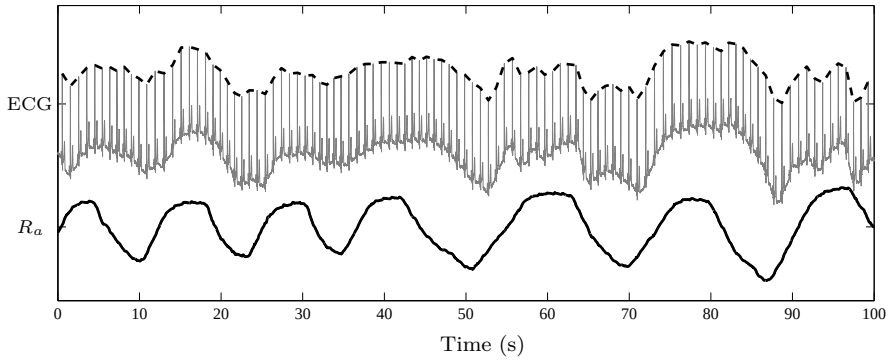

Figure 1.7: *Example of a simultaneous ECG and respiratory effort R_a measured in the abdomen using a respiratory belt. The dashed line indicates the variations in the R-peak amplitude of the ECG. Note the effect of respiration on the baseline of the ECG as well as on its amplitude.*

mechanical effects are more pronounced in the standard lead II [150]. Figure 1.7 shows an example of an ECG segment where the respiratory cycle is clearly affecting the amplitude of the R-peak.

Recently, it was shown that the mechanical interaction of respiration with the ECG affects not only the amplitude of the characteristic waves, but also the positions of the R-peaks [77]. The authors compared the heart rate derived from lead I with the one obtained using lead II, and they found that the "error" contained in lead II is closely correlated with the breathing cycle. This becomes relevant when comparing heart rates derived from different leads.

1.5 The autonomic nervous system

The autonomic nervous system (ANS) is the part of the central nervous system in charge of the unconscious regulation of several functions such as heart rate and respiratory rate. It has the remarkable capability of quickly modifying not only these two variables, but also the functioning of several organs within the body. This is done in order to adapt the organism to external stimuli and maintain homeostasis. The ANS is regulated by different centers in the hypothalamus, spinal cord, brain stem, and the cerebral cortex [84]. Therefore, when one of these centers, especially those located in the brain, does not function properly or is reached by abnormal discharges, such as epileptic seizures, the well-functioning of the ANS is compromised.

The ANS is subdivided into two branches: *the sympathetic nervous system* and the *parasympathetic nervous system*, which have opposite actions. For instance, the sympathetic nervous system is also known as the fight-or-flight response, while its counterpart the parasympathetic branch is responsible for the rest-and-digest actions. The sympathetic nervous system reacts to stressful situations in which fast adaptations are needed, as in obstructive sleep apnea events for example. During such events patients stop breathing due to a mechanical obstruction and the sympathetic nervous systems comes to the rescue causing the patients to wake up and start breathing again. On the other hand, the parasympathetic nervous system helps patients to recover after an acute situation like apnea or epileptic seizures by slowing down functions such as heart rate and respiration. Both branches of the ANS work in tandem, and they play a key role in the regulation of the heart rate and respiration, as will be described next.

1.5.1 Heart rate control

As mentioned before, the heart rate is determined by the firing rate of the SA node, and this on its turn is modulated by the ANS. The SA node and other parts of the heart are connected to different sympathetic and parasympathetic nerve fibers. In fact, the sympathetic nerves are connected to all regions of the heart, while the parasympathetic nerves are mainly connected to the SA and AV nodes. As a consequence, an increased sympathetic activation will result in an improved conductivity and rhythmicity of the heart. An increased parasympathetic activation, on the other hand, will reduce the rhythm of the SA node, and it will increase the delay between the activation of the AV node and the propagation of the electrical impulse to the ventricles. In other words, pronounced parasympathetic activation reduces heart rate. This activation is also called *vagal* activation, because the parasympathetic system communicates with the heart through the *vagus* nerve.

The reaction of the heart to sympathetic activation might take up to 15 seconds, while the effect of the parasympathetic activation is much faster, typically within one second. These two activations compete with each other, but they are balanced in a way that it allows the organism to quickly react and adapt to all circumstances. This balance is call the *autonomic tone* or *sympathovagal balance*, and it is regularly used to determine autonomic influences at particular moments in time. If an autonomic dysfunction takes place, a distortion in the balance between sympathetic and parasympathetic activations is expected. As a result, patients with autonomic dysfunctions might not be able to react or recover after stressful situations that can lead to life-threatening distortions of the cardiac cycle. In addition, it is known that overflow of sympathetic

activation is strongly related to the development of cardiovascular diseases [131]. Different conditions that are known to be associated with increased sympathetic activation include sleep apnea, obesity, hypertension, heart failure, and epilepsy [46, 59].

1.5.2 Respiratory control

Respiration has two main goals. One is to transport oxygen into the lungs from where it is distributed to the entire body, and the other is to facilitate the transport of carbon dioxide (CO_2) out of the body. This is done by means of different mechanisms in charge of extracting the oxygen from the lungs and transporting it to the different cells in the organism, and at the same time extracting their CO_2 and releasing it into the atmosphere.

The respiratory rate is mainly controlled by the autonomic nervous system, however, it is possible to override its effect by consciously changing the breathing pattern. The control center consists of a group of neurons located in the brainstem, especially in the medulla oblongata and pons [84]. These centers have afferent input from central and peripheral chemoreceptors, and receive descending control from higher brain systems such as the prefrontal cortex, amygdale and insula. The main role of these mechanisms in the respiratory control is to adjust the rate of alveolar ventilation in order to supply the body with the demands of oxygen and CO_2.

Different mechanisms can affect the respiration and they include sleep apnea and epilepsy. In apnea, mechanical blockages or dysfunctions of the respiratory centers can take place [60, 147], while in epilepsy the respiratory centers can be involved during an epileptic seizure, causing changes in the respiratory pattern [149, 176, 182].

1.6 Cardiorespiratory interactions

The mechanical interactions between respiration and the morphology of the ECG was discussed in the previous section. However, this is not the only relation that exists between cardiac and respiratory activities. Several studies have focused on understanding the different ways in which respiration affects the heart rate [17, 19, 28, 120, 220], and three different forms of cardiorespiratory coupling have been identified. These forms are the well-known respiratory sinus arrhythmia (RSA) [28], cardiorespiratory phase synchronization [19], and time-delay stability [18]. All of them occur at different

time scales, they are all independent from each other, and they can be related to different physiological mechanisms. For instance, it has been observed that the strength of each one of these forms of cardiorespiratory coupling changes across sleep stages, which suggests that heart rate and respiration respond to different autonomic controls [18]. The reason why these two mechanisms have a coupled response can be explained by the fact that central respiratory neurons directly influence cardiac neurons and that the inflation of the lungs causes a reflex cardiac response [120]. As a result, the pulmonary gas exchange and the regulation of central functions like arousal, are optimized [79].

1.6.1 Respiratory sinus arrhythmia

Respiratory sinus arrhythmia (RSA) is the most studied form of cardiorespiratory coupling [120] and it accounts for the rhythmical fluctuations of the heart rate in phase with the breathing cycle. This is to say, during inspiration the heart rate increases, and it decreases during expiration. Consequently, the RSA is, in a short-term, the main modulator of the heart rate, and it appears as a pronounced component in the spectrum of the latter. For this reason, the RSA has been extensively used as an indicator of vagal activity, since it is believed that respiration is mainly mediated by the parasympathetic nervous system [89]. However, this is still a matter of debate, since the complete physiological role of RSA is not yet fully understood [198].

Different factors are known to affect the magnitude of the RSA. For example, it is known that RSA is more pronounced during sleep, especially during slow wave sleep [79]. However, in patients suffering from obstructive sleep apnea (OSA), an increased sympathetic tone has been found to be responsible for a reduced RSA. As a result, OSA is considered a risk factor for developing different cardiovascular diseases. A different study showed that children who died of sudden infant death syndrome had a significantly lower RSA and HRV when compared against control subjects [110]. This was also found in epileptic children who subsequently died of sudden unexpected death in epilepsy (SUDEP) [182]. Therefore, it is clear that the magnitude of the RSA can be seen as a strong indicator of an underlying pathology.

1.6.2 Cardiorespiratory phase synchronization

The second form of coupling is the cardiorespiratory phase synchronization (CRPS) [19]. In this case, clusters of heart beats occur in phase with a consecutive number of breathing cycles. In comparison with the RSA, CRPS acts on a larger time scale, and it is believed that it is the result of an adjustment

of the cardiac and respiratory controls [19]. Furthermore, sleep studies have shown that CRPS is modified across sleep stages, and it does not respond to increases in the parasympathetic tone as its counterpart the RSA [17, 18, 19]. In addition, the sleep pattern displayed by CRPS seems to be independent on age, and this is not the case for RSA. These findings show that CRPS and RSA respond to different physiological factors, and they are both two independent forms of cardiorespiratory coupling.

1.6.3 Time-delay stability

The last form of coupling is the time-delay stability (TDS), recently studied in [18]. TDS accounts for the "stable" time delay between the occurrence of transient changes in the respiration and the reaction of the heart rate to these changes. A stronger TDS coupling is found during periods of constant time delay, and in normal subjects these periods are expected to be longer. This type of coupling has also been studied during sleep [18], and it was found to be larger during light sleep and wake than during rapid eye movement (REM) and deep sleep. These findings show once again, that this is an independent form of coupling since during deep sleep RSA and CRPS display both an opposite effect with respect to TDS.

Whereas RSA is the most studied form of coupling, CRPS and TDS have been proven to be coexisting forms of cardiorespiratory coupling. Therefore, it is important to keep these in mind when analyzing the interactions between respiration and heart rate. One way of doing this, is by not limiting this analysis to the well-defined high frequency (HF) band of heart rate variability (HRV) [51].

1.7 Ambulatory monitoring

Different ways of monitoring the cardiac activity exist in an ambulatory setting, and they can be classified in two main categories: continuous and intermittent systems. In continuous monitoring, the ECG signal is recorded during days or even weeks, and the goal is to be able to detect arrhythmias and abnormalities in the heart rhythm that cannot be detected during physical examination [114]. An example of such abnormalities are tachycardias produced by unprovoked epileptic seizures [24]. In those cases the continuous monitoring of the ECG signal is needed. This requires the patient to carry a system, e.g. Holter monitor, during their normal activities for several days. This is of course, the best way to capture isolated events that can help for the diagnosis or follow-up

of a certain disease. Nevertheless, the level of obtrusiveness of such systems can be sometimes untolerated by many patients [51].

Other ambulatory systems that appear to be more tolerated by patients are the intermittent systems. They can record the ECG signal whenever an event occurs, or at specific moments in time. Other intermittent systems used to monitor cardiac condition take discrete measurements of heart rate and other physiological parameters related to the condition of the heart. One example of this type of systems is the one developed during the Recap project [212]. In this project, the compliance of patients to the system was evaluated, and it was found that by encouraging patients and educating them, it was possible to not only monitor their physiological parameters daily, but also improve their self-care behavior. The latter is very important especially in cases where the outcome depends on the routine of the patients. This study showed that in order to develop a home monitoring system it is necessary to understand the effect that different technologies have on the outcome of the monitoring period. In this context, if a continuous monitoring ECG system is envisioned, it is necessary to first make the patients aware of the importance of their adherence to the system, and to identify the most appropriate technologies for each particular case. The main results of the Recap project are described in Appendix A.

In general, the type of ambulatory tools that can be used to accurately monitor a given condition is still determined by the representations of the disease itself. For instance, a continuous monitor is needed to detect unprovoked seizures, while an intermittent system can be used to monitor patients suffering from heart failure. The latter is due to the fact that symptoms of decompensation can be observed days before hospitalization or death.

Part I

Algorithms

Chapter 2

Machine Learning Techniques

Machine learning is a multidisciplinary field that draws upon several areas such as artificial intelligence, computer science, information theory, philosophy, and statistics, to address the capability of a computer program to learn. This learning capability can be associated with the automatic improvement of the performance of some computational tasks through experience [138]. Obviously, the improvement is strongly related to the way the learning algorithm is designed, which involves the definition of experience, the representation of both the examples and the target function, and the task that needs to be performed.

Two main machine learning tasks can be defined, namely *supervised* and *unsupervised* learning [88]. In supervised learning, an algorithm is presented with a set of training points for which the target is known. This "desired target" allows the algorithm to solve problems like function estimation and classification. In unsupervised learning on the other hand, no target is known and the algorithm learns the underlying structure of the input data. This leads to the solution of problems like clustering and principal component analysis.

This chapter describes the supervised and unsupervised techniques used in this work. In addition, it presents two methods for large scale data, to be used when the amount of training examples presented to the learning algorithms exceeds the limits of the computational resources (e.g. memory). Section 2.3.2 presents an original contribution of this work, namely a modification to the fixed-size algorithm proposed in [187].

2.1 Supervised learning techniques for classification

Supervised learning techniques aim to find a function $f(\cdot)$ that can predict a "known" target variable y from a set of input training data x. All training examples are given as a pair (x, y), where $y = f(x)$, and the main goal is to apply the learned function $f(\cdot)$ to new, unseen examples, and still obtain "reasonably good"[1] predictions.

One supervised learning technique is *classification*, where the target y corresponds to class labels, and the function $f(\cdot)$ is a decision boundary that best discriminates between the given classes in the training set. In addition, classification aims to find a good decision boundary that is able to *generalize* the learned concept on new data examples. This is to say, the label y_t for a new example x_t can be predicted reasonably well. Several ways to find these decision boundaries exist. For instance, linear predictors or more elaborated methods that lead to nonlinear discriminants can be used. In this section, three different approaches are discussed, namely linear discriminant analysis, support vector machines, and least-squares support vector machines.

2.1.1 Linear discriminant analysis

Linear discriminant analysis (LDA) for binary classification, projects a set of multivariate zero-mean data points $\{x_i\}_{i=1}^{N}$, with $x_i \in \mathbb{R}^d$, into a one-dimensional subspace (i.e. a line) as

$$z = w^T x, \tag{2.1}$$

based on which the best discrimination between two classes can be achieved [68]. This can be done by maximizing the *Rayleigh* quotient [187]

$$J(w) = \frac{w^T \Sigma_B w}{w^T \Sigma_W w}, \tag{2.2}$$

where Σ_B measures the between-class scatter defined as

$$\Sigma_B = (\mu_1 - \mu_2)(\mu_1 - \mu_2)^T, \tag{2.3}$$

with μ_1 and μ_2 the means of the input data in class 1 and class 2, respectively. The within-class scatter is defined as $\Sigma_W = \Sigma_1 + \Sigma_2$, where Σ_1 and Σ_2 are

[1] This performance depends on the tolerance to deviations or errors from a predefined target.

the covariance matrices of both classes, which are considered to be equal, i.e. $\Sigma_1 = \Sigma_2$. This can be seen as a limitation in many applications where $\Sigma_1 \neq \Sigma_2$ is often the case. However, LDA has been proven to outperform other classifiers in different real life problems [188].

Note that by maximizing (2.2), the between-class variances are maximized, and the within-class variances are minimized. Furthermore, in [83], it was shown that the solution to (2.2) corresponds to the following generalized eigenvalue problem

$$\Sigma_B w = \lambda \Sigma_W w, \tag{2.4}$$

where

$$\lambda = \frac{w^T \Sigma_B w}{w^T \Sigma_W w}. \tag{2.5}$$

Since Σ_W is a non-singular matrix, (2.4) can be written as

$$\Sigma_W^{-1} \Sigma_B w = \lambda w, \tag{2.6}$$

and after replacing Σ_B by the expression in (2.3), the solution to the problem can be found as

$$w = \Sigma_W^{-1}(\mu_1 - \mu_2). \tag{2.7}$$

Once w is computed, it is possible to obtain the projections of all the points in the dataset, and the means and covariances of the two classes in the new one-dimensional subspace. With these, a new point can be classified by means of the Mahalanobis distance.

It is known that LDA is very sensitive to outliers and highly unbalanced problems [88], due to the estimation of the sample means and covariances. However, it is possible to overcome these limitations by pre-processing the input data and carefully selecting the training set.

Another disadvantage of LDA is that it performs poorly when the decision boundary is highly non-linear, or when the covariance matrices of the classes are not equal, or highly dissimilar. In these cases, other algorithms like support vector machines (SVM) can be used as alternative.

2.1.2 Support vector machines

The main goal of support vector machines (SVM) is to find a hyperplane that best separates the different classes in the dataset, by maximizing the margin, or the distance between the hyperplane and the points closest to the decision boundary [203, 187]. In other words, if all the points closest to the hyperplane satisfy $|w^T x_i + b| = 1$, then the margin between the classes is $2/\|w\|_2$.

Linear SVMs were originally formulated for separable data in [202], then for non-separable cases in [53], and finally extended to non-linear versions in [203]. These extensions use the well-known *kernel trick*, which allows to map the input data into a high-dimensional feature space, where a linear classification is then performed. The main advantage of the kernel trick is that there is no need to actually compute the high-dimensional feature space. Instead, the linear or nonlinear mapping $\varphi(x)$ of an input data point x into the feature space is done by means of a positive definite kernel function $K(x, z)$, with $x, z \in \mathbb{R}^d$.

The nonlinear classification problem in SVM is formulated as the maximization of the margin subject to some tolerance to misclassification, and it is written as

$$\begin{aligned} \min_{w,b,\xi} J_P(w,\xi) &= \frac{1}{2} w^T w + c \sum_{i=1}^{N} \xi_i \\ \text{s.t.} \quad & y_i[w^T \varphi(x_i) + b] \geq 1 - \xi_i, \\ & \xi_i \geq 0, \quad i = 1, ..., N, \end{aligned} \tag{2.8}$$

where ξ_i are the slack variables introduced to relax the constraints, c is a regularization constant, and N the total number of points in the dataset. Note that the minimization of $w^T w/2$ is equivalent to the maximization of $2/\|w\|_2$.

After formulating the Lagrangian of the primal problem in (2.8), and characterizing the saddle point of the Lagrangian, the dual formulation becomes the following quadratic programming problem [187]:

$$\begin{aligned} \max_{\alpha} J_D(\alpha) &= -\frac{1}{2} \sum_{i,j=1}^{N} y_i y_j K(x_i, x_j) \alpha_i \alpha_j + \sum_{i=1}^{N} \alpha_i \\ \text{s.t.} \quad & \sum_{i=1}^{N} \alpha_i y_i = 0, \quad 0 \leq \alpha_i \leq c, \quad i = 1, ..., N, \end{aligned} \tag{2.9}$$

where $K(x_i, x_j) = \varphi(x_i)^T \varphi(x_j)$. Points corresponding to $\alpha_i \neq 0$ are the *support vectors*. In this way, sparsity is added to the problem, and only the distances between those points and the hyperplane are maximized.

The final formulation of the nonlinear SVM classifier is then

$$y(x) = \text{sign} \left[\sum_{i=1}^{N} \alpha_i y_i K(x, x_i) + b \right]. \tag{2.10}$$

The bias term b can be obtained from the KKT conditions for optimality as indicated in [187].

Different kernel functions can be used, for instance the linear kernel

$$K(x, x_i) = x_i^T x, \tag{2.11}$$

the polynomial kernel

$$K(x, x_i) = \left(\tau + x_i^T x\right)^d, \tag{2.12}$$

where d is the degree of the polynomial and $\tau \geq 0$, and the RBF kernel

$$K(x, x_i) = \exp\left(-\frac{\|x - x_i\|_2^2}{\sigma^2}\right). \tag{2.13}$$

2.1.3 Least-squares support vector machines

Least-squares SVM (LS-SVM) [189] simplifies the classification problem in (2.8) by defining the cost function as a least squares problem, and equality constraints instead of inequality constraints. The primal LS-SVM formulation is written as

$$\begin{aligned} \min_{w,b,e} J_P(w, e) &= \frac{1}{2} w^T w + \gamma \frac{1}{2} \sum_{i=1}^{N} e_i^2 \\ \text{s.t.} \quad & y_i[w^T \varphi(x_i) + b] = 1 - e_i, \quad i = 1, ..., N, \end{aligned} \tag{2.14}$$

where γ is a regularization constant. After formulating the Lagrangian and satisfying the KKT conditions for optimality, the dual problem corresponds to the following linear system [187]:

$$\left[\begin{array}{c|c} 0 & y^T \\ \hline y & \Omega + \mathbf{I}/\gamma \end{array}\right] \left[\begin{array}{c} b \\ \hline \alpha \end{array}\right] = \left[\begin{array}{c} 0 \\ \hline 1_N \end{array}\right], \tag{2.15}$$

where $\alpha = [\alpha_1, \ldots, \alpha_N]^T$ are the Lagrange multipliers, $y = [y_1, \ldots, y_N]^T$, $1_N = [1, \ldots, 1]^T$, and $\Omega_{ij} = y_i y_j K(x_i, x_j)$, $i, j = 1, \ldots, N$. From this, it is clear that a set of linear equations is solved instead of a quadratic programming problem as in (2.9). Another important difference between this formulation and the traditional SVM is the lack of sparseness. Here all points take values of $\alpha_i \neq 0$, which means that all points in the training set contribute to the construction of the model.

Different studies have shown that for some applications the performance of LS-SVM is comparable to or better than the one of SVM [139, 187, 197]. For further reading on SVM and LS-SVM, the reader is referred to [170], [186], [187], and [203].

2.2 Unsupervised learning techniques

Unsupervised learning techniques like principal component analysis (PCA) and clustering deal with unlabeled data, and aim to find the underlying structure of a given input set. In principal component analysis (PCA), for example, an input dataset $\mathcal{D}$ is transformed into a set of uncorrelated variables, also known as *principal axes* [103]. These variables correspond to the leading eigenvectors of the covariance matrix of $\mathcal{D}$, and they span a subspace that retains maximum variance of the data. A clear disadvantage of this technique is that only "linear" structures can be represented by these principal axes. However, these limitations are overcome by a natural "nonlinear" extension of PCA, namely kernel PCA [169]. This nonlinear PCA first applies a (possibly) nonlinear transformation φ to $\mathcal{D}$, in order to map the input space to a high-dimensional feature space $\mathcal{F}$. The mapping is performed by means of Mercer kernels K, which are functions that generate positive semidefinite kernel matrices [137]. Classical PCA is then applied in this kernel induced space, and the resulting principal axes lie now in $\mathcal{F}$, rather than in the input space. Both versions of PCA have been widely used in many applications such as dimensionality reduction, data compression, denoising, and pattern recognition.

The second unsupervised technique, described in section 2.2.3 is clustering, where $\mathcal{D}$ is partitioned into substructures, called *clusters* [98]. A cluster consists of objects that are similar to each other and dissimilar to objects in other clusters. Examples of clustering algorithms are the well-known k-means, hierarchical clustering, and spectral clustering. However, they all have disadvantages that are relevant for the type of applications explored in this research. For instance, k-means and hierarchical clustering assume an underlying distribution of the data [174], and spectral clustering requires to re-train the model every time a new data point is added. Therefore, it is of paramount importance to use a more advanced clustering algorithm that achieves good performances when other algorithms seem to fail. One such example is kernel spectral clustering (KSC), which was originally proposed in [10], and will be described later in this section.

2.2.1 Principal component analysis

Given a set of N zero mean data points $\{x_i\}_{i=1}^N$ with $x_i \in \mathbb{R}^d$, PCA finds projected variables $w^T x$ with maximal variance

$$\begin{aligned} \max_w \mathrm{Var}(w^T x) = &\quad w^T \mathrm{C} w \\ \text{s.t.} &\quad w^T w = 1, \end{aligned} \tag{2.16}$$

where C is the sample covariance matrix defined as

$$C = \frac{1}{N} \sum_{i=1}^{N} x_i x_i^T. \tag{2.17}$$

The solution to the problem in (2.16) can be found after computing its Lagrangian as follows

$$\mathcal{L}(w; \lambda) = \frac{1}{2} w^T C w - \lambda (w^T w - 1), \tag{2.18}$$

with λ the Lagrange multiplier. After satisfying the conditions for optimality $\partial \mathcal{L}/\partial w = 0$ and $\partial \mathcal{L}/\partial \lambda = 0$, the solution to this problem corresponds to the eigenvalue problem

$$Cw = \lambda w. \tag{2.19}$$

The matrix C is positive semidefinite, which leads to $\lambda \geqslant 0$, and the eigenvector w corresponding to the largest eigenvalue indicates the direction of maximal variance of the projection variables.

2.2.2 Kernel principal component analysis

The formulation of KPCA used in this work corresponds to the LS-SVM formulation [188]. This formulation is based on primal-dual interpretations, where the dual problem is equivalent to the original formulation of KPCA described in [169]. The reason to use the LS-SVM formulation relies on the fact that it allows to perform estimations in the primal space, which is an advantage for large scale applications, where solving the dual problem for a very large dataset becomes infeasible [187]. This formulation also allows to perform out-of-sample extensions. In Section 2.3 two large scale methodologies will be discussed, and the LS-SVM formulation of KPCA will be described below.

Given a set of N data points $\{x_i\}_{i=1}^N$ with $x_i \in \mathbb{R}^d$, the projected variables $w^T \varphi(x_i)$, and $\varphi(\cdot) : \mathbb{R}^d \to \mathbb{R}^{d_h}$ the mapping to a high-dimensional feature

space $\mathcal{F}$ of dimension d_h, the *primal problem* can be defined as

$$\begin{aligned} \max_{w,e} J_P(w,e) &= \gamma \frac{1}{2}\sum_{i=1}^{N} e_i^2 - \frac{1}{2} w^T w \\ \text{s.t.} \quad e_i &= w^T(\varphi(x_i) - \hat{\mu}_\varphi), \quad i = 1,\ldots,N, \end{aligned} \tag{2.20}$$

with γ a positive regularization constant, and $\hat{\mu}_\varphi = (1/N)\sum_{i=1}^{N}\varphi(x_i)$. The Lagrangian of this primal problem is defined as

$$\begin{aligned} \mathcal{L}(w,e;\alpha) = \quad & \gamma \frac{1}{2}\sum_{i=1}^{N} e_i^2 - \frac{1}{2} w^T w \\ & - \sum_{i=1}^{N} \alpha_i (e_i - w^T(\varphi(x_i) - \hat{\mu}_\varphi)), \end{aligned} \tag{2.21}$$

with Karush-Kuhn-Tucker (KKT) conditions for optimality

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial w} = 0 \rightarrow & w = \sum_{i=1}^{N} \alpha_i(\varphi(x_i) - \hat{\mu}_\varphi) \\ \frac{\partial \mathcal{L}}{\partial e_i} = 0 \rightarrow & \alpha_i = \gamma e_i, \quad i = 1,\ldots,N \\ \frac{\partial \mathcal{L}}{\partial \alpha_i} = 0 \rightarrow & e_i - w^T(\varphi(x_i) - \hat{\mu}_\varphi) = 0, \quad i = 1,\ldots,N. \end{cases} \tag{2.22}$$

After eliminating the primal variables e, w and defining $\lambda = 1/\gamma$, the dual formulation becomes

$$\Omega_c \alpha = \lambda \alpha, \tag{2.23}$$

where Ω_c is the centered kernel matrix with entries defined as

$$\Omega_{c,ij} = (\varphi(x_i) - \hat{\mu}_\varphi)^T(\varphi(x_j) - \hat{\mu}_\varphi), \qquad i,j = 1,\ldots,N. \tag{2.24}$$

The mapping $\varphi(\cdot)$ might be unknown but based on the Mercer theorem it can implicitly be defined by means of a kernel function $K(x_i,x_j) = \varphi(x_i)^T\varphi(x_j)$. This kernel function is then used to compute a kernel matrix $\Omega_{ij} = K(x_i,x_j)$, and its centered version $\Omega_c = M_c \Omega M_c$, with the centering matrix

$$M_c = I - \frac{1}{N} 1_N 1_N^T, \tag{2.25}$$

and $1_N = [1,1,\ldots,1]^T$. In this work the RBF kernel function is used.

The eigenvalue problem (2.23) satisfies the KKT conditions (2.22), and every eigenvalue-eigenvector pair of the centered kernel matrix can be seen as a candidate solution of (2.20). The direction of maximal variance in the high-dimensional space, i.e. *the first principal component*, will be determined by the eigenvector corresponding to the largest eigenvalue of Ω_c. The total amount of components that can be extracted is determined by the size of the kernel matrix which is given by the number of training points.

The projection of an input data point x onto a subspace spanned by the most relevant principal axis in the feature space, also known as the score variable, can be obtained from (2.20) as

$$\begin{aligned} e(x) &= w^T(\varphi(x) - \hat{\mu}_\varphi) \\ &= \sum_{i=1}^{N} \alpha_i \Bigg(K(x_i, x) - \frac{1}{N}\sum_{j=1}^{N} K(x_j, x) - \frac{1}{N}\sum_{j=1}^{N} K(x_j, x_i) \\ &\qquad + \frac{1}{N^2}\sum_{j=1}^{N}\sum_{k=1}^{N} K(x_j, x_k) \Bigg). \end{aligned} \tag{2.26}$$

For denoising applications, it is necessary to map the transformed data points from the feature space back into the input space. Given a point ϕ in the feature space $\mathcal{F}$, the *pre-image* problem consists of finding a point x in the input space such that $\phi = \varphi(x)$. Pre-image is recognized as an ill posed problem, where the solution does not always exist, and if it exists, it is not always unique [137][117]. This problem becomes of primary interest in denoising applications, and several methodologies have been proposed to find approximate pre-images by solving the following optimization problem

$$x^\star = \underset{x}{\operatorname{argmin}} \, \|\phi - \varphi(x)\|^2, \tag{2.27}$$

where $x^\star$ is the approximate pre-image of ϕ. In this work an iterative fixed point algorithm for RBF kernels [137] is used.

A well known challenge in unsupervised learning techniques such as PCA and kPCA is the selection of the model parameters. For example the number of components n_c is required in both PCA and kPCA, while the kernel parameter is needed only in kPCA. This problem is tackled in this work and a model selection approach is proposed and described in Chapter 3.

2.2.3 Kernel spectral clustering

Kernel spectral clustering [10] is based on weighted kernel principal component analysis (PCA)[6] and primal-dual LS-SVM formulations discussed in [187].

Problem definition

The LS-SVM formulation for kPCA discussed above was modified in [6] by introducing a weighting scheme that links kPCA with different spectral clustering algorithms. This weighted formulation defines a new primal problem where a symmetric positive definite, diagonal matrix V is introduced into the problem defined in (2.20) without centering, as follows

$$\begin{aligned} \min_{w,e} J_P(w,e) &= \gamma \frac{1}{2N} \sum_{i=1}^{N} v_i e_i^2 - \frac{1}{2} w^T w \\ \text{s.t.} \quad e_i &= w^T \varphi(x_i), \quad i = 1, \ldots, N, \end{aligned} \tag{2.28}$$

where $\{v_i\}_{i=1}^N$ are the diagonal entries of the weighting matrix V. This is to say, $V = \text{diag}[v_1, \ldots, v_N]$. It has been shown in [6] that each eigenvalue/vector pair of

$$V\Omega\alpha = \lambda\alpha \tag{2.29}$$

is a candidate solution of (2.28). In addition, given that the kernel matrix Ω can be seen as a similarity matrix, where $k(x_i, x_j)$ corresponds to the similarity between the points x_i and x_j, the degree of each data point can be computed as

$$\deg(x_i) = \sum_{j=1}^{N} k(x_j, x_i), \tag{2.30}$$

which is equal to the sum of the columns (or rows) of Ω. These degrees are used to compute a degree matrix $D = \text{diag}(\deg(x_1), \ldots, \deg(x_N))$ [50], and if $V = D^{-1}$ is used as the weighting matrix, the eigenvalue problem in (2.29) is equivalent to a binary NCut problem [6]. However, in order to be able to obtain more than a binary partition of the data, the formulation proposed in [10] is necessary. This new formulation uses $k-1$ binary partitions, with k the

total number of clusters, and it extends (2.28) as follows

$$\min_{w^{(l)},e^{(l)},b_l} \quad \frac{1}{2N}\sum_{l=1}^{k-1}\gamma_l {e^{(l)}}^T D^{-1} e^{(l)} - \frac{1}{2}\sum_{l=1}^{k-1} {w^{(l)}}^T w^{(l)}, \tag{2.31}$$

$$\text{subject to} \quad \begin{cases} e^{(1)} = \Phi w^{(1)} + b_1 1_N, \\ e^{(2)} = \Phi w^{(2)} + b_2 1_N, \\ \vdots \\ e^{(k-1)} = \Phi w^{(k-1)} + b_{k-1} 1_N, \end{cases}$$

where 1_N is the all-ones vector, γ_l are positive regularization constants, $e^{(l)}$ is the compact form of the projections $e^{(l)} = [e_1^{(l)}, \ldots, e_N^{(l)}]$, and $\Phi = [\varphi(x_1), \ldots, \varphi(x_N)]^T$. Note that by taking $q^{(l)} = \text{sign}(e_i^{(l)})$ one can get $k-1$ binary cluster indicators q.

It was proven in [10] that the Karush-Kuhn-Tucker (KKT) optimality conditions of the Lagrangian of (2.31) are satisfied by the eigenvalue problem defined as

$$D^{-1} M_D \Omega \alpha^{(l)} = \lambda_l \alpha^{(l)}, \quad l = 1, \ldots, k-1, \tag{2.32}$$

where $\lambda_l = N/\gamma_l$, and $\Omega_{ij} = \varphi(x_i)^T \varphi(x_j) = K(x_i, x_j)$. The matrix M_D is a weighting centering matrix

$$M_D = I_N - \frac{1}{1_N^T D^{-1} 1_N} 1_N 1_N^T D^{-1}, \tag{2.33}$$

b_l are the bias terms that result from the optimality conditions

$$b_l = \frac{1}{1_N^T D^{-1} 1_N} 1_N^T D^{-1} \alpha^{(l)}, \quad l = 1, \ldots, k-1, \tag{2.34}$$

and I_N the identity matrix. The effect of M_D and b_l is to center the kernel matrix Ω by removing the weighted mean from each column. As a result, the eigenvectors are zero-mean. Given this and due to the fact that the eigenvectors are piecewise constant, it is possible to use the eigenvectors corresponding to the first $k-1$ eigenvalues to partition the dataset into k clusters. This approach corresponds to the error-correcting output code (ECOC) used in [10], where a codeword $c_p \in \{-1, 1\}^{n_e}$, with $n_e = k-1$, is defined to represent each cluster $\mathcal{A}_p$, where $p = 1, \ldots, k$. The codebook $\mathcal{C} = \{c_p\}_{p=1}^k$ is computed after binarizing the eigenvector matrix $[\text{sign}(\alpha^{(1)}), \ldots, \text{sign}(\alpha^{(n_e)})]$, and finding the k most common occurrences. After that, in order to assign a given data point

x_i to any of the clusters, the Hamming distances between its codeword and the codebook are computed. Then, x_i is assigned to the cluster $\mathcal{A}_p$ for which the Hamming distance is the lowest. A different encoding scheme can use n eigenvectors to find 2^n clusters, however, the ECOC showed to be more robust against errors introduced by noise [10].

Using (2.31), the score variables can be expressed in terms of the eigenvectors $\alpha^{(l)}$ as

$$e_i^{(l)} = {w^{(l)}}^T \varphi(x_i) + b_l = \sum_{j=1}^{N} \alpha_j^{(l)} K(x_i, x_j) + b_l. \tag{2.35}$$

To illustrate this methodology, the toy example shown in Figure 2.1 is used. It is clear that the number of clusters is $k = 3$, hence the clustering information is contained in the two eigenvectors corresponding to the two highest eigenvalues of (2.32). These two eigenvectors, $\alpha^{(1)}$ and $\alpha^{(2)}$, are also indicated in Figure 2.1, and it is clear that for each element in the dataset, a codeword can be obtained as $[\alpha_i^{(1)}\ \alpha_i^{(2)}]$. For this particular toy example, the first 50 elements of the set can be characterized by the code [-1 -1], the next 50 elements by [1 -1], and the last 50 by [1 1]. This example uses the RBF kernel with kernel parameter $\sigma^2 = 0.21$. The selection of this parameter has a direct influence on the partition of the data. For instance, when σ^2 is too small, the similarity between the points drops to zero and the kernel matrix is then close to the identity matrix. The opposite case occurs when σ^2 is too big, and all the points look similar to each other. In these extreme situations, the algorithm is unable to partition the data properly. Therefore, the selection of the kernel parameter needs to be done based on a suitable criterion. In this work, the balanced line-fit (BLF) criterion proposed in [10] is used. This criterion exploits the fact that the eigenvectors and the corresponding projections of the data points are piecewise, and it measures the co-linearity between the projections of points in the same cluster. These projections are computed using a validation set $\mathcal{D}^v$. Their collinearity with respect to a clustering $\{\mathcal{A}_1, \ldots, \mathcal{A}_k\}$ for $k > 2$ is defined as

$$\text{linefit}(\mathcal{D}^v, k) = \frac{1}{k} \sum_{p=1}^{k} \frac{k-1}{k-2} \left(\frac{\varsigma_1^{(p)}}{\sum_l \varsigma_l^{(p)}} - \frac{1}{k-1} \right), \tag{2.36}$$

with $\varsigma_1^{(p)} \geqslant \cdots \geqslant \varsigma_{k-1}^{(p)}$ are the ordered eigenvalues of

$$C_{\tilde{Z}}^{(p)} = \frac{1}{|\mathcal{A}_p|} {\tilde{Z}^{(p)}}^T \tilde{Z}^{(p)}, p = 1, \ldots, k, \tag{2.37}$$

where the matrix $\tilde{Z}^{(p)} \in \mathbb{R}^{|\mathcal{A}_p| \times (k-1)}$ represents the zero-mean score variables for the validation points assigned to cluster $\mathcal{A}_p$. When $k = 2$ the collinearity

measure is defined as

$$\text{linefit}(\mathcal{D}^v, k) = \sum_{p=1}^{2} \left(\frac{\breve{\varsigma}_1^{(p)}}{\breve{\varsigma}_1^{(p)} + \breve{\varsigma}_2^{(p)}} - \frac{1}{2} \right), \tag{2.38}$$

with $\breve{\varsigma}_l^{(p)}$ the ordered eigenvalues of $\breve{Z}^{(p)} \in \mathbb{R}^{|\mathcal{A}_p| \times 2}$, which corresponds to the zero-mean score variables of the validation points assigned to cluster $\mathcal{A}_p$ when $k = 2$.

To illustrate the collinear property of the projections of the data on the space spanned by the eigenvectors of (2.32), Figure 2.2 shows the score variables for all the points in the dataset. Note that the collinerarity of the points in the score variable space is directly related to the performance of the algorithm, which on its turn is determined by the selection of the kernel parameter σ^2.

The BLF criterion also considers the balance of the clusters. Therefore a balance measure is defined as

$$\text{balance}(\mathcal{D}^v, k) = \frac{\min\{|\mathcal{A}_1|, \ldots, |\mathcal{A}_k|\}}{\max\{|\mathcal{A}_1|, \ldots, |\mathcal{A}_k|\}}, \tag{2.39}$$

where $|\mathcal{A}_1|$ is the number of elements in cluster 1. When $\text{balance}(\mathcal{D}^v, k)$ is equal to 1 the clusters have the same number of points, and when it is equal to 0, the clusters are extremely unbalanced. Using expressions (2.36), (2.38), and (2.39), the BLF criterion is defined as

$$\text{BLF}(\mathcal{D}^v, k) = \eta \text{linefit}(\mathcal{D}^v, k) + (1 - \eta)\text{balance}(\mathcal{D}^v, k), \tag{2.40}$$

with $0 \leqslant \eta \leqslant 1$ a user-defined parameter that determines the importance of the balance measure with respect to the collinearity.

Extensions to new data points

One of the goals of a machine learning technique, is to be able to apply a learned concept to new data. For KSC in particular, this means that the model should be able to cluster new data. This is the so-called *out-of-sample extension*, which is implemented using the projections of the mapped new points onto the $\alpha^{(l)}$ eigenvectors obtained with the training dataset.

For a test set of N_t data points $\{x_m^\text{t}\}_{m=1}^{N_\text{t}}$, a test kernel matrix $\Omega^\text{t} \in \mathbb{R}^{N_t \times N}$, with entries $\Omega_{mj}^\text{t} = K(x_m^\text{t}, x_j)$ is computed. This matrix can be seen as a similarity measure between the test and the training points. To assign a test point x^t to one of the k clusters, its projections onto the $\alpha^{(k-1)}$ eigenvectors,

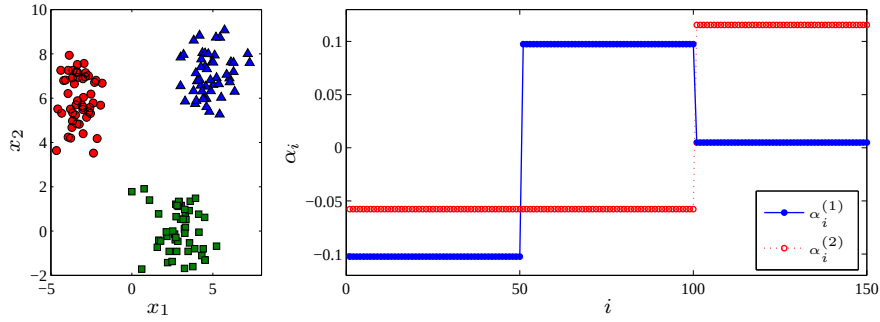

Figure 2.1: *(left) Dataset with 150 data points and 3 clusters. (right) eigenvectors corresponding to the 2 largest eigenvalues of* (2.32)*, which contain the clustering information. Note that if the eigenvectors are binarized it is possible to obtain a codeword for each cluster, where the first 50 elements will be assigned to a cluster with code [-1 -1], the next 50 will go to cluster [1 -1], and the last 50 to [1 1].*

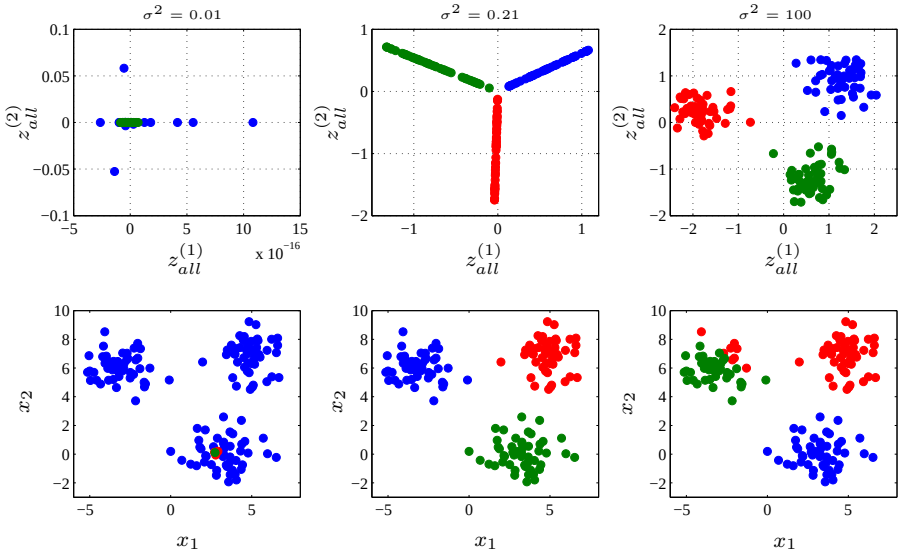

Figure 2.2: *Score variables for all the points in toy example dataset (top) and clustering results (bottom) for different values of the kernel parameter σ^2. Note that for the optimal value of $\sigma^2 = 0.21$, the partition of the data is perfect, and the collinearity of the score variables is highest.*

or so-called score variables, are first computed as

$$z^{(l)}(x^t) = \sum_{i=1}^{N} \alpha_i^{(l)} K(x_i, x^t) + b_l \qquad l = 1, \ldots, k-1, \tag{2.41}$$

and then binarized as $\text{sign}(z^{(l)}(x^t))$. A codeword for each test point is obtained and compared with the codes obtained for the training set, using the Hamming distance. This allows to associate each test point to the most similar cluster.

2.3 Large scale problems

From the description of KPCA and KSC it is clear that it is necessary to evaluate the kernel function $K(x, x_i)$, $i = 1, \ldots, N$, in all the training examples and to compute its eigendecomposition. This is a problem since the computational cost of these algorithms directly depends on the size of the kernel matrix, and in large scale applications this becomes a limitation.

Different methodologies have been proposed to tackle this issue, for example in [196] a probabilistic approach is used to reduce the number of training vectors. The size of the kernel matrix can also be reduced by means of low rank approximation methods such as Nyström [218], and incomplete Cholesky factorization (ICF) where a set of approximated eigenvectors are computed [9], [11]. These methodologies aim to find an approximated kernel matrix $\Omega \in \mathbb{R}^{M \times M}$ where $M \ll N$. Using the Nyström approximation method, the reduced set of support vectors is selected at random but this does not guarantee a close representation of the underlying distribution of the data. A more suitable selection of support vectors is achieved using the fixed-size method (FS) which actively selects the training examples using an iterative procedure. FS and ICF are used in this work and they will be discussed below.

2.3.1 Incomplete Cholesky factorization

Any symmetric, positive definite matrix $A \in \mathbb{R}^{N \times N}$ can be decomposed as $A = HH^T$, where H is a lower triangular matrix, called the Cholesky factor of A. This decomposition is not always unique or it may not even exist when A is positive semidefinite. However, it is still possible to compute a lower triangular matrix $G \in \mathbb{R}^{N \times Q}$ with $Q < N$ and $N - Q$ columns of G equal to zero, such that $A \approx GG^T$. This incomplete Cholesky factorization [83] allows to compute a "close" approximation of the exact Cholesky factor H, where $\|A - GG^T\|_2^2 \leq \eta$, with η a user-defined error threshold. The matrix G is

formed in an iterative way by performing symmetric permutations of the rows and columns of A, and by keeping up to Q pivots for which the sum of the remaining $N - Q$ pivots is lower than η. Therefore, the larger the value of η, the less pivots are selected, and the worse the approximation is. The set of remaining $N-Q$ data points can be split into validation and test sets. This low rank approximation of a positive semidefinite matrix has been used to reduce the computational complexity of interior point methods in SVM training [67], to compute approximated kernel ICA contrast functions [8, 14], and to find a reduced set of representative training points for sparse kernel spectral clustering [9], [11].

The incomplete Cholesky factorization is used to find an approximate Cholesky factor G of the kernel matrix Ω, such that $\Omega \approx GG^T$. From the factorization of the kernel matrix it is possible to compute an approximated set of eigenvectors as described in [9, 11] in the framework of sparse kernel spectral clustering. This approximation is summarized in Algorithm 1.

Algorithm 1. Approximated eigenvectors using ICF [11]

Input: Dataset $\mathcal{D} = \{x_i\}_{i=1}^N$, with $x_i \in \mathbb{R}^d$, N the total number of data points, a positive definite kernel function $K(x_i, x_j)$, kernel bandwidth σ^2, error parameter η.

Output: Approximated eigenvectors $\alpha^{(l)}$, for $l = 1 \ldots, N$, subset of pivots $\mathcal{D}_{\text{pv}} = \{x_n\}_{n=1}^Q$, validation set $\mathcal{D}^\nu = \{x_m^\nu\}_{m=1}^{N_\nu}$ and test set $\mathcal{D}^{ts} = \{x_j^{ts}\}_{j=1}^{N_{ts}}$, where $N = Q + N_\nu + N_{ts}$.

1. Compute the approximated Cholesky factor G of the kernel matrix defined by $\mathcal{D}$, and obtain the subset of pivots $\mathcal{D}_{\text{pv}} = \{x_n\}_{n=1}^Q$.
2. Compute the left eigenvectors U and singular values Λ of the matrix G
3. Compute the eigendecomposition in (2.45) and approximate the eigenvectors $\alpha^{(l)} = U\rho^{(l)}$.
4. Split the remaining $N - Q$ data points into validation $\mathcal{D}^\nu$ and test $\mathcal{D}^{ts}$ sets.

The singular value decomposition (step 2) of the triangular matrix $G \in \mathbb{R}^{N\times Q}$ for $Q < N$ is:

$$G = U\Lambda V^T \tag{2.42}$$

where the columns of $U \in \mathbb{R}^{N\times Q}$ correspond to the left singular vectors, $\Lambda \in \mathbb{R}^{Q\times Q}$ contains the singular values in the diagonal, and the columns of $V \in \mathbb{R}^{Q\times Q}$ are the right singular vectors of G. From this decomposition, the incomplete Cholesky factorization of Ω can then be written as:

$$\Omega \approx U\Lambda^2 U^T. \tag{2.43}$$

This decomposition can be used for instance in kPCA, to rewrite (2.23) as:

$$M_c U \Lambda^2 U^T \alpha^{(l)} = \lambda_l \alpha^{(l)}, l = 1, \ldots, n_c, \tag{2.44}$$

where M_c is the centering matrix defined in (2.25) and n_c the number of principal components to compute.

Premultiplying (2.44) by U^T results in

$$\begin{aligned} U^T M_c U \Lambda^2 U^T \alpha^{(l)} &= \lambda_l U^T \alpha^{(l)} \\ U^T M_c U \Lambda^2 \rho^{(l)} &= \lambda_l \rho^{(l)}, \end{aligned} \tag{2.45}$$

with $\rho^{(l)} = U^T \alpha^{(l)}$ and $l = 1, \ldots, n_c$.

Remark 1 *Note that the complexity of the eigendecomposition has been reduced because the training matrix is now of size $Q \times Q$, and the eigenvectors $\alpha^{(l)} \in \mathbb{R}^N$ can be approximated by $U^T \rho$, with $\rho^{(l)} \in \mathbb{R}^Q$, where $Q < N$.*

2.3.2 Fixed-size LS-SVM

Another way to select a representative training set is the fixed-size method, which was originally proposed in [187]. It actively selects the M training samples that best represent the underlying distribution of the dataset by maximizing the quadratic Renyi entropy defined as

$$H_R = -\log \int p(x)^2 dx, \tag{2.46}$$

which is related to kernel PCA and density estimation [81] through

$$\int \hat{p}(x)^2 dx = \frac{1}{N^2} 1_N^T \Omega 1_N, \tag{2.47}$$

where Ω is the kernel matrix.

In this work a slight modification to the original implementation presented in [187] is proposed under the name FS2. Here, the initial support vectors are not selected at random, instead, they correspond to the centroids of k clusters, computed using k-means. It is well-known that the k-means algorithm is sensitive to the initialization. However, since the resultant centroids are used as initial points, and even if the initialization did not lead to the best clustering partition, the iterative procedure of the maximization of the Renyi entropy takes care of finding the underlying structure of the data. In other words, the goal of k-means is to initialize FS2 with a more suitable data than the one obtained at random.

Remark 2 *Note that the centroids of the clusters obtained with k-means, already represent a certain structure in the data, which allows to initialize the search of support vectors with a relatively high value of entropy.*

The second modification to the original algorithm avoids the random replacement of candidate support vectors. Here, the candidate support vector that indicates the maximum similarity with the rest of the "training points" is replaced with another vector in the remaining set. In this way, the quadratic Renyi entropy converges faster. Algorithm 2 summarizes the proposed approach.

For example, for the set of 1500 data points shown in Figure 2.3, 50 points were selected as training set using FS2 and the regular fixed-size algorithm. It is clear that for this particular example, after only one iteration, FS2 already reaches a higher entropy value than the one obtained by the regular fixed-size algorithm after 4000 iterations.

Algorithm 2. Modified Fixed-size method (FS2)

Input: Dataset $\mathcal{D} = \{x_i\}_{i=1}^{N}$, with $x_i \in \mathbb{R}^d$, N the total number of data points, a positive definite kernel function $K(x_i, x_j)$, the number of support vectors M, and total number of iterations.

Output: Set $\mathcal{D}_{\text{SV}} = \{x_n\}_{n=1}^{M}$ of M support vectors that maximize the quadratic Renyi entropy, validation set $\mathcal{D}^{\nu} = \{x_m^{\nu}\}_{m=1}^{N_\nu}$ and test set $\mathcal{D}^{ts} = \{x_j^{ts}\}_{j=1}^{N_{ts}}$, where $N = M + N_\nu + N_{ts}$.

1. Apply the k-means algorithm to the dataset with $k = M$. The centroids of the clusters are selected as the initial set of possible support vectors $\mathcal{D}_{\text{SV}} = \{x_n\}_{n=1}^{M}$.
2. Compute the kernel matrix $\Omega \in \mathbb{R}^{M \times M}$ for the initial set of support vectors.
3. Compute the quadratic Renyi entropy of Ω using (2.46).
4. Sum the columns (or rows) of Ω and select the support vector that corresponds to the column (or row) with the highest value. This support vector corresponds to the point that is most similar to the rest of the points in the subset. Replace it by a randomly selected data point from the remaining $N - M$ data points.
5. Compute (2.46) for the subset $\mathcal{D}_{\text{SV}}$ with the new candidate. If the entropy increases, keep the data point in the subset and repeat step 4. If the entropy decreases select another random point from the remaining examples and repeat.
6. Repeat 4 and 5 until the number of iterations is exceeded or until the entropy has converged.
7. Split the remaining $N - M$ data points into validation $\mathcal{D}^{\nu}$ and test $\mathcal{D}^{ts}$ sets.

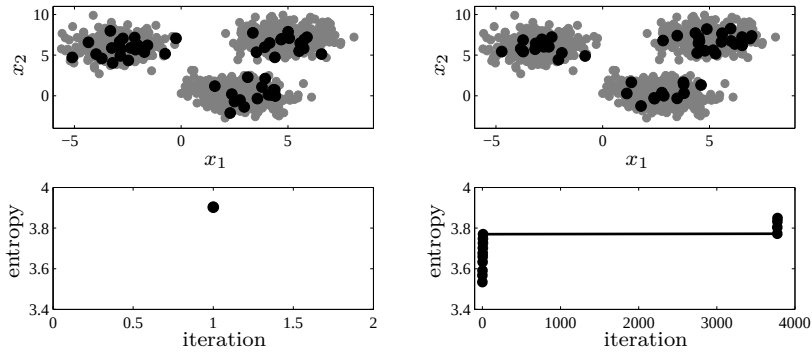

Figure 2.3: *(top) Dataset of 1500 data points indicated in grey, and 50 training points (black) selected using FS2 (left), and the regular fixed-size algorithm (right). (bottom) Entropy values for each iteration step. Note that with FS2 (left), the first iteration already reached a higher entropy value than the one achieved by regular fixed-size algorithm (right) after almost 4000 iterations.*

2.4 Conclusions

This chapter was a theoretical introduction to the machine learning techniques used in this research. First the supervised techniques for classification were discussed. Next, principal component analysis and its kernel version were presented, together with the formulation for kernel spectral clustering. The last part of the chapter dealt with two different methodologies that can be used to tackle large scale problems. These two methodologies are the incomplete Cholesky factorization and a modification of the fixed-size LS-SVM method. The latter methodology is one of the contributions of this work, and it was concluded that it finds a better subset of training points than the conventional fixed-size algorithm. This is to say, a highest value of entropy can be achieved in a lower number of iterations. These two large-scale techniques will be used and evaluated in the next chapter.

Chapter 3

Model Selection for KPCA

This chapter presents one of the main contributions of this research, namely a methodology for the automatic selection of the model parameters for KPCA. As mentioned in the previous chapter, this is a well-known challenge in unsupervised learning and different heuristics [94, 137] and (semi-) theoretical models have been proposed to tackle it. For instance, in [7] a general loss function is included in the formulation of the problem and the model selection is carried out by means of a heuristic. A different approach is proposed in [105], where a noise level is estimated using parallel analysis. This method uses a null-dataset to estimate the level of noise. It does have some disadvantages, firstly the null-dataset requires many permutations, and secondly, the cost function does not always indicate a clear maximum. Moreover, in order to have good noise level estimation in low dimensions, it is necessary to rotate the input data into a higher dimensional space. This last approach defines the first principal components as the meaningful information in the data, and the rest of components are associated to noise. This denoising concept is related to the minimum description length (MDL) principle [163], which finds the minimum number of components needed to fully represent the intrinsic information in the data. The separation of the noise components is performed in the eigenvalue distribution, which is also the way in which the proposed approach works. Here, a new way to estimate the level of noise is presented. It does not require rotations nor permutations, it can handle low dimensionality, and it is observed that the objective function has a global maximum for all the given examples. This new approach is called model selection based on distance distributions (MDD), and it uses the statistical properties of the distribution of distances of the input data to generate a noise distribution, based on the assumption that uniformly distributed data generate unimodal (Euclidean)

distance distributions. This property has been studied in [30, 91, 113] and used in [122] in the framework of visualization and exploration of high-dimensional data.

This chapter is organized as follows. In Section 3.1 the problem of manifold learning by estimating the noise in different dimensionality is discussed. The criterion proposed in this study to select the model parameters n_c and σ^2 is presented in Section 3.2. Experimental results using toy examples and real life problems are discussed in Section 3.4, and conclusions are presented in Section 3.5.

3.1 Noise and dimensionality

It is well known that the main problem in principal component analysis is to separate the "components" containing the information, from those describing the noise in the dataset. One way of doing this is by comparing the eigenvalues computed from the data with the ones from a set of uncorrelated variables (i.e., null-dataset), as proposed in parallel analysis [92], and later used in its kernel version [105]. This idea of spectral comparison is also used in this study, but with a novel approach to estimate the noise eigenvalues. Details on this approach will be described later. For now, the focus will be on the problems that are encountered when the null-data is created from permutations of the input dataset as in [92, 105]. This will allow to get more insights on the problem of noise estimation.

To find the noise spectrum, the input dataset can be first permuted several times along the dimensions, in order to destroy any manifold structure in the dataset [105]. After each permutation, a centered kernel matrix is computed with its corresponding eigenvalues as in (2.23). Finally, the noise spectrum is then characterized by the 95th percentile of the obtained eigenvalue distribution.

As an example, one can take a structured dataset (see Figure 3.1 top) $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^{100}$, with $\mathbf{x}_i \in \mathbb{R}^2$, and another dataset with the same structure in 50 dimensions: $\mathcal{D}' = \{\mathbf{x}'_i\}_{i=1}^{100}$, with $\mathbf{x}'_i \in \mathbb{R}^{50}$. To create $\mathcal{D}'$, $\mathcal{D}$ is projected onto $\mathbb{R}^{50}$ where 48 dimensions equal zero. This new dataset is $\mathcal{D}^{50} = \{\mathbf{x}_i^{50}\}_{i=1}^{100}$, with $\mathbf{x}_i^{50} = [\mathbf{x}_i, 0, \ldots, 0]^T$ where $\mathbf{x}_i^{50} \in \mathbb{R}^{50}$. Then, a rotation matrix $\mathrm{W} \in \mathbb{R}^{50\times 50}$ of the form:

$$\mathrm{W} = \begin{bmatrix} \mathbf{w} & 0 & \ldots & 0 \\ 0 & \mathbf{w} & \ldots & 0 \end{bmatrix}, \tag{3.1}$$

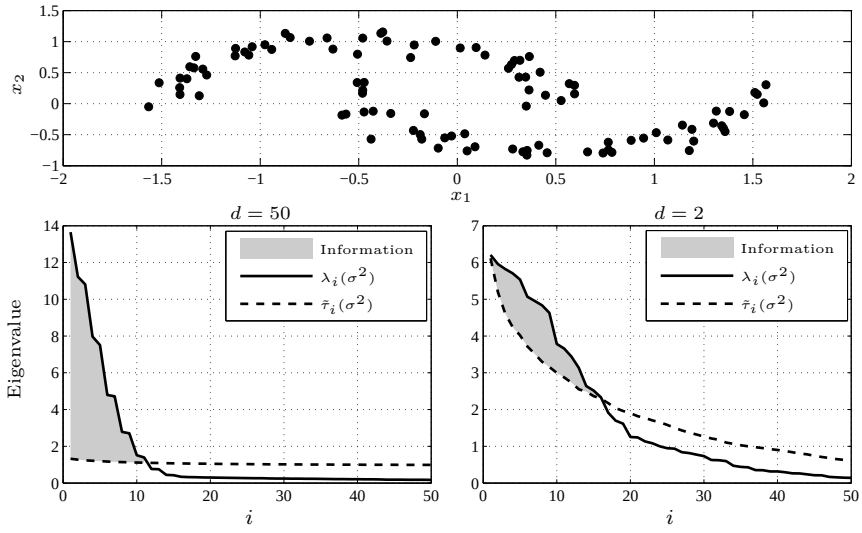

Figure 3.1: *(top) Structured dataset $\mathcal{D}$, where $N = 100$, and $d = 2$. (bottom) Data and noise spectra of the dataset in a* 50*, and in a* 2 *dimensional space. The first 50 eigenvalues λ_i of a centered kernel matrix computed for a given σ^2, are displayed. The noise levels, $\hat{\tau}_i$ are defined as the 95th percentiles of the eigenvalue distributions after* 100 *permutations of the input data. Note that after permuting the data in 2 dimensions, the noise level is highly similar to the eigenvalue spectrum of the data. This indicates the presence of structure in the permuted dataset. In 50 dimensions this is no longer the case, instead, the noise level resembles the spectrum of a random dataset.*

with $\mathbf{w} = [w(1), \ldots, w(25)]^T$ the 25-point symmetric hamming window vector, is used to generate the high-dimensional dataset $\mathrm{D}' = \mathrm{WD}^{50}$, with $\mathrm{D}' = [\mathbf{x}'_1, \ldots, \mathbf{x}'_{100}]$, and $\mathrm{D}^{50} = [\mathbf{x}^{50}_1, \ldots, \mathbf{x}^{50}_{100}]$.

The spectrum of $\mathrm{D}' \in \mathbb{R}^{50}$ for a given σ^2 is shown in the bottom left panel of Figure 3.1. There, a clear difference between the eigenvalues of the structured data (solid line) and the ones of the permuted data (dashed line) can be observed. This implies that for 50 dimensions, the permutations successfully eliminate the existing structure contained in the set $\mathcal{D}$. This, however, is no longer the case in lower dimensions, as it is shown in the bottom right panel of Figure 3.1, where a dataset in $d = 2$ is used. This phenomenon, also know as *curse of dimensionality* [25], is due to the fact that the higher (lower) the dimensionality, the higher (lower) the entropy, and the harder (easier) it is to find any shape or any correlation between the dimensions.

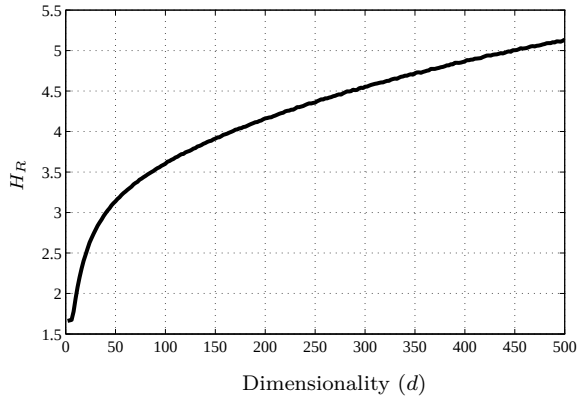

Figure 3.2: *Quadratic Renyi entropy H_R of a structured dataset of* 100 *points contained in different dimensions. Note that as the dimensionality increases, the harder it is to find any structure in the dataset.*

Another way of observing this effect, is by measuring the quadratic Renyi entropy defined as

$$H_R = -\log \int p(x)^2 dx. \tag{3.2}$$

In [81], (3.2) is related to kernel PCA and density estimation through

$$\int \hat{p}(x)^2 dx = \frac{1}{N^2} 1_N^T \Omega 1_N, \tag{3.3}$$

where Ω is the kernel matrix. This is shown in Figure 3.2, where the quadratic Renyi entropy H_R for the given dataset increases (i.e., more uniformity) as a function of the dimensionality d.

The fact that any structure in a dataset cannot be easily destroyed in lower dimensions, makes algorithms like kernel parallel analysis [105] (or a method based on *random data points*) unsuccessful in estimating noise in low dimensionality spaces. This limitation is overcome with the method proposed in this study, which is able to get an estimation of the noise level, even in lower dimensions. This method is called model selection based on distance distributions (MDD).

Before going into the details of MDD, however, the influence of dimensionality on the distribution of distances in a dataset needs to be discussed. It has been shown in [91, 30, 113, 122] that for uniformly distributed (i.e. unstructured) points, the probability distribution of the distances between the points is unimodal. Furthermore, the statistical moments of the distribution change

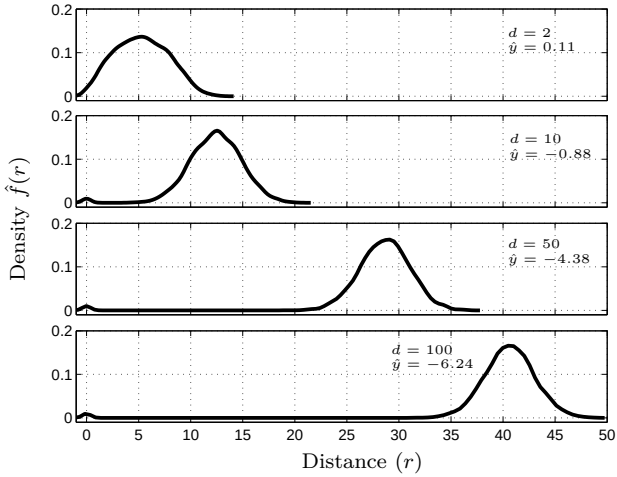

Figure 3.3: *Density functions of distances r for a dataset of 100 uniformly distributed points in spaces of 2, 10, 50 and 100 dimensions. Note that the skewness $\hat{y}$ of the distribution decreases as d increases.*

as a function of the dimensionality. This is indicated in Figure 3.3 where the distribution of $N = 100$ uniformly distributed data points in $[0, 1]^d$ clearly shows a unimodal structure, and its statistical moments change for increasing dimensionality[1]. Therefore, the main assumptions made in this study, are that the entropy for unstructured datasets (i.e. noise or uniformly distributed data points) is higher than the one for structured datasets, and that noise is characterized by a unimodal distance distribution, as if it were generated by a uniformly distributed dataset.

3.2 MDD: model selection based on distance distributions

The assumption made at the end of the previous section is essential for the method presented here, which can be used to select the number of components n_c and the RBF kernel parameter σ^2. This method is called Model selection based on Distance Distributions (MDD).

[1]The density function is computed using kernel smoothing density estimation, and the smoothing parameter is computed using *Silverman's rule*, defined in [181].

Given a training set $\mathcal{D}^t = \{x_n^t\}_{n=1}^{N_t}$, $x_n^t \in \mathbb{R}^d$, and a validation set $\mathcal{D}^v = \{x_m^\nu\}_{m=1}^{N_\nu}$, $x_m^\nu \in \mathbb{R}^d$, the Euclidean distances between $x_i \in \mathcal{D}^t$ and $x_j \in (\mathcal{D}^t \cup \mathcal{D}^v)$, $r_{ij} = \|x_i - x_j\|_2$, with $i = 1, \ldots, N_t$ and $j = 1, \ldots, N_t + N_v$, are computed. This results in a distribution of distances $\hat{f}_{\mathcal{D}}(r)$. From this distribution, the first four statistical moments, namely, the mean μ, standard deviation σ, skewness $\hat{y}$, and kurtosis β, are obtained and used to generate a second distribution[2] $\hat{f}_{\mathcal{N}}(\hat{r})$, with $0 \leq \hat{r} \leq r_{max}$, and r_{max} the maximum distance contained in the dataset. This second distribution is generated in order to mimic randomization of the input data points, and estimate the distance distribution of a noisy dataset. Values drawn from $\hat{f}_{\mathcal{N}}(\hat{r})$ are used as entries to a new distance matrix $\hat{R} \in \mathbb{R}^{N_t \times N_t}$, and they are organized in such a way that the conditions of a metric are satisfied. These conditions for arbitrary points x_i, x_j, a, are defined as

$$\begin{aligned}
&1. \quad \hat{r}(x_i, x_j) \geq \quad 0 \\
&2. \quad \hat{r}(x_i, x_j) = \quad 0, \ \text{iff} \ \ i = j \\
&3. \quad \hat{r}(x_i, x_j) = \quad \hat{r}(x_j, x_i) \\
&4. \quad \hat{r}(x_i, a) \leq \quad \hat{r}(x_i, x_j) + \hat{r}(x_j, a).
\end{aligned} \tag{3.4}$$

The first condition guarantees non-negativity of the metric, and it is satisfied by selecting points that are only included in $\hat{f}_{\mathcal{N}}(\hat{r})$. The second one is fulfilled by making $\hat{r}(x_i, x_i) = 0$ for $i = 1 \ldots, N$, and the symmetry of $\hat{R}$ will guarantee the third condition. The last condition is the so-called *triangle inequality* and needs to be satisfied in a recursive fashion and again, only values selected from $\hat{f}_{\mathcal{N}}(\hat{r})$ will be used in the matrix. Details on how to generate this new distance matrix are presented below.

3.2.1 Computation of the distance matrix

When a set of values is drawn from the distribution $\hat{f}_{\mathcal{N}}(\hat{r})$, these values can be organized in the matrix $\hat{R}$ in such a way that the entry $\hat{R}_{ij} = \hat{r}(x_i, x_j)$ is always lower or equal than the minimum value of the pairwise sum of the columns i and j. This is:

$$\hat{R}_{ij} \leq \min(\hat{R}_i + \hat{R}_j), \tag{3.5}$$

[2]The Pearson system is used to generate the distribution. In matlab this corresponds to the function *pearsrnd*, contained in the Statistics toolbox.

where $\hat{R}_i$ and $\hat{R}_j$ correspond to the ith and jth column of the matrix $\hat{R}$. Once this condition is satisfied, one can say that the matrix $\hat{R}$ satisfies the triangle inequality $\forall ij$. The next example shows how to generate this matrix.

For a dataset of N data points, both distance matrices R (original), and $\hat{R}$ (generated) have dimensionality $N \times N$. It is clear that the diagonal consists of only zeros, and that they are both symmetric matrices. Therefore, it is sufficient to generate $N_r = (N^2 - N)/2$ data points $\{\hat{r}_i\}_{i=1}^{N_r}$ from $\hat{f}_{\mathcal{N}}(\hat{r})$, where $\hat{r}_1 \geq \hat{r}_2 \geq \cdots \geq \hat{r}_{N_r}$. These values can then be organized in the off-diagonal entries of the lower triangular matrix of $\hat{R}$:

$$\hat{R}_L = \begin{bmatrix} 0 & & & & & \\ \hat{r}_1 & 0 & & & & \\ \hat{r}_2 & \hat{r}_N & 0 & & & \\ \hat{r}_3 & \hat{r}_{N+1} & \hat{r}_{2N-2} & \cdots & & \\ \vdots & \vdots & \vdots & \vdots & 0 & \\ \hat{r}_{N-1} & \hat{r}_{(2N-3)} & \hat{r}_{(3N-4)} & \cdots & \hat{r}_{N_r} & 0 \end{bmatrix}.$$

Once this matrix has been created, the condition in (3.5) is satisfied.

Keep in mind that a distance matrix usually does not satisfy the order imposed in the approach described above. This order is introduced here because it allows to speed up the creation of $\hat{R}$. However, to guarantee that the eigenvalue distribution of a kernel matrix computed using $\hat{R}$ is not affected by this ordering, a distance matrix generated in an iterative way from $\hat{f}_{\mathcal{N}}(\hat{r})$ was computed. During the generation of this matrix, each entry of $\hat{R}$ was checked in each iteration. The eigenvalue distribution of the kernel using this matrix was then compared with the one obtained from the ordered distance matrix. This was done for different datasets and different number of points, and it was found that both distributions do not differ significantly ($p > 0.05$). Hence, by using the fast and efficient approach described in this section, it is possible to keep the eigenvalue distribution and at the same time satisfy the different conditions in (3.4). It is necessary to emphasize that both matrices were created using the same unimodal distribution.

3.2.2 Maximization of the information content

Once $\hat{R}$ is created, its entries will be used in an RBF kernel as

$$K_{\mathcal{N}}(x_i, x_j) = \exp\left(-\frac{\hat{r}(x_i, x_j)^2}{2\sigma^2}\right), \tag{3.6}$$

and the eigenvalue spectrum $\hat{\tau}$ for a given σ^2, is estimated from

$$\Omega_{\mathcal{N}_c}\nu = \hat{\tau}\nu, \tag{3.7}$$

with $\Omega_{\mathcal{N}_c} = M_c\Omega_{\mathcal{N}}M_c$, M_c as defined in (2.25), and $\Omega_{\mathcal{N}_{ij}} = K_{\mathcal{N}}(x_i, x_j)$.

This artificial spectrum characterizes the noise level and can be compared with the one of the original data λ_i (2.23). Keep in mind that the noise is estimated using a validation set, which avoids overfitting of the training data. The number of components for a given σ^2, that contain information about the structure in the dataset will then be selected as

$$n_c(\sigma^2) = \max_{\lambda_i - \hat{\tau}_i > 0} i, \tag{3.8}$$

and the kernel parameter σ^2 can be obtained after maximizing the "structural information" I_s expressed by n_c with respect to the noise, defined as

$$I_s(\sigma^2) = \sum_{i=1}^{n_c(\sigma^2)} \lambda_i - \hat{\tau}_i. \tag{3.9}$$

This is not the same as maximizing the variance explained by the first n_c components $\sum_{i=1}^{n_c} \lambda_i$ in the sense that here the estimated noise information (unstructured information) contained in each component is also taken into account. This criterion of maximizing the difference in variance between the data and the estimated noise is also used in [105] under the name of energy. The main difference between the estimation of I_s and energy relies on the definition of noise. In this work a novel procedure to estimate the noise is proposed, which works in both high- and low-dimensionality. This is not the case for the noise estimation proposed in [105], which only works if the data is contained in a high-dimensional space (see section 3.1). For low-dimensional spaces, there exist the possibility that the noise spectrum estimated with [105] will indicate a larger variance than the one in the input data. Besides, when the SNR is very small, the spectrum of the data still indicates the presence of structure. The proposed methodology is summarized in Algorithm 3, and illustrated in the following example.

The datasets indicated in Figure 3.4 represent the same structure but with different signal-to-noise ratios (SNR). If the distance distribution for each example is calculated, it is observed that the behavior tends to be unimodal as soon as more noise is added to the dataset. This is shown in Figure 3.5. Note that the eigenvalue spectrum for the case with SNR=1, still indicates the presence of information related to the original dataset. Now, if the distance distribution of the last case is compared against a unimodal one, which

is representing uniformly distributed points, the differences in noise spectra become more evident (see Figure 3.6).

Algorithm 3. MDD-KPCA

Input: Training dataset $\mathcal{D}^t = \{x_n^t\}_{n=1}^{N_t}$, with $x_n^t \in \mathbb{R}^d$, validation set $\mathcal{D}^\nu = \{x_m^\nu\}_{m=1}^{N_\nu}$, with $x_m^\nu \in \mathbb{R}^d$, a positive definite kernel function $K(x_i, x_j)$, search grid of p possible values of the kernel parameter $\mathcal{S} = \{\sigma_s^2\}_{s=1}^{p}$, with $\sigma^2 \in \mathbb{R}$.

Output: Number of components n_c, kernel parameter σ^2, and denoised or reconstructed dataset $\mathcal{D}^\star = \{x_i^\star\}_{i=1}^{N}$, with $\mathcal{D}^\star = \mathcal{D}^{t\star} \bigcup \mathcal{D}^{\nu\star}$, and $N = N_t + N_\nu$.

1. Compute the distances between $\mathcal{D}^t$ and $\mathcal{D}^t \bigcup \mathcal{D}^\nu$, and the distance distribution $\hat{f}_{\mathcal{D}}(r)$. Obtain its statistical moments, and define the distance range $[0, r_{max}]$.
2. Generate a distribution $\hat{f}_{\mathcal{N}}(r)$ from the information computed in 1, and construct the distance matrix $\hat{R} \in \mathbb{R}^{N_t \times N_t}$ satisfying (3.4).
3. Compute the kernel matrix $\Omega_{\mathcal{N}} \in \mathbb{R}^{N_t \times N_t}$ for a given σ_s^2, and the eigenvalues $\breve{\gamma}_l$, $l = 1, \ldots, M$ using (3.7).
4. Repeat steps 2 and 3, hundred times and define the noise level $\hat{\gamma}_l$ as the 95th percentile of the distribution of each eigenvalue $\breve{\gamma}_l$.
5. For σ_s^2 compute the kernel matrix Ω, and the eigenvalues λ and eigenvectors α of (2.23).
6. Find the number of components $n_c(\sigma_s^2)$ from (3.8) and the amount of information $I_s(\sigma_s^2)$ contained in n_c from (3.9).
7. Repeat from step 2 until all the values from S are taken.
8. Select the n_c and σ^2 that maximize I_s.
9. Use the eigenvectors $\alpha^{(l)}(\sigma^2)$, $l = 1, \ldots, n_c$ to compute the denoised datasets $\mathcal{D}^{t\star}$ and $\mathcal{D}^{\nu\star}$ after applying the algorithm proposed in [137].

In the last case, one can see that by creating a unimodal distance distribution from the statistical moments of the original one, and by limiting the values to the maximum distance in the real set, as done in MDD, a certain structure is imposed to the noise. This can be observed in the first eigenvalue of the spectrum of the noise. From the spectrum of the estimated noise, it is also possible to see that the power is now spread among all the other components, as expected for uniform noise. The main difference between this spectrum and the one produced by data permutation [105], is that for MDD the estimated

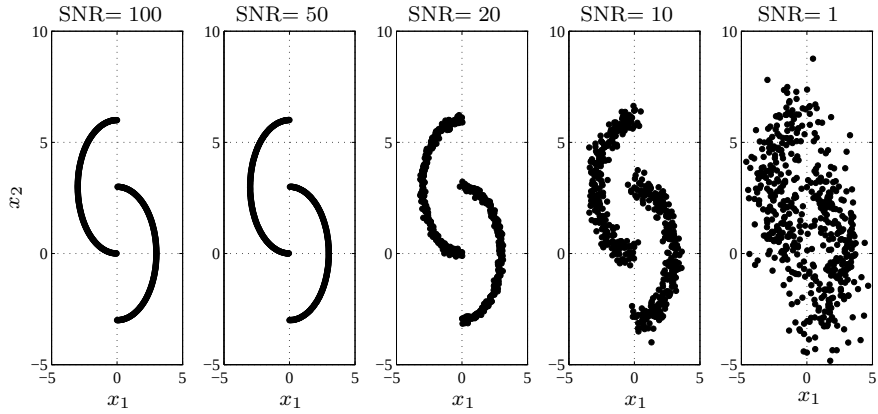

Figure 3.4: *Toy datasets with different signal to noise ratio.*

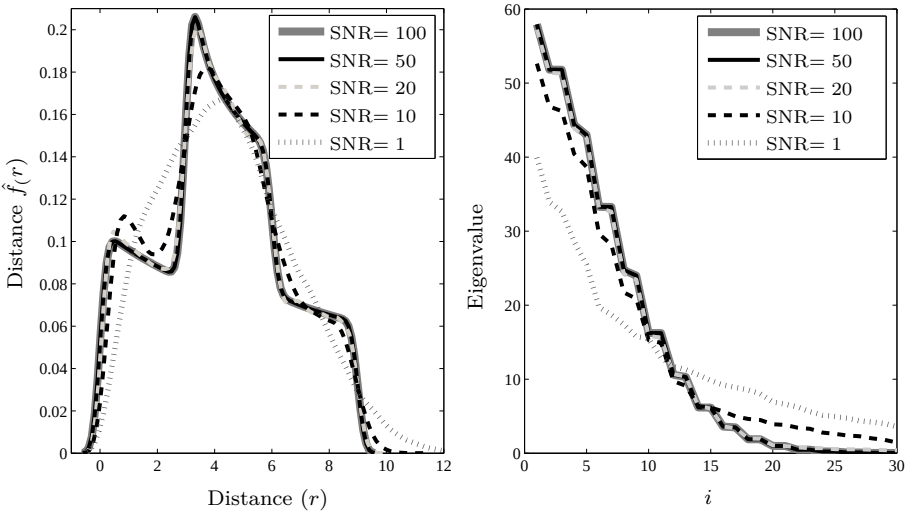

Figure 3.5: *Distance distributions (left) of datasets with different SNR. (right) Eigenvalue spectra for each case, with* $\sigma^2 = 0.8$.

noise does not follow the same pattern of the real data. Therefore, only when the data is permuted to estimate the noise, those two spectra may encounter each other in more places. Remember that the four statistical moments of the distance distribution are needed to estimate the noise. This was discussed in Section 3.1, and it is related to the fact that important information about the dimensionality of the data is contained in the values of skewness $\hat{y}$, and kurtosis β. For low dimensionality problems, using the first two statistical moments of

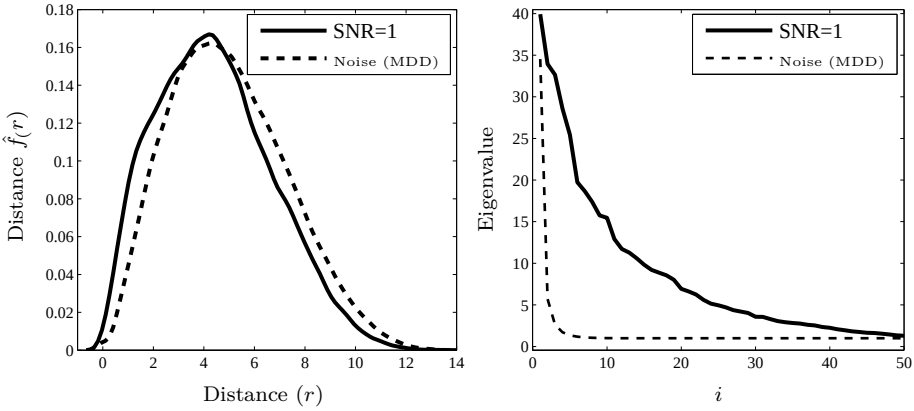

Figure 3.6: *Distance distributions (left) of the dataset with SNR=1, and the generated one using the statistical moments of the first one, as proposed in MDD. (right) Eigenvalue spectra, with* $\sigma^2 = 0.8$.

the original distance distribution is enough, since with less dimensions the distribution is closer to a normal one.

Now, concerning the first component, as mentioned before, components are only selected when the noise is not overestimated. In other words, when the first component of the noise is lower than the one of the real data.

An aspect to be considered here is the fact that the distance matrix is generated from a probability distribution, which requires the implementation of a Monte Carlo simulation. Note that the original distance values are not re-used or re-sampled as in bootstrapping methods, on the contrary, new distance values are drawn from $\hat{f}_{\mathcal{N}}(\hat{r})$, which implies a data generation process characteristic of Monte Carlo methods. Here, the eigenvalue spectrum is defined as the 95th percentile confidence interval of each eigenvalue distribution generated after q initializations. In this study $q = 100$. Another aspect to keep in mind, is that as the number of data points increases, the estimation of the statistical moments of the distance distribution becomes more accurate. Therefore, deviations in this estimation need to be considered when the size of the dataset is very small.

The proposed methodology is illustrated in the following example.Given is the dataset shown in Figure 3.7(a) with $N = 200$, where $N = N_t + N_\nu$. The distance distribution computed from the data, and one generated from its statistical moments, are indicated in Figure 3.7(b). After generating a distance matrix $\hat{R}$, and computing its eigenvalues $\hat{\tau}$, hundred times, the 95th percentile of

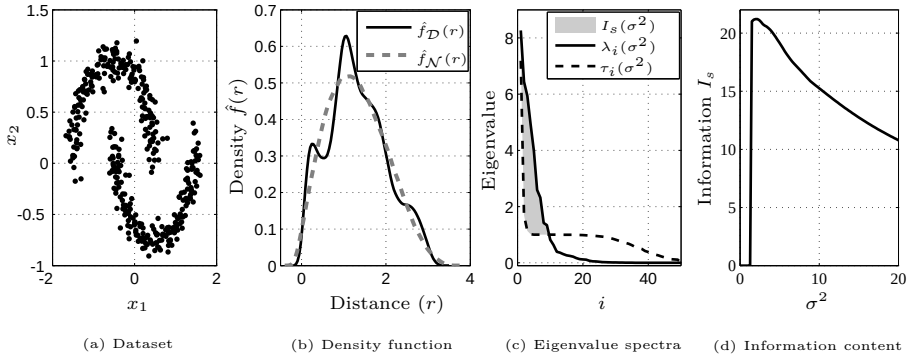

(a) Dataset (b) Density function (c) Eigenvalue spectra (d) Information content

Figure 3.7: *Selection of number of components n_c and σ^2. The dataset consists of $N_t = 100$ and $N_\nu = 100$, (b) density function of the distances in the data $\hat{f}_{\mathcal{D}}(r)$ (solid line), and the noise $\hat{f}_{\mathcal{N}}(r)$ (dashed line). (c) Eigenvalue spectrum of the data λ_i(solid line) and the noise $\hat{\tau}_i$ (dashed line). The shaded area corresponds to the structural information $I_s(\sigma^2)$ explained by the first nine components. (d) Information content as a function of σ^2.*

the eigenvalue distribution is obtained and is represented by the dashed line in Figure 3.7(c). The shaded area can be interpreted as the amount of information contained in the first nine components. Figure 3.7(d) shows the information content I_s as a function of the kernel parameter σ^2.

It is also possible to maximize I_s for a known number of components. There are some applications in which n_c is already given [217], and one only needs to estimate the best σ^2.

3.3 Large scale KPCA

The selection of the model parameters is not the only challenge in unsupervised learning. A different problem arises when large-scale datasets need to be analyzed. When N is large, building a kernel matrix Ω becomes impractical. To overcome this problem, low rank approximations of the kernel matrix such as the Nyström method [218] and Incomplete Cholesky factorization (ICF) [83] have been used [9, 11, 14, 187]. They reduce the size of the eigenvalue problem and approximate the *full-size* eigenvectors. As a result, the model can be trained with a smaller set, and the projections of the remaining data points can be predicted. This problem was already discussed in Section 2.3, and here both large scale methodologies, namely the modification of the fixed-size

algorithm FS2 (see 2.3.2), and ICF (see 2.3.1) will be used in combination with the model selection procedure presented in Section 3.2. Their implementation is summarized in Figure 3.8.

It is worth noting that the methodology based on ICF can be used when the σ^2 is given and one only wants to find the optimal number of components n_c. It could also be used to optimize the kernel parameter but its high computation time will make it impractical. This is due to the fact that for different values of σ^2 a different set of pivots is selected and a new approximated Cholesky factor is computed. By doing so, a new distance matrix must be obtained which requires higher computation times as the number of pivots increase. With this in mind, it is better to use ICF for cases when σ^2 is given.

3.4 Experimental Results

This section describes some experimental results, which were carried out in MATLAB R2012a on an Intel® Core™ i7, 3.4GHz, 7.8 GB RAM, running Ubuntu 12.04 LTS. The RBF kernel function was used for all the experiments, and a grid search for the bandwidth parameter $\mathcal{S} = \{\sigma_s^2\}_{s=1}^{p}$ was defined for each dataset separately, around the rule of thumb[3]

$$\hat{\sigma}^2 = 0.1 \times d \times \mathbb{E}\left[\mathrm{Var}[\mathcal{D}']\right], \tag{3.10}$$

where $\mathcal{D}' = \{x'_j\}_{j=1}^{d}$, $x'_j \in \mathbb{R}^N$, and $\mathrm{Var}[\mathcal{D}'] = [Var[x'_1], Var[x'_2], \ldots, Var[x'_d]]$, where each entry corresponds to the variance of one dimension in the dataset. The factor 0.1 was selected experimentally.

Two methods were used for comparison, the maximization of the Shannon entropy of the kernel matrix [217] (Shannon), and kernel parallel analysis (kPA) [105]. The first method only allows to find the kernel parameter σ^2. Therefore, the number of components n_c was selected using the methodology proposed in this chapter (MDD), based on the maximization of the structural information content I_s. With the second method, however, it is possible to find both model parameters simultaneously, and the results were obtained by means of a publicly available code[4] proposed in [105]. A comparison between the two algorithms for large scale KPCA is also presented in this study.

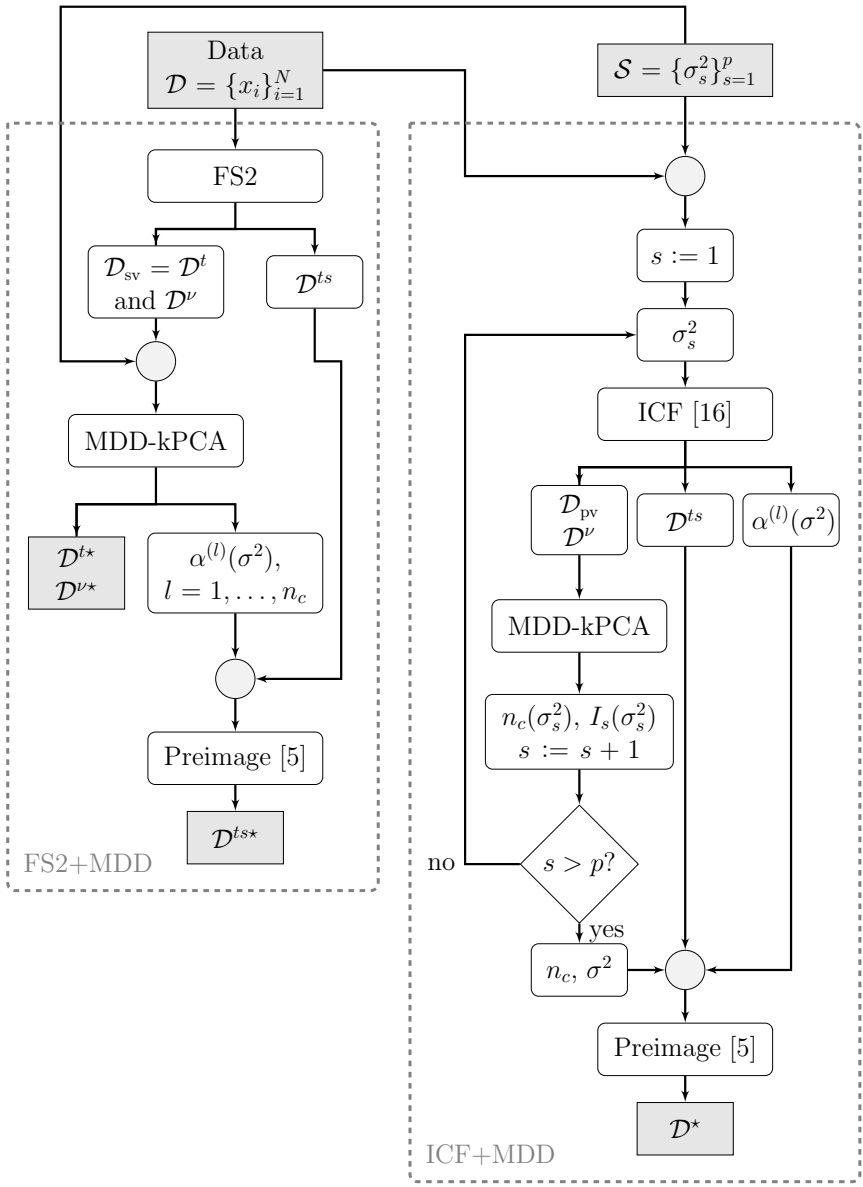

Figure 3.8: *Implementation of large scale KPCA with the proposed model selection algorithm MDD.*

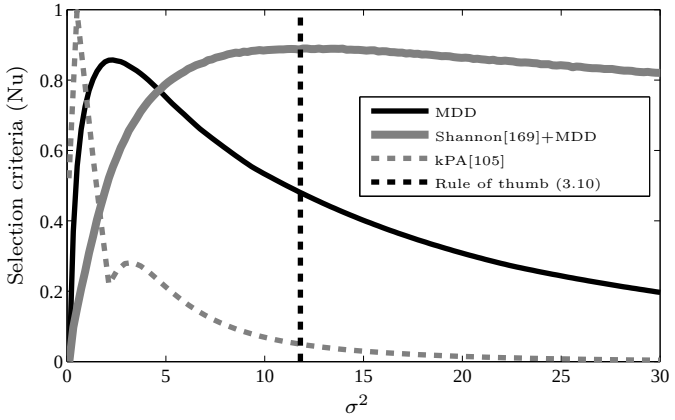

Figure 3.9: *Model selection criteria for different values of* σ^2*. Note that the kPA criteria has two local maxima, while the methodology proposed in this work reveals a global one. The rule of thumb of* (3.10) *gets a value close to the maximum of the curve of the Shannon entropy. This rule of thumb can be taken as an upper bound for the definition of the search grid of* σ^2*.*

3.4.1 Toy examples

Three artificial datasets were created after adding noise (SNR= 10) to predefined clean patterns. These patterns were then used to evaluate the performance of the algorithms after denoising. The first comparison was made between the denoised results obtained using all the points, and using the two large scale methodologies, namely the modified fixed-size (FS2) and incomplete Cholesky factorization (ICF). Each model selection algorithm was tested using these three methodologies.

The first dataset consists of two semicircles of $N = 500$ points contained in a space of two dimensions. For the first large scale approach proposed in this work, namely FS2+MDD, the training set was selected by maximizing the quadratic Renyi entropy with a subset of 50 support vectors. The remaining 450 points were split into validation set $N_\nu = 250$ and test set $N_{ts} = 200$. The selected model corresponds to $\sigma^2 = 2.19$ and $n_c = 9$. For the second approach (ICF+MDD), a set of pivots was determined using ICF with $\eta = 10^{-2}$ and $\sigma^2 = 1.13$ computed using (3.10). The resultant number of components was $n_c = 24$, for 70 pivots and $N_\nu = 250$. The definition of the parameter η can help to limit the number of pivots that are selected in cases where the

[3] used in the tutorial of http://www.esat.kuleuven.be/stadius/lssvmlab/
[4] http://www2.imm.dtu.dk/~kwjo/index.php?page=code

Table 3.1: *Comparison of different model selection criteria and two large-scale techniques, on toy examples. The training sets for FS2 consisted of* 50 *data points. Results are indicated as MSE±standard deviation and* (σ^2, n_c).

		All points	**FS2**	**ICF [16]**
Semicircles N=500 d=2	**MDD**	0.063 ± 0.07 $(0.8, 33)$	0.04 ± 0.05 $(2.19, 9)$	0.09 ± 0.17 $(1.13, 24)$
	Shannon[217]+MDD	0.155 ± 0.18 $(11.6, 7)$	0.12 ± 0.14 $(11.8, 4)$	$--$ $--$
	kPA[105]	3.04 ± 2.64 $(3.11, 1)$	2.75 ± 2.16 $(12.91, 4)$	0.23 ± 0.28 $(1.13, 6)$
Square N=2000 d=2	**MDD**	0.0033 ± 0.003 $(0.28, 22)$	0.007 ± 0.009 $(0.18, 10)$	0.0025 ± 0.0022 $(0.13, 34)$
	Shannon[217]+MDD	0.0059 ± 0.0052 $(0.66, 12)$	0.011 ± 0.011 $(0.66, 6)$	$--$ $--$
	kPA[105]	$- \pm -$ $(-, 0)$	$- \pm -$ $(-, 0)$	$- \pm -$ $(0.13, 0)$
Rings N=3000 d=3	**MDD**	0.021 ± 0.030 $(1.25, 71)$	0.037 ± 0.041 $(2.51, 13)$	0.051 ± 0.10 $(1.25, 71)$
	Shannon[217]+MDD	0.040 ± 0.055 $(2.51, 45)$	0.037 ± 0.041 $(2.51, 13)$	$--$ $--$
	kPA[105]	$- \pm -$ $(-, 0)$	$- \pm -$ $(-, 0)$	$- \pm -$ $(1.25, 0)$

σ^2 is small. However, in this work its value was defined experimentally and no optimization was performed. The model selection criteria are indicated in Figure 3.9, and the denoised datasets for all the criteria are shown in Figure 3.10. Table 3.1 presents the mean squared error (MSE) measured between the original clean test patterns and the denoised test data. The results obtained using the methodologies of [105] and [217] are also indicated. The algorithm based on the maximization of the Shannon entropy was not tested together with ICF because the kernel parameter was already defined using (3.10).

The other two toy examples with their corresponding denoised sets are indicated in Figure 3.10. The second toy example corresponds to a square of 2000 points in a 2D space, and the third one consists of two rings of 3000 points contained in a three dimensional space. As in the previous experiment, 50 support vectors were selected from each dataset by means of FS2, and $N_\nu = 1000$. For

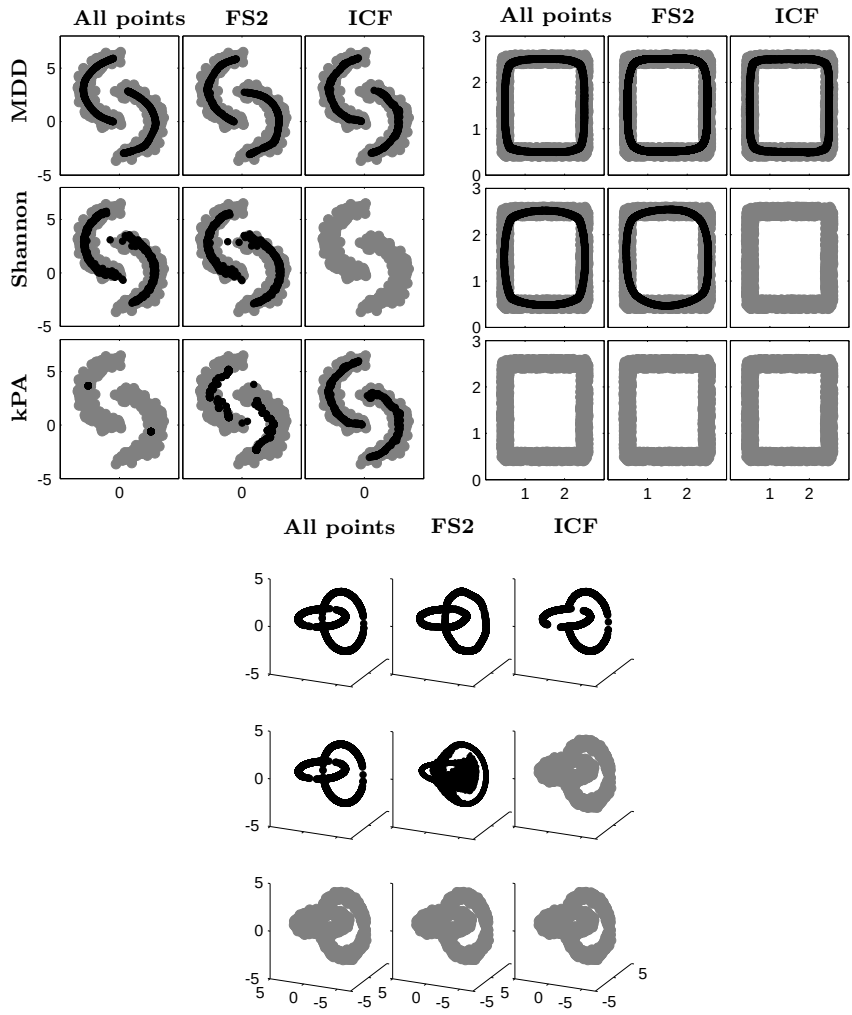

Figure 3.10: *Denoised datasets (black dots) after model selection. The gray dots represent the input dataset. The first three columns correspond to the semicircles, the next three to the square example, and the last three to the rings. For visualization purposes, the original rings are not indicated in MDD and Shannon. Note that there are no results for Shannon entropy with ICF, because the kernel parameter is already defined, and the number of components is determined using MDD. For the square and the rings, kPA was not able to find any component.*

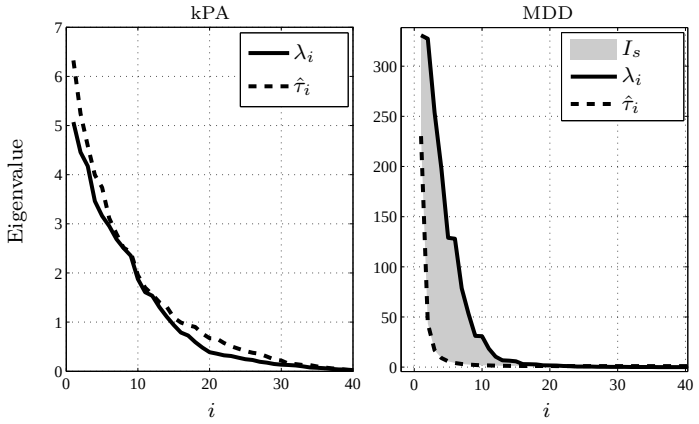

Figure 3.11: *Eigenvalue spectra for the square example using kPA and MDD. Note that no component is selected using kPA because the estimated noise is always larger than the information contained in the data.*

ICF+MDD, the kernel parameter was computed using (3.10), and 98 pivots were selected for the square dataset, and 493 for the 3D-Rings example, with $N_\nu = 1000$. The model parameters and the MSE are shown in Table 3.1. The estimated noise from kPA tends to be overestimated in low dimensions (see Section 3.1), in other words the level of noise can be larger than the level of variance explained by the components. For this reason, this method failed to select any components for the square and 3D-Rings datasets. Figure 3.11 shows the eigenvalues obtained for the square example using kPA and MDD. Note that for kPA the noise estimation is always larger than the information contained in the data, which confirms what was discussed in section 3.1.

From Figure 3.10 and Table 3.1 it is clear that the best denoised results are obtained using MDD, and all the points of the dataset. However, the performance is not compromised when FS2 is used for large scale applications.

The computation times for the different model selection methodologies are indicated in Figure 3.12. For this comparison, the kernel parameter was set to the "optimal" value, the size of the training set was changed and a mean computation time was obtained after running the algorithms 50 times on each dataset. Note that the computation times of ICF+MDD are lower than for FS2+MDD, this is due to the fact that the set of approximated eigenvectors is computed from a set of pivots, which is normally much smaller than the training set. The simulations with FS2+Shannon+MDD are not indicated since the σ^2 is fixed and the algorithm that is used in this study to compute n_c is the same as

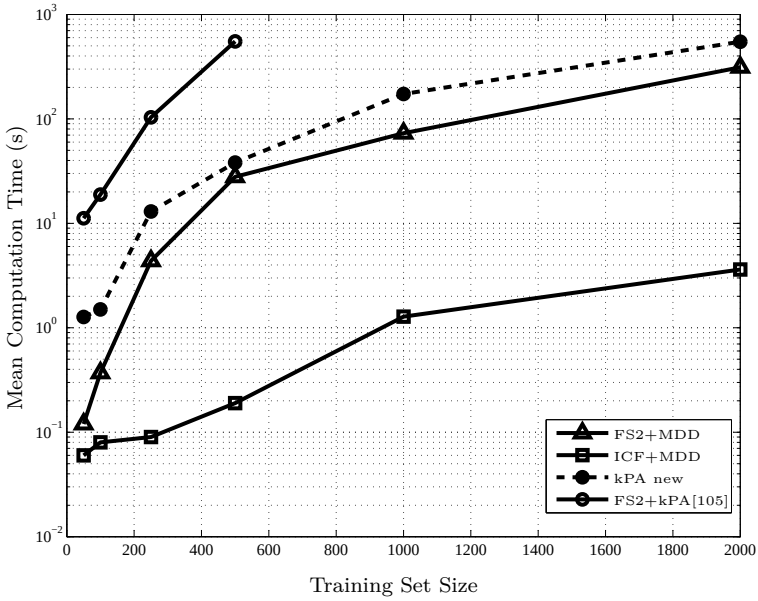

Figure 3.12: *Mean computation time on toy examples. The dashed line indicates the time for the algorithm of kPA implemented in this study, which performs better than the one in [105] (open circles). The main reason why FS2+MDD is much slower than ICF+MDD is the Monte Carlo simulations that are required to compute the noise level.*

for FS2+MDD. In addition, the publicly available code used to implement kPA is unstable when the dataset is larger than 500 points. For this reason, a new implementation was used for these particular cases and it was shown to be faster and more efficient than the public code (see Figure 3.12). The disadvantage of the public implementation is that all the 100 permutations of the input data are stored *all* in the memory, and after this, the kernel matrices and eigendecompositions are computed. In order to make the code more efficient, each permutation was performed separately, and once its kernel matrix and eigenvalue decomposition were computed, the input (permuted) matrix was replaced by a new permutation. The fact that the 100 input matrices are not stored in the memory, makes the algorithm more efficient, and able to handle more than 500 points.

In these toy examples, a certain ratio between the number of training and validation points was used. For instance, for the square example this ratio was set to 50/1000. In order to evaluate how sensitive the proposed algorithm is

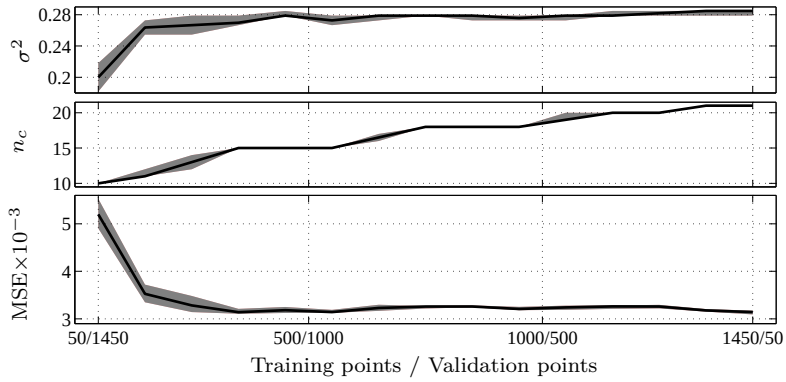

Figure 3.13: *Model parameters and MSE (on test set) for different ratios of training and validation points. The solid black line indicates the median for 50 independent runs, and the shaded area indicates the 25th and 75th percentiles. Note that from 350 training points on, the generalization remains almost constant. The value of n_c closely depends on the size of the kernel matrix.*

to this ratio, the training set was selected using FS2 with SV going from 50 to 1450 points in steps of 100. The validation set was selected at random and it changed from 1450 to 50 data points in steps of 100. For each pair training-validation, 50 runs were performed where the MSE was computed on the test set. Figure 3.13 shows the median and the 25th and 75th percentiles of the MSE values and the selected model parameters for different training sizes. It can be observed that when the model is trained using 350 points or more, the same kernel parameter ($\sigma^2 \approx 0.28$) can be selected, and a similar MSE can be obtained. However, the number of components n_c increases as the size of the kernel matrix increases. Therefore, the ratio between the size of training and validation sets influences the sparsity of the solution. Furthermore, it is remarkable that after selecting the training set and randomizing the validation set several times, no high variations were found. Hence, the sensitivity of the proposed approach to changes in training sizes and changes in validation data, is very low.

3.4.2 Denoising of digits

The second type of application handles the denoising of two datasets, namely the University of California at Irvine (UCI) *multiple features*, and the US Postal

Table 3.2: *Comparison of different model selection criteria and two large-scale techniques, on the UCI dataset, with* $N = 218$ *and* $d = 240$. *The training sets for FS2 consisted of* 180 *data points. Results are indicated as MSE*$\pm$*standard deviation and* (σ^2, n_c).

	All points	**FS2**	**ICF [16]**
MDD	0.08 ± 0.01 $(13.48, 35)$	0.09 ± 0.01 $(13.35, 29)$	0.08 ± 0.05 $(1.8, 77)$
Shannon[23]+MDD	0.10 ± 0.01 $(17.82, 33)$	0.11 ± 0.01 $(19.07, 27)$	$--$ $--$
kPA[7]	0.10 ± 0.01 $(14.93, 16)$	0.11 ± 0.01 $(9.06, 13)$	0.08 ± 0.01 $(1.8, 82)$

Service (USPS) handwritten datasets. The UCI dataset consists of 218 digit images of 15×16 pixels, and the USPS dataset of $N = 9298$ images of 16×16 pixels. In these examples, Gaussian noise was added to the original digits and a subset of 180 images from the UCI set and 2000 from the USPS, were selected as training set, using FS2. The validation sets were of 20 and 1000 images respectively. The SNR for the UCI dataset was 8.22, and 5.51 for USPS. The MSE was measured between the original digits without noise and the denoised images. For the UCI dataset, a comparison between different model selection, and large scale methodologies was performed. These results are presented in Table 3.2, while Figure 3.14 shows the search grid for the model parameters. Note that the solution found by MDD using all the points is indicated by the gray circle, and it is close to the minimum MSE. The performances for the USPS dataset are presented in Table 3.3. In this case, only the results for the large scale methods are presented due to the large size of the dataset. Some denoised test images of the UCI dataset are presented in Figure 3.15, which also shows the denoised test images using the state-of-the-art denoising algorithm proposed by Kong et al., 2013 in [112]. Note that MDD does not lead to a sparse solution when compared with FS2+kPA. In this particular case, permuting the data points as in [105] will result in a good approximation of the noise. Another way of approximating the noise is by generating several random datasets of uniformly distributed points within the limits of the input dimensions. By doing this a slightly better performance compared with kPA can be achieved, and it corresponds to $n_c = 15$, $\sigma^2 = 14.78$ for the UCI dataset with MSE 0.09 ± 0.01, and to $n_c = 33$, $\sigma^2 = 11.72$ for the USPS dataset with MSE 0.03 ± 0.01.

As mentioned in Section 3.2, the four statistical moments of the original distance distribution must be used, because they all contain valuable information

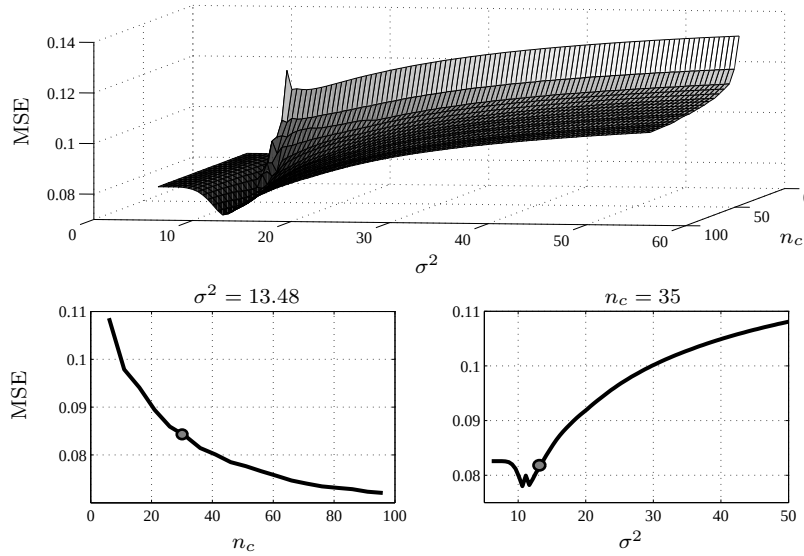

Figure 3.14: *(top) Grid search for the model parameters. The bottom panels indicate the MSE for the parameters identified by MDD using all the points in the UCI dataset. The gray circles indicate the selected parameters. Note that the selected σ^2 produces results with a performance close to the minimum of the curve. However, when looking at the number of components for the given kernel parameter, the best model tends towards non sparse solutions. According to the curve, the best solution corresponds to $\sigma^2 = 12.65$ and $n_c = 91$.*

Table 3.3: *Size of the training set, model parameters indicated as (σ^2, n_c), and MSE±standard deviation on denoised test data and original USPS patterns.*

	USPS ($N = 9298$, $d = 256$)			
	Training	σ^2	n_c	MSE
FS2+MDD	2000	9.88	313	0.03 ± 0.01
ICF[11]+MDD	1361	117.35	272	0.09 ± 0.06
Shannon[217]+MDD	2000	13.26	30	0.04 ± 0.04
FS2+kPA[105]	2000	13.26	28	0.04 ± 0.04
FS2+MDD [a]	2000	11.72	33	0.03 ± 0.01

[a] Results of noise estimation after generating, several times, a dataset of uniformly distributed points within the limits of the input dimensions.

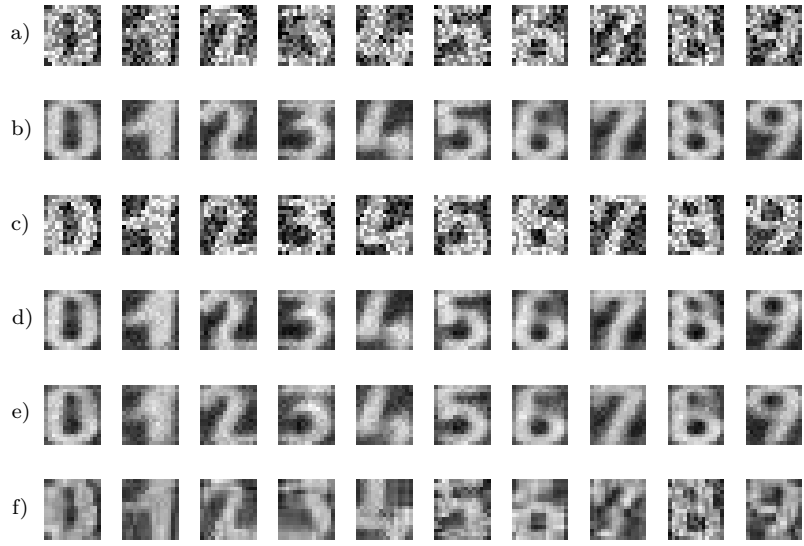

Figure 3.15: *a) Noisy test UCI dataset; b) denoised after FS2+MDD with* $\sigma^2 = 13.35$ *and* $n_c = 29$*; c) ICF+MDD with* $\sigma^2 = 1.8$ *and* $n_c = 77$*; d) FS2+Shannon+MDD with* $\sigma^2 = 19.07$ *and* $n_c = 27$*; e) FS2+kPA with* $\sigma^2 = 9.06$ *and* $n_c = 13$*; f) Kong et al., 2013[112].*

about the structure and dimensionality of the data. For example, if in this particular application only the first two moments of the original distance distribution are used to generate the noise distance matrix, the corresponding eigenvalue distribution for the optimal MDD σ^2 does not allow to select any component. This can be observed in Figure 3.16a(solid gray line), where it is also clear that after the first component, the noise is always bellow the data. If now, instead of looking at the optimal MDD σ^2, the information content I_s is computed for each kernel parameter, it is observed that for some cases ($I_s = 0$) the noise may contain more information than the original data (see Figure 3.16b, solid gray line). Here, the advantage of using kurtosis and skewness becomes clear. For low dimensionality problems, using the first two moments will produce similar results than when using all four. This is because the lower the dimensionality of the data, the closer the artificial distribution is to a normal one.

A different approach to generate the artificial distance matrix is by using the observed distance distribution. Figure 3.16a(dashed black line) indicates the corresponding eigenvalues for the optimal MDD σ^2. Note that the first eigenvalues are larger, and this is due to the fact that more data structure is

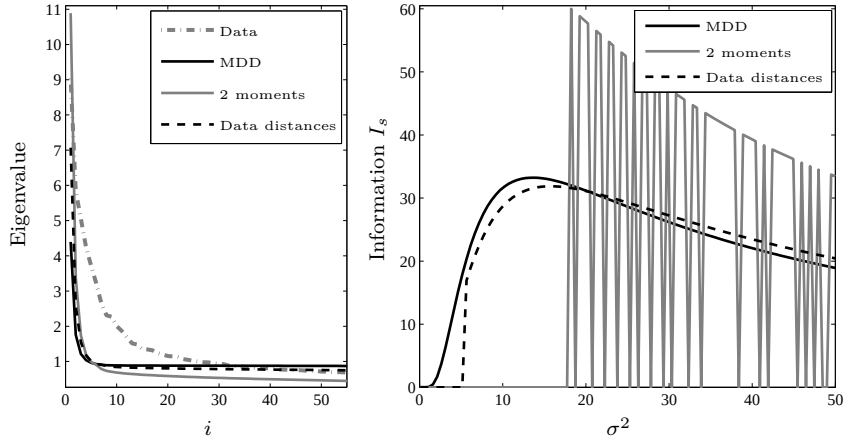

(a) (b)

Figure 3.16: *Eigenvalue distributions (a) and information (b) for different noise estimation approaches. "2 moments" generates the noise matrix using only the first two statistical moments of the real distance distribution. This method does not provide a valid approximation for different values of* σ^2*. "Data distances" indicates the results when the entries of the noise matrix are drawn from the observed distance distribution. The last case underestimates the noise contained in the data.*

still contained in the distribution of distances. Additionally, the noise tends to be underestimated. In this example, 8 more components are needed for data denoising, which causes a drop in performance: 0.11 ± 0.03. If I_s is maximized for this case (see Figure 3.16b, dashed black line), the resulting MSE is 0.1 ± 0.01 with $n_c = 32$ and $\sigma^2 = 13.37$. These differences in performance are not critical, which indicates that creating the noise matrix using the observed distance distribution, can be seen as an alternative keeping in mind that the noise will tend to be underestimated.

Going back to the results obtained with the proposed approach, for the USPS dataset, it is observed that the MSE values obtained with ICF+MDD are larger than for all the other algorithms. This can be due to the fact that the kernel parameter σ^2 is computed using the rule of thumb in (3.10), and that the user defined error threshold η is not optimized. The reasons for this are discussed at the end of Section 3.3.

The performance of the algorithms was evaluated for different SNRs. Figure 3.17 shows the MSE and the standard deviations, for different SNR applied to

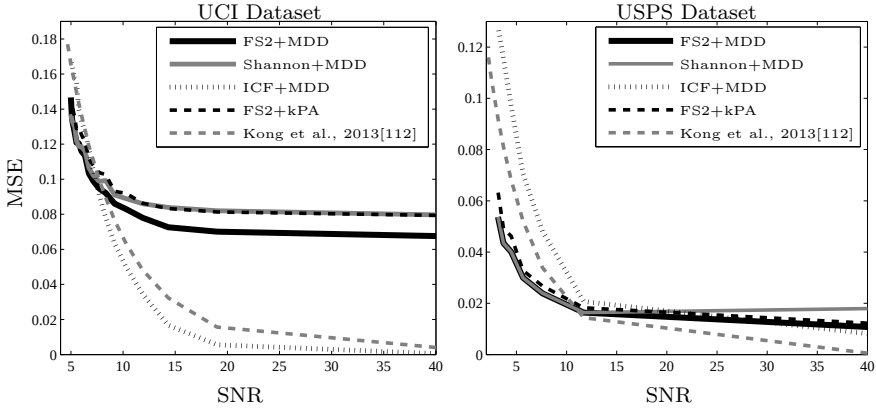

Figure 3.17: *MSE for different signal to noise ratios of the UCI and USPS datasets. Lowest errors are obtained with ICF+MDD and the method in [112]. However, when SNR is lower, these two methods tend to perform worse than the other algorithms. Note that kPA and Shannon indicate low performances even for high dimensional datasets.*

the UCI and USPS datasets. The results for this application were compared with the algorithm in [112]. From the figure, it is clear that MDD and the large scale methods described in this work outperform the other algorithms, and lower errors are obtained when the SNR increases. It is also clear that kPA and Shannon methods indicate lower performances even for high-dimensional data. Additionally, for low SNR, the method in [112] performs similar to ICF+MDD, which rapidly decreases performance as the SNR gets worse. This can also be observed in the denoised test images in Figure 3.15.

3.5 Conclusions

This chapter presented a criterion to select the model parameters σ^2 for RBF kernels and n_c in kernel PCA. The main contribution of this work relies on the estimation of the noise by generating an artificial distance distribution, based on the original dataset. This new distribution allows the noise to be estimated without generating any artificial dataset. It was shown that this method outperforms other criteria for model selection, especially in lower-dimensional spaces, where other methods seem to fail. Based on the findings presented in this chapter, MDD can be used both in low- and high-dimensional spaces, but its advantages are better exploited in low dimensions.

The strength of the methodology proposed here is that it is able to retrieve the model parameters in any dimensional space, and that it can be useful in many applications. One of this applications is the derivation of the respiratory signal from the ECG, which will be discussed in Section 4.4.

A clear drawback of MDD-KPCA is the implementation of a Monte Carlo simulation, which can be seen as a bottleneck when online applications are considered. Future studies may consider to avoid this implementation. Furthermore, the application of MDD in different unsupervised methods like clustering, and in supervised techniques like regression might be explored.

Chapter 4

ECG pre-processing and Analysis

This chapter describes the steps in the pre-processing of the ECG signals, and it introduces the basis for the construction of the tachogram and the analysis of the heart rate variability. In addition, it presents some of the contributions of this research, such as a novel algorithm for the detection of ECG artifacts based on the autocorrelation function, and an adaptation to the well-known Pan-Tompkins algorithm for the detection of the QRS complexes. The performance of each of these algorithms is evaluated and presented in this chapter. Furthermore, the methodology used for the construction of the tachogram is described. After that, the basic measures used for the quantification of the heart rate variability are discussed. Finally, three different approaches to derive respiratory information from the ECG are introduced and compared on a public ECG dataset.

4.1 ECG Artifacts

The diagnostic capabilities of the electrocardiogram (ECG) can be reduced by the presence of artifacts in the signal. These contaminants can be caused by electrode motion, contact noise and electromyography among others [51]. Due to these undesired features, the output of different algorithms can be corrupted. For example when classifying segments of the signal, or deriving respiration from the ECG, an artifact can be interpreted as an anomaly. Therefore, in order to obtain accurate conclusions from any ECG analysis, it is very important to filter and eliminate any type of erroneous sections [73].

In principle this can easily be done when short signals are analyzed, because the length of the measurements allows for manual inspection of the noise. However, in practice, long-term recordings of several patients are studied at the same time, which requires automated algorithms to detect and remove artifacts.

Different methodologies have been proposed to enhance the quality of the ECG, such as [26, 165, 178]. A disadvantage of those methods, however, is that they do not provide information of the location of the corrupted sections. Therefore, even though the quality of the signal is enhanced, some artifacts are still present, and can have a negative influence on the analysis. An approach to find the specific location of disturbances is presented in [36], where the signal is evaluated with a moving window. One shortcoming of this method is the use of rules like the mean of many points in the past, or the standard deviation of the signal, to identify contaminants. Furthermore, the method assumes that no abrupt alterations occur, which poses a problem, since artifacts may as well cause abrupt changes in the ECG.

In this section these problems are addressed. First, the different sources of noise and artifacts in the ECG will be described. Next, two methods to estimate the noise level in the ECG signal will be proposed, and later they will be evaluated on a real dataset.

4.1.1 Sources of noise and artifacts

Typical examples of noise and artifacts contaminating the ECG, can be classified as follows [73]:

- *Power line interference* is noise at 50Hz for the datasets collected in Leuven and 60Hz for datasets collected in the US. Its amplitude can go up to 50% of the peak-to-peak amplitude of the ECG signal.

- *Electrode contact noise* can be caused by a loss contact between the electrodes and the skin of the patient. These can occur intermittently and they are measured as drastic changes in the ECG signal with amplitudes larger than the peak-to-peak ECG amplitude.
- *Motion artifacts* caused by movement of the patient or motion of the electrode with respect to the skin that results in impedance changes. These changes can be observed as baseline variations in the ECG.
- *Electromyogram (EMG) artifacts* are caused by muscle contractions, and appear as zero-mean Gaussian noise.
- Baseline drift is a low frequency noise that is mainly caused by respiratory movements, or perspiration, which cause impedance changes.
- *Instrumentation noise* that can be added by the hardware recording the ECG.
- *Electrosurgical noise* is added by other equipment present near the ECG monitor device.

Figure 4.1 shows four of these typical ECG contaminants. It is clear that some artifacts can be removed after applying the correct filtering technique. For instance, the power line interference in Figure 4.1a can be removed by means of a band-stop filter, and the baseline drift in Figure 4.1c can be corrected using a high-pass filter. However, a different approach needs to be used for artifacts that affect the ECG morphology beyond repair. This is the case illustrated in Figure 4.1b, where it is obvious that the only option left is to detect those segments and remove them from further analysis. A more difficult case is indicated in Figure 4.1d. Note that the artifact destroyed the morphology of the P- and T-waves, but one can still find the location of the R-peaks. Therefore, segments like this one could be useful for heart rate variability analysis, but not for the detection of T-wave alternans [164].

Provided that for some applications, the elimination of "all" ECG contaminants is required, while for other cases some noise can be allowed, it is necessary to estimate the contamination level of the ECG. This problem will be described next.

4.1.2 Noise level estimation

The type of ECG application determines the level of contamination that is allowed in the signal (i.e. signal-to-noise ratio or SNR). For instance, in

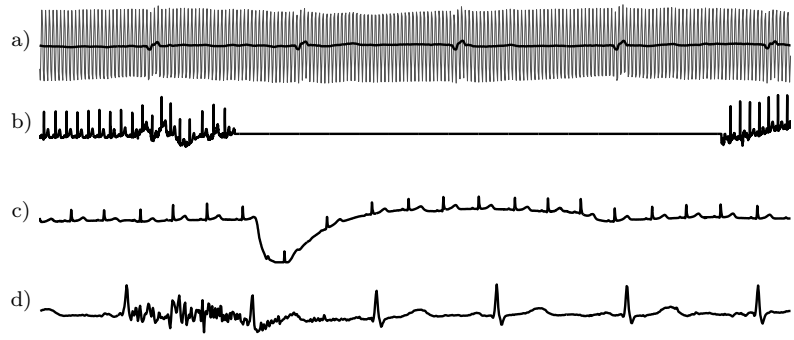

Figure 4.1: *ECG segments contaminated by artifacts. a) In gray the original ECG contaminated by power line interference at 50Hz is indicated, and in black its filtered version. b) Contact noise. c) electrode motion. d) EMG artifact.*

the detection of T-wave alternans [164] it is crucial to clearly identify and segment the T-waves. Therefore, baseline drifts, power line interference, motion and EMG artifacts must be removed from the analysis. This is because any alteration of the T-wave at the level of micro-volts can significantly affect the results. In this application it is necessary to achieve a low level of contamination in the ECG. Conversely, in heart rate variability, this level can be higher due to the fact that the only characteristic wave that needs to be detected is the R-peak. The R-peak is the "easiest-to-detect" point in the ECG due to its large amplitude with respect to the baseline. For this reason, low amplitude artifacts, like the one observed in Figure 4.1d, that do not mask the position of the QRS complex, can be included in the analysis.

In order to estimate the contamination level in the ECG, measures such as the RMS power in an isoelectric region, or the ratio between the R-peak amplitude and the noise amplitude in the isoelectric region can be used [51]. However, these measures make some assumptions of the properties (e.g. normally distributed) of the noise that are not always satisfied. Other ways of estimating the SNR of the ECG use PCA [140], ICA [16], cubic-splines & state-space formulations [136], and time-frequency representations [13, 177, 184]. Even though all these methods achieve good performances in removing baseline drifts, estimating in-band noise, and recovering the morphology of the ECG, they have clear disadvantages. In PCA and ICA for example, the selection of the number of components characterizing the noise is completely arbitrary. In addition, these methods cannot deal with changes in the dynamics of the noise [166]. On the other hand, the performance of methods estimating baseline noise is

significantly reduced when other artifacts, like those introduced by movement are included [74].

In this work, two different methods to estimate the level of noise in the ECG are proposed. The first one is based on the autocorrelation function (ACF) of the ECG, and detects ECG segments that are highly corrupted by noise or artefacts. The idea behind it, is the characterization of ECG segments by means of the ACF, and the use of their similarity to distinguish between them. This works because the ACF shows repetitive patterns, and hence also disturbances, i.e. artefacts, in the ECG. The ACF has been used in previous studies using ECG signals [116, 161, 190, 195], but never to locate artefacts. The second method uses the model selection algorithm MDD, described in Section 3.2.

The advantage of these two methodologies is that a level of noise is assigned to each ECG segment, which allows to identify the highly contaminated segments and those with an "acceptable"[1] level of noise. With this is mind, it is possible to determine how relevant each ECG segment is for the final results. In other words, the level of contamination can be used as a confidence level, or as a weight when training, validating or testing a model. This will become clear later in this chapter. First the two proposed methodologies will be discussed.

Autocorrelation function (ACF)

Each ECG recording is first divided into segments of length N. Then, the low and high frequency components are attenuated, by means of a band pass Butterworth filter with cut-off frequencies of 1Hz and 40Hz. The frequency bands are selected based on the low- and high-frequency limits of a diagnostic ECG [51] and on the fact that variations due to baseline drifts should not be considered.

After this preprocessing step, the ACF of each ECG segment is obtained by computing the inverse Fourier transform of its corresponding power spectral density (PSD). Note that the Fourier transform of a non-stationary signal is used, which normally should be avoided because of low frequency resolution. In other words, the Fourier transform should not be used to study e.g. beat abnormalities or rhythm changes [51]. However, the goal at this point is to detect artifacts of the signal, rather than deviations from healthy ECGs. Therefore, in this particular application, the resolution of the Fourier transform is sufficient. This can be observed in Figure 4.2, where the power spectra of two ECG segments are compared. The figure shows that a contaminated segment contains a broader spectrum when compared with a clean one. This behavior

[1] Acceptable is related to the application and it is typically indicated by the user

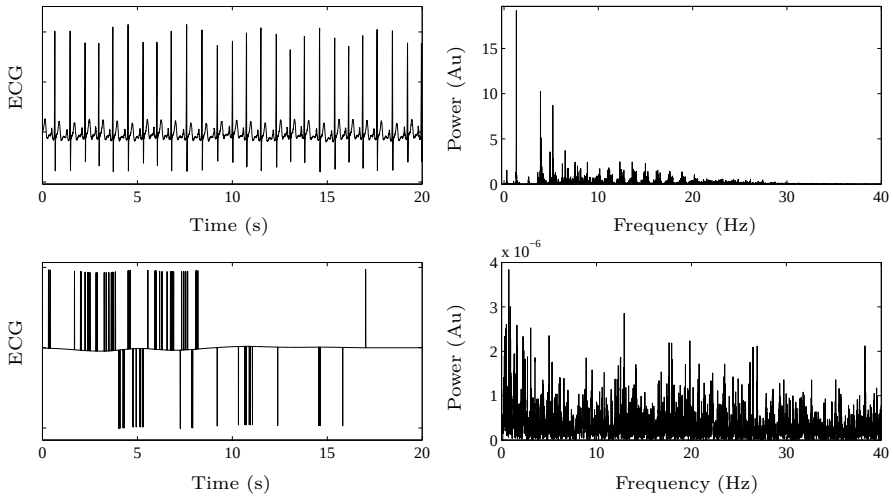

Figure 4.2: *ECG segments (left) and their corresponding power spectra (right). Note that the frequency content of an artifact is broader (bottom) compared to the one of a clean signal (top).*

is expected as the artifact should behave more like a random sequence with a broader frequency content. A similar assumption is made in [115], where motion artifacts are removed from photoplethysmographic data, using time and frequency analysis.

Once the ACFs are computed, the contaminated segments need to be localized. Figure 4.3 shows some examples of one-minute ECG segments and their corresponding ACF. Note that when an ECG is corrupted by an artifact (column 3 and 4), a clear change in its autocorrelation is observed. This means that the ACFs that contain disturbances are "dissimilar" from the noise-free segments. To find these differences, the following procedure is proposed. First, the pairwise cosine similarity between the ACFs is computed as

$$\text{Similarity} = \cos\theta_{1,2} = \frac{A_1^{\mathrm{T}} A_2}{\parallel A_1 \parallel \parallel A_2 \parallel}, \tag{4.1}$$

where A_1 and A_2 correspond to the ACFs of segments 1 and 2 respectively, which are two vectors with an angle $\theta_{1,2}$ between them, $\parallel . \parallel$ defines the two-norm and $A_1^{\mathrm{T}} A_2$ the dot product between the vectors. From all the pairwise similarities, a similarity matrix K can be computed such that the entries $k_{ij} = 1$ for $i = j$, $k_{ij} = k_{ji}$, and $k_{ij} = \cos\theta_{i,j}$ for $i \neq j$. The next step is to compute the *degree* of each ACF. This was discussed in Section 2.2.3, where the degree

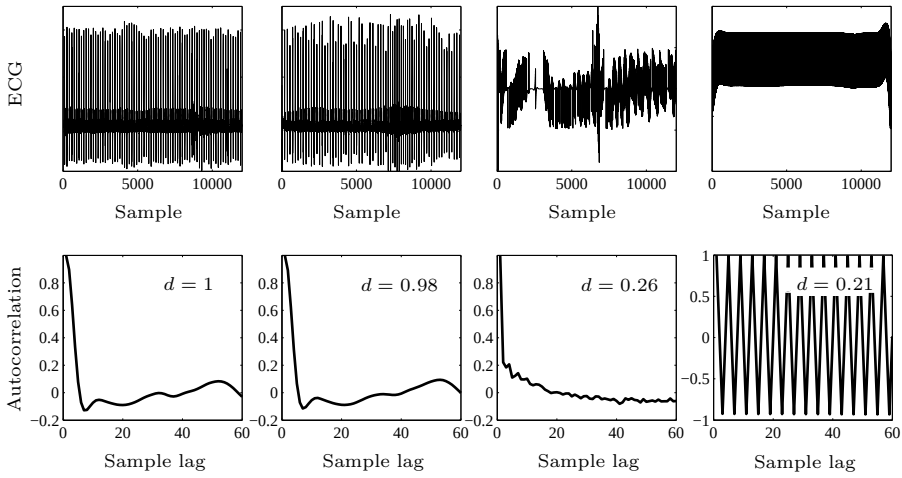

Figure 4.3: *Autocorrelation functions (ACFs) of four different ECG segments of one minute. The first two columns correspond to normal minutes, while the other two to minutes with artifacts. The normalized degree values d, for the ACFs are indicated. Note that the similarity between the contaminated segments and the rest is significantly low compared to the ones of normal minutes.*

was defined as

$$d_i = \sum_{j=1}^{N} k_{ji}, \tag{4.2}$$

which is equal to the sum of the columns (or rows) of K. In this way, segments that are significantly different, i.e. artifacts, are represented by low degree values. When ECG segments are clean, their ACFs are similar, and they get high degree values. This can be seen in Figure 4.3, where two clean (left) and two extremely contaminated minutes (right) are plotted together with their ACFs and their degree values. This work proposes to use these degrees as "weights" to indicate the contamination level of the ECG segments. Figure 4.4 shows all ACFs of a long-term ECG signal segmented on a minute-by-minute basis. The figure also indicates the normalized degrees for all the ACFs. The procedure described above is summarized in Algorithm 4 (steps 1-4), and steps 5 and 6 are proposed to automatically discriminate between "clean" and contaminated segments. This will be clarified further in this section.

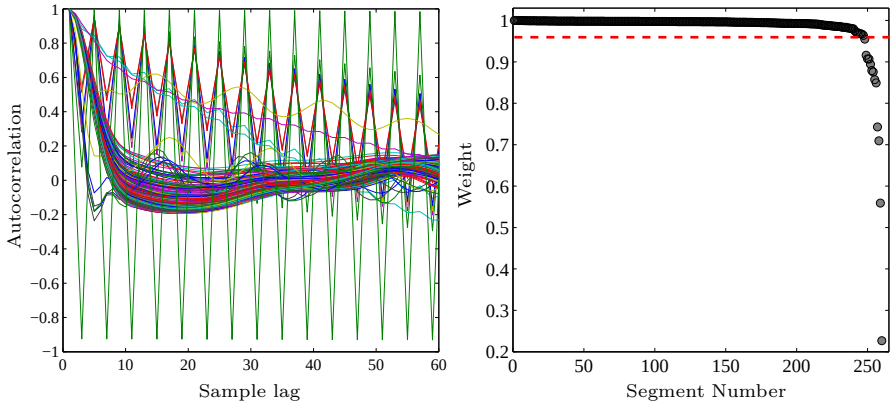

Figure 4.4: *Autocorrelation functions (left) of a long term ECG signal segmented on a minute-by-minute basis, and the degrees or weight values (right) for each ECG segment computed using step 4 of Algorithm 4. The dashed red line indicates the threshold of 95%.*

Algorithm 4. Artifact detection algorithm using the ACF

Input: ECG signal and length N. In this work $N = 60s$.

Output: Weight for each ECG segment and 95th percentile of the weights.

1. Segment the ECG into epochs of length N.
2. Compute the ACF of each segment using the inverse Fourier transform of its corresponding PSD
3. Calculate the cosine similarity (4.1) between each pair of ACFs. This results in the similarity matrix K between all ACFs.
4. Compute the degree values using (4.2) or by summing all the columns of the similarity matrix K. These degrees correspond to the weights, which can be normalized in order to obtain a value between 0 and 1.

Optional

5. Calculate the 95th percentile of the weights.
6. Label the segments with weights outside the 95th percentile, as contaminated.

MDD

In Section 3.2, the model selection criterion for KPCA denoising, based on distance distributions (MDD), was described. Here, this criterion is used to assign a value of information content, I_s, to each ECG segment. First, the preprocessing step described before is applied, namely, the ECG is segmented into epochs of length N, and then filtered using a band-pass filter with cut-off frequencies of 1Hz and 40Hz. Next, the eigenvalue distributions of the data and

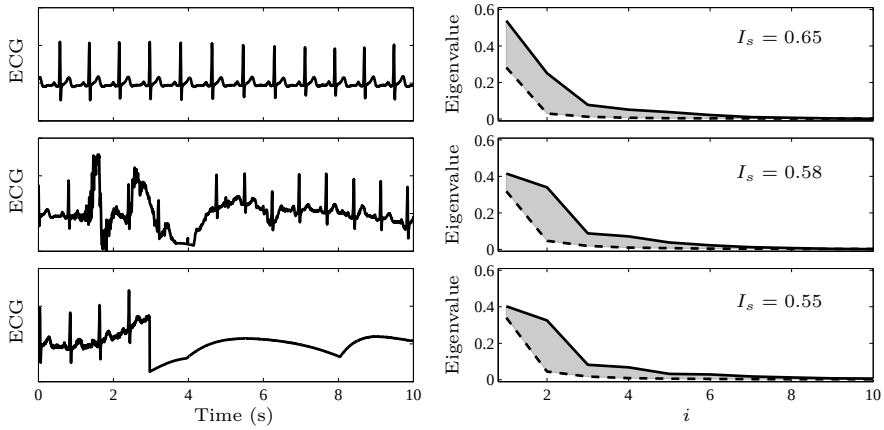

Figure 4.5: *ECG segments (left) of 10s and eigenvalue distributions (right) of the data (solid line) and noise (dashed line). The shaded area between the eigenvalue distributions indicates the amount of information I_s contained in the data. Even though I_s is lower for contaminated ECG segments, it is not a very sensitive measure to discriminate artifacts.*

the noise are computed and the value of I_s is obtained. These distributions are computed using the rule of thumb for the kernel parameter, which in this case is the variance of the ECG signal in each segment. Figure 4.5 shows three different ECG segments with their corresponding information content. Note that the values of I_s decrease as the contamination level increases. With this in mind, MDD can be used to identify the contamination level of the ECG, in the same way that the degrees of the ACF were used. However, from Figure 4.5 it is clear that I_s is not a sensitive measure to distinguish between normal and contaminated segments. It is worth noting that in order to compute the degrees, the similarity between all ACFs must be calculated, while I_s can be obtained for each segment independently. This means that the first approach based on ACF requires long-term ECG signals, with high SNR in order to determine the contamination level. With MDD, however, this is no longer needed, because the estimation of noise is computed for each segment separately. Nevertheless, if an approach based on percentiles is implemented, a longer ECG signal is again required. Algorithm 5 summarizes this approach.

Apart from offering a low sensitivity towards artifacts, MDD requires the computation of a kernel matrix, the size of which is determined by the length of the ECG signal. For instance, if segments of 1 minute are used with a sampling frequency of 200Hz, then a kernel matrix of 12000×12000 must be computed. As a result, the eigenvalue decomposition of such a matrix becomes unfeasible,

especially when low power, and low computational resources are available (e.g. in home monitoring systems). Hence, it is necessary to reduce the length of the signal (e.g. 10s), or characterize the ECG using a more compact representation like the ACF.

Algorithm 5. Artifact detection algorithm using MDD

Input: ECG signal and length N. In this work $N = 60s$.
Output: Amount of information I_s contained in each ECG segment, and 95th percentile of all values of I_s.

1. Segment the ECG into epochs of length N.
2. Compute the ACF of each segment using the inverse Fourier transform of its corresponding PSD
3. Compute the kernel parameter σ_s^2 as the variance of the ACF for each segment.
4. Compute the pairwise distances between all the points in the ACF, and the distance distribution $\hat{f}_{\mathcal{D}}(r)$. Obtain its statistical moments, and define the distance range $[0, r_{max}]$.
5. Generate a distribution $\hat{f}_{\mathcal{N}}(r)$ from the information computed in 4, and construct the distance matrix $\hat{R} \in \mathbb{R}^{N_t \times N_t}$ satisfying (3.4).
6. Compute the kernel matrix $\Omega_{\mathcal{N}} \in \mathbb{R}^{N_t \times N_t}$, and the eigenvalues $\breve{\tau}_l$, $l = 1, \ldots, M$ using (3.7).
7. Repeat steps 5 and 6, hundred times and define the noise level $\hat{\tau}_l$ as the 95th percentile of the distribution of each eigenvalue $\breve{\tau}_l$.
8. For σ_s^2 compute the kernel matrix Ω, and the eigenvalues λ and eigenvectors α of (2.23).
9. Find the number of components $n_c(\sigma_s^2)$ from (3.8) and the amount of information $I_s(\sigma_s^2)$ contained in n_c from (3.9).

Optional

10. Repeat steps 2 to 9 for all ECG segments.
11. Calculate the 95th percentile of the weights.
12. Label the segments with weights outside the 95th percentile, as contaminated.

4.1.3 Artifact detection algorithms

Considering the two methods (ACF and MDD) described above, one way to remove highly contaminated ECG segments is by retaining the 95th percent of the most similar segments, which in the case of ACF are the vertices with the higher weights (see Figure 4.4), and in MDD are the segments with the largest I_s. In percentile analysis, the 95 percentile is an often used value to detect outliers. Algorithms 4 and 5 summarize the proposed approaches to detect artifacts in the ECG using ACF and MDD respectively, and the 95th percentile threshold. Note, however, that this threshold can vary depending on the total SNR of the signal. For example, when the signal is completely free of

artifacts, 5% of "clean" segments will be removed or labeled as contaminated, and when the SNR is very low, contaminated segments will be considered clean. To overcome this, a clustering technique can be applied on the weights, or on the values of I_s. Here, the kernel spectral clustering approach, described in Section 2.2.3, and the well-known k-means clustering algorithm are used. These two approaches are used to split the weights and I_s into clusters, and their performances are compared with the one obtained with the thresholding approach.

Data

In order to evaluate the two proposed approaches, and to compare their performance with the ones obtained with the clustering techniques, a real dataset is used. It consists of 16 single lead ECG recordings extracted from polysomnographic measures of 16 different patients of the sleep laboratory of the University Hospital Leuven (UZ Leuven), Belgium. The signals were sampled at 200Hz and the recording time ranged from 4 to 11 hours. In total 152 hours and 12 minutes were extracted. Each recording was manually annotated by a medical doctor with experience on interpreting polysomnographic signals, including ECG. The total number of minutes in the dataset is 9132 and only 294 of them contain artifacts, which were manually labeled and verified by a medical doctor.

Results

The following approaches are evaluated: ACF + 95th percentile, ACF + k-means, ACF + KSC, MDD + 95th percentile, MDD + k-means, and MDD + KSC. Due to inter-subject variability in the ECG and ACF, the algorithms are applied to each ECG signal separately. To compare the performance of the algorithms, statistical measures often used in classification known as sensitivity, specificity, and accuracy are computed. These values are obtained for each algorithm and for each patient separately, and the mean values are indicated in Table 4.1. Note that ACF + 95th percentile outperforms all the other tested methodologies, because accuracy, specificity and sensitivity are in general higher for all patients. The possible reasons for this will be discussed later in this subsection.

For completeness, all the one-minute segments are put together and the algorithms are applied to all the signals at the same time. The best results are again produced by ACF + 95th percentile threshold and correspond to a sensitivity of 93.2% and specificity of 90.4%.

Table 4.1: *Comparison of the performance of the proposed algorithms using the 95th percentile threshold, k-means and KSC. Mean values and standard deviations are indicated.*

Method		Sensitivity	Specificity	Accuracy
ACF	95th percentile	**0.96 ± 0.05**	**0.90 ± 0.03**	**0.90 ± 0.03**
	k-means	0.89 ± 0.13	0.82 ± 0.13	0.82 ± 0.13
	KSC	0.92 ± 0.11	0.75 ± 0.15	0.75 ± 0.14
MDD	95th percentile	0.95 ± 0.08	0.34 ± 0.28	0.93 ± 0.02
	k-means	0.75 ± 0.21	0.54 ± 0.14	0.54 ± 0.14
	KSC	0.56 ± 0.30	0.55 ± 0.27	0.56 ± 0.29

To determine if the duration of the segments affects the results, the experiments are repeated with N ranging from 3 to 60 seconds. A better differentiation of the artifacts is obtained after analyzing segments of 3 seconds, the lower bound of the epochs. This is no surprise, since as shorter segments are analyzed, less information from normal ECG is included in the contaminated signals. In other words, when the selected epochs contain, proportionally to the window length, more artifact than normal ECG, its ACF differs more from that of a normal signal. Figure 4.6 shows a one minute ECG segmented into epochs of 5 seconds. These results indicate that when the window is reduced, the resolution of the algorithm is improved. The ACFs of the epochs are also indicated in the figure, where it is possible to observe the differences between corrupted and clean signals. These findings can be used in online monitoring systems to discard short periods of contaminated segments with a delay of 5 seconds. Nevertheless, it is important to keep in mind that the lower the window length, the more sensitive the algorithm is to small changes in the signal. This can be seen in Figure 4.7, where the performance of Algorithm 4 on one ECG signal, and using different window lengths (i.e. different values of N) is indicated. These results were obtained using a threshold of 95th percentile, but it is clear that even though all artifacts are detected, more "clean" segments are labeled as contaminated. To overcome this limitation a different threshold must be defined according to the window length.

There are reasons why the clustering algorithms do not perform as good as the 95th percentile threshold. On the one hand, this is a highly unbalanced problem, since in most of the cases, the artifacts constitute a lower percentage of the signal. In other words, the number of segments in the "clean" group must be

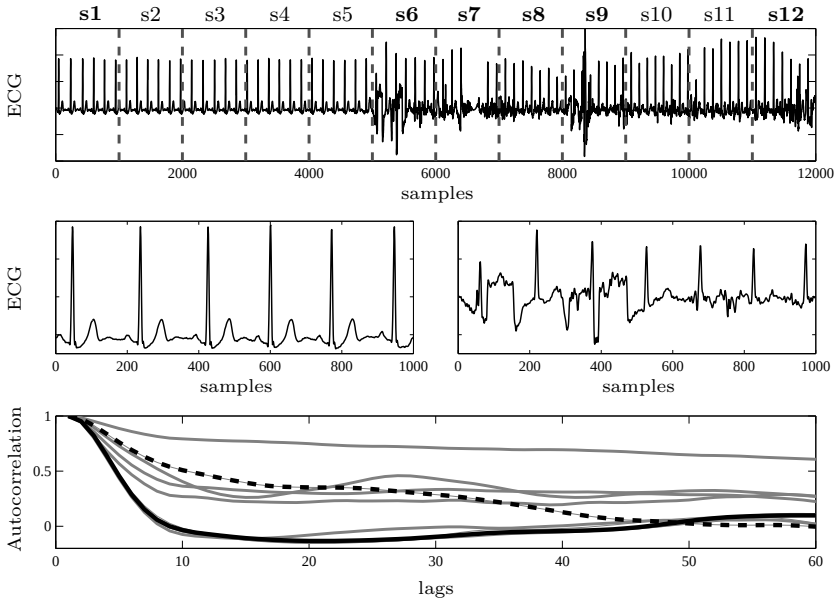

Figure 4.6: *One minute ECG segmented into epochs of 5 seconds (top). Note that segments S6, S7, S8, S9 and S12 are contaminated. Segments S1 (middle left) and S6 (middle right) are indicated. Note that the ACFs (bottom) are significantly different between S1 (tick solid black line) and S6 (dashed line). The other curves correspond to the ACFs of the other segments.*

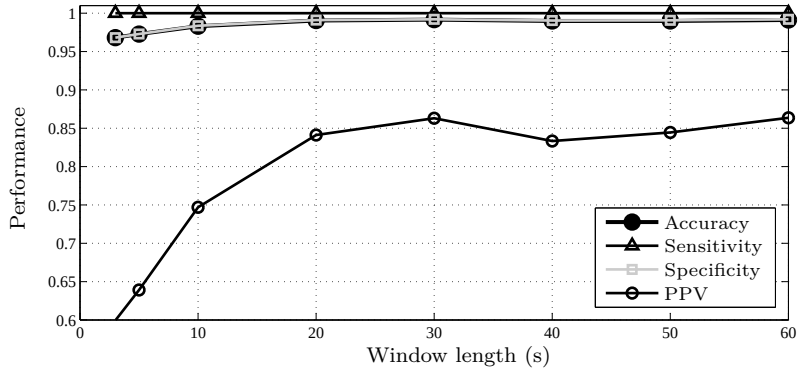

Figure 4.7: *Performance of Algorithm 4 applied to one ECG signal and for different window lengths.*

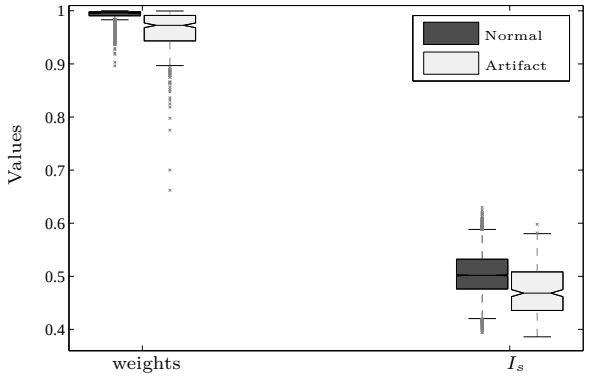

Figure 4.8: *Weights and information content I_s for all the 9132 ECG segments. The weight values of the segments containing artifacts are significantly lower ($p < 0.05$) than the ones of normal segments. This difference is not observed for I_s.*

significantly higher than the amount of segments in the "contaminated" group. KSC uses a model selection criterion that takes these unbalanced problems into account [10]. However, special attention must be paid to the definition of the parameter that determines how balanced the clusters must be. A bad selection of this parameter can produce misleading results.

On the other hand, the k-means clustering finds globular[2] structures in the data. When such structures are not present or are not clearly separable, this technique tends to perform poorly. In addition, this algorithm tends to converge to local minima, which is one of the well known disadvantages of k-means.

With regard to the use of MDD for artifact detection, several disadvantages were observed. First, it is clear from Table 4.1 that all clustering algorithms performed worse for MDD than for ACF. This can be explained by the fact that more variability is obtained in the values of I_s, which reduces the possibility to correctly discriminate the contaminated segments from the clean ones. To demonstrate this, all the segments are taken together and the values for the weights and I_s are compared between clean segments, and the ones contaminated by artifacts. This comparison is shown in Figure 4.8, where it is clear that a better distinction can be obtained using the weights of the ACF.

The variability observed in I_s may be explained by the fact that no ECG reference is used, and no comparison between previous segments in the signals is needed to compute the values of I_s, which is the case to compute the weights.

[2]Globular refers to a spherical shape

Moreover, the kernel parameter used to estimate the noise in each segment was computed independently, which can pose a problem when two different kernels are compared. In other words, two kernels computed using different kernel parameters are normally not describing the same feature space. Therefore, comparisons become unreliable. For all these reasons it is possible to conclude that MDD is not useful for artifact detection in the ECG.

On the other hand, ACF + 95th percentile using $N = 60s$ allows to find segments contaminated by artifacts with a higher accuracy. It is easy to implement and interpret, due to the fact that it uses well-known tools such as the autocorrelation function and percentile analysis. Furthermore, the algorithm does not require computationally expensive operations, e.g. eigenvalue decomposition. Nevertheless, in order for the algorithm to be implemented online, a training phase is required, where the 95th percentile is defined. This can be avoided if a dynamic threshold is implemented.

Another advantage of the algorithm is that it allows to locate the contaminated segments in the ECG. This information is not provided by algorithms that enhance the overall quality of the signal. In addition, other artifact detection algorithms require more ECG channels or extra information, to identify and remove artifacts. This is not the case for this methodology, since it only requires a long term ECG signal.

This algorithm is implemented in the graphical user interface described in Appendix B, which also allows for manual change of the hard threshold used in this comparison, and it offers other functionalities.

4.2 QRS detection and the construction of the tachogram

A crucial step in the study of the ECG, especially when analyzing heart rate variability, is the location of the QRS complexes. These complexes are the most prominent waves in the ECG, and they contain an enormous amount of information about the state of the heart. For this reason, the detection of the QRS complex constitutes the basis for a large number of ECG algorithms [111].

During the last two decades, several algorithms have been proposed to first enhance and then detect the QRS complexes [61]. For instance, techniques like wavelet decomposition [180, 191], Hilbert transform [12], empirical mode decomposition [192] and filter banks [4] have been used to enhance the QRS complexes. Once the quality of the ECG allows for a better delineation of the QRS complexes, algorithms based on threshold definitions [12, 71, 148, 151] are

regularly used [61]. One of these algorithms is the well-known Pan-Tompkins [151], which has been proven to outperform many more elaborated algorithms [61]. In this research, the Pan-Tompkins algorithm is used together with a novel approach to flatten the ECG and enhance the QRS complexes.

4.2.1 Flattening the ECG

After filtering out power line interferences and other artifacts from the ECG, this research proposes to compute the upper (U_{ecg}) and lower (L_{ecg}) ECG envelopes (see Figure 4.9) by means of the secant method. This method finds the segment with the largest slope, within a window of length t. Figure 4.9 illustrates this method, and Algorithm 6 describes step-by-step the computation of U_{ecg} and L_{ecg}.

Once U_{ecg} and L_{ecg} are obtained, they are used to obtain a flattened version of the ECG defined as: $F_{ecg} = U_{ecg} - L_{ecg}$. Figure 4.10 shows an ECG segment with its upper and lower envelope using t =150ms. The flattened ECG is then used to find the possible locations of the QRS complexes by means of the Pan-Tompkins algorithm. The value of t was defined experimentally as the value that offers the best performance of the R-peak detection algorithm for several ECG signals.

Algorithm 6. Computation of lower and upper ECG envelopes

Input: ECG signal and length t.
Output: Upper U_{ecg} and lower L_{ecg} ECG envelopes.

1. Compute the slopes m of all lines connecting the initial point with all points in a window of length t.
2. For U_{eng} and L_{ecg} select the line with the largest m and lowest m respectively.
3. Set the selected line as a segment of the corresponding U_{eng} or L_{ecg} envelope.
4. Define as the initial point, the last point connected by the selected line.
5. Repeat steps 1 to 7 until the end of the signal is reached.

4.2.2 The Pan-Tompkins algorithm

The original Pan-Tompkins algorithm [151] first uses a band-pass filter with cutoff frequencies of 5Hz and 15Hz in order to maximize the QRS energy. Next, the slope information in the signal is captured by calculating the derivative of the signal, the result is then squared to eliminate all negative values. After that, a moving window integrator of 150ms is used to smooth the resultant filtered ECG. Finally, this result is used to find the position of the R-peaks by means of adaptive thresholding [151].

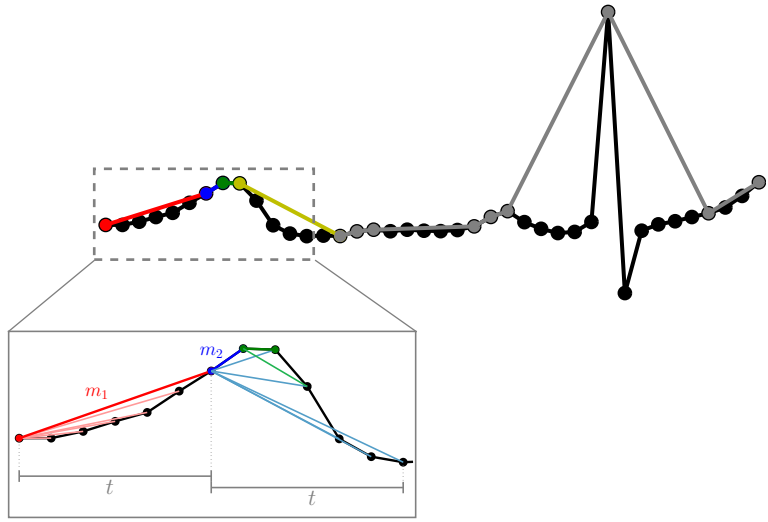

Figure 4.9: *Computation of U_{ecg} using the secant method. In red, the secant curves in a window of length t are indicated. The curve with the largest slope m_1 is selected as the first segment of the envelope. The lines indicated in blue are the new segment candidates. Note that the largest envelope m_2 corresponds to the line connecting the first two points of the curve. Therefore, the new secant curves (green) will be drawn from the second point connected by the line with slope m_2.*

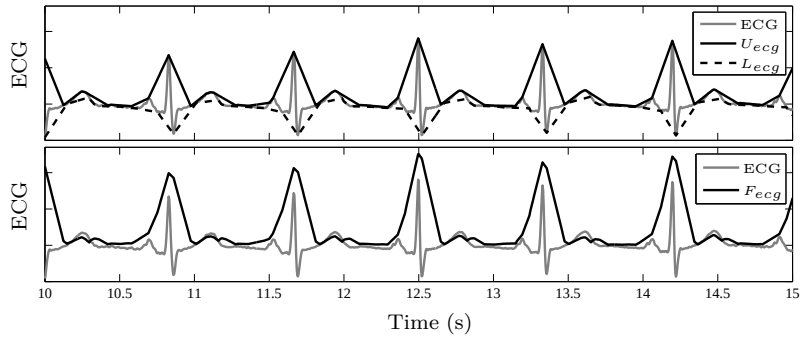

Figure 4.10: *Computation of the flattened ECG signal. (top) upper (U_{ecg}) and lower (L_{ecg}) envelope computed using the secant method on 150ms. (bottom) Resultant flattened ECG ($F_{ecg} = U_{ecg} - L_{ecg}$) that is used to search for the R-peaks. Remember that the locations found in F_{ecg} are used to determine the final R-peak positions in the original ECG signal.*

In this research, the filtering, derivative, squaring and integration steps are replaced by the flattening methodology described before. Therefore, the adaptive thresholding step of the Pan-tompkins algorithm is applied on F_{ecg}. Once the positions of the R-peaks are identified, the original ECG is used to find the final location of the QRS complexes.

Remember that in order to avoid errors in the estimation of the location of the R-peaks, the sampling frequency of the ECG must satisfy the general requirements described in [129]. It is known that optimal values for the sampling frequency should be between 250Hz and 500Hz or even higher. All experiments performed in this research[3] satisfy this requirement, and when the sampling frequency was below the lower limit of the recommended range, the signals were up-sampled.

4.2.3 Construction of the tachogram

The tachogram, or RR-interval time series [129] is a key tool in the analysis of heart rate variability, and in the study of the ECG in general. It consists of the time differences between successive R-peaks, and it exhibits changes in the regulation of the cardiac control mechanism of the autonomic nervous system [51]. This mechanism acts on the sinoatrial (SA) node [84] by accelerating and decelerating the heart. Hence, changes in the RR-intervals are strongly related to changes in the modulation of the autonomic nervous system. For this reason, all the R-peak positions used in the construction of the tachogram must correspond to heart beats originated from the SA node (i.e. normal beats). As mentioned in Chapter 1, other heart beats can be included in the analysis if no special care is taken to remove them from the RR-interval time series. Examples of these undesired beats are arrhythmic episodes and ectopic beats. Furthermore, erroneous peaks introduced by the R-peak detection algorithm, and missing peaks can also lead to the wrong conclusions. Therefore, they must be identified and corrected.

Despite the fact that the Pan-Tompkins algorithm regularly produces satisfactory results in the detection of the QRS complexes [61], it is always necessary to perform a post-processing step. For this purpose, the search-back procedure proposed in [54] was slightly modified and used in all the studies performed in this research. This modified search-back procedure is proposed next.

[3]This applies for all studies where the R-peak detection was required.

Correction of ectopic beats

Given the initial RR-interval time series $\{RR_i\}_{i=1}^{N}$, with N the total number of RR intervals, a median filter of 5 samples, denoted by $\{\overline{RR}_i\}_{i=1}^{N}$, is computed. Then, an RR-interval RR_i is labeled as ectopic if one of the following conditions is satisfied:

$$\begin{aligned} &\left(RR_i < 0.85 \times \overline{RR}_i \quad \wedge \quad RR_{i+1} > 1.15 \times \overline{RR}_i\right) \\ &\qquad\qquad\qquad or \\ &\left(RR_i > 1.15 \times \overline{RR}_i \quad \wedge \quad RR_{i+1} < 0.85 \times \overline{RR}_i\right). \end{aligned} \tag{4.3}$$

If this is the case, then RR_i is replaced by $(RR_i + RR_{i+1})/2$. The different values in (4.3) were selected experimentally as the ones for which the R-peak detection performed best in several ECG signals.

Correction of erroneous or missing peaks

Each RR-interval was compared against $\overline{RR}_i$. If it was lower or equal than 20% of the median value [54], then it was merged with the previous RR-interval. On the other hand, if $RR_i \geq 1.8 \times \overline{RR}_i$, then a new peak was added at a distance $\overline{RR}_i$ from the previous peak. Then, the procedure was repeated, until $0.2 \times \overline{RR}_i \leq RR_i \leq 1.8 \times \overline{RR}_i$.

Algorithm 7 summarizes the entire procedure to detect the R-peaks and to construct the tachogram.

4.2.4 Test on the MIT/BIH arrhythmia Physionet data

In order to evaluate the performance of the QRS detection algorithm implemented in this study, the MIT/BIH arrhythmia public database [141] was used. This dataset is made up of 48 ECG recordings of about 30min each, that were recorded from patients of the Boston's Beth Israel Hospital between 1975 and 1979. The recordings consists of two-channel ECG signals measured with a sampling frequency of 360Hz. Two independent cardiologist annotated all heart beats in all the recordings, and made the distinction between normal and abnormal beats. In total, 110 122 heart beats were annotated, and 89 133 were labeled as normal.

Algorithm 7. R-peak detection and computation of the tachogram or RR-interval time series

Input: ECG signal.
Output: R-peak positions and tachogram $\{RR_i\}_{i=1}^{N}$.

1. Compute the flattened ECG as $F_{ecg} = U_{ecg} - L_{ecg}$ using Algorithm 6.
2. Find the position of the peaks in F_{ecg} using the adaptive thresholding step of the Pan-Tompkins algorithm.
3. Use the position of the peaks obtained in 2 to find the position of the R-peaks in the original ECG signal.
4. Compute the initial RR-interval time series $\{RR_i\}_{i=1}^{N}$ as the time differences between consecutive R-peaks.
5. Compute $\{\overline{RR_i}\}_{i=1}^{N}$ using a median filter of 5 samples on RR_i.
6. Find ectopic beats using the conditions in (4.3), and replace them by $(RR_i + RR_{i+1})/2$.
7. If $RR_i \leq 0.2 \times \overline{RR_i}$ then $RR_{i-1} = RR_{i-1} + RR_i$ and remove RR_i.
8. If $RR_i \geq 1.8 \times \overline{RR_i}$ then $RR_i = \overline{RR_i}$, and repeat from 7 until RR_N is reached.

Table 4.2 shows the results obtained using the flattening algorithm and the adaptive thresholding of the Pan-Tompkins algorithm. The total sensitivity (Se) and positive predictive value (PPV) for all the heart beats in the dataset, including normal and abnormal beats, are 97.89% and 99.98%, respectively. If the search-back procedure is implemented to remove ectopic and erroneous peaks, a sensitivity of 99.99% is then reached for all normal heart beats. These results are comparable with those in the literature, especially with the Pan-Tompkins algorithm, which also reaches a PPV of 99% [61]. In addition, the algorithm proposed in this section performs better than Pan-Tompkins in the detection of atrial premature complexes [151].

Table 4.2: *Performance of the QRS detection algorithm on the Physionet MIT/BIH arrhythmia dataset. All heart beats are include, normal and abnormal. Sensitivity (Se) and positive predictive value (PPV) are indicated.*

Record	Heart beats	TP	FP	FN	Se	PPV
100	2273	2273	0	0	1	1
101	1865	1860	1	5	0.99	0.99
102	2187	2187	0	0	1	1
103	2084	2084	0	0	1	1
104	2229	2228	11	1	0.99	0.99
105	2572	2568	7	4	0.99	0.99
106	2027	1618	0	409	0.79	1

Table 4.2 – *Continued from previous page*

Record	Total	TP	FP	FN	Se	PPV
107	2137	2122	1	15	0.99	0.99
108	1763	1508	7	255	0.84	0.99
109	2532	2532	5	0	1	0.99
111	2124	2123	0	1	0.99	1
112	2539	2539	0	0	1	1
113	1795	1784	0	11	0.99	1
114	1879	1852	11	27	0.98	0.99
115	1953	1953	0	0	1	1
116	2412	2390	3	22	0.99	0.99
117	1535	1535	1	0	1	0.99
118	2275	2279	0	9	0.99	1
119	1987	1987	2	0	1	0.99
121	1863	1863	1	0	1	0.99
122	2476	2476	0	0	1	1
123	1518	1514	3	4	0.99	0.99
124	1619	1619	0	0	1	1
200	2601	2578	8	23	0.99	0.99
201	1963	1856	2	107	0.94	0.99
202	2136	2094	3	42	0.98	0.99
203	2980	2580	0	400	0.86	1
205	2656	2656	0	0	1	1
207	1862	1862	0	0	0.88	0.97
208	2955	2942	0	13	0.99	1
209	3005	3005	0	0	1	1
210	2650	2587	0	63	0.97	1
212	2748	2748	0	0	1	1
213	3251	3251	0	0	1	1
214	2262	2250	0	12	0.99	1
215	3363	3363	0	0	1	1
217	2208	2204	0	4	0.99	1
219	2154	2154	0	0	0.99	1
220	2048	2031	0	17	0.99	1
221	2427	2397	0	30	0.98	1
222	2483	2441	0	42	0.98	1
223	2605	2605	0	0	1	1
228	2053	2052	16	1	0.99	0.99
230	2256	2256	0	0	1	1
231	1573	1158	0	415	0.73	1
232	1780	1780	0	0	1	1
233	3079	3079	0	0	1	1

Table 4.2 – *Continued from previous page*

Record	Total	TP	FP	FN	Se	PPV
234	2753	2753	0	0	1	1
Total	110122	107893	130	2229	**0.97**	**0.99**

4.3 Heart rate variability

Variations in the time differences between *normal*[4] heart beats are determined not only by factors such as age, physical activity, and life style [51], but also by the control mechanisms of the ANS [84]. These mechanisms correspond to the two branches of the ANS, namely, the sympathetic and parasympathetic nervous systems. As it was mentioned in the introduction, these two systems have competing effects on the heart rate, and their relative changes can be studied from the variability of the normal-to-normal (NN) intervals, or the so called RR-intervals. It is worth noting that a cardiac cycle begins with the atrial systole, which corresponds to the beginning of the P-wave. Therefore, to measure the time difference between each heart beat, it would be necessary to measure the time intervals between consecutive P-waves and not between consecutive R-peaks [84]. However, the detection of the P-wave is very difficult due to its low amplitude in the ECG. This is why an easier-to-detect wave, namely the R-peak, is used in the analysis of heart rate variability (HRV) [51].

Different methods to quantify the variations of the heart rate have been proposed in the literature [3, 29, 51, 129, 75], and they can be divided into two main groups: time domain and frequency domain methods. In addition, some of these methods can be used on a short-term (of about 2 to 5 minutes) or on a long-term recording (24h), and their interpretation varies according to the period of time under investigation. In this section, the different time and frequency domain measures of HRV that were used in this research, will be described.

4.3.1 Time domain analysis

Time domain analysis of heart rate variability includes the computation of very easy-to-compute and easy-to-interpret indices. For instance, the mean and standard deviation of the heart rate in a given time interval. Despite the simplicity of these indices, they have been proven to be a robust and powerful

[4]Heart beats generated from the sinoatrial (SA) node

Table 4.3: *Indices of heart rate variability (HRV) calculated on the normal-to-normal (NN) intervals, or the so-called RR-interval time series. Adapted from [38, 129].*

Time domain		
Index	**Description**	**Computation**
μ_R	Mean	$\mu_R = \mathrm{E}[RR_i]$
std(RR)	Standard deviation of the RR series	$\mathrm{std}(RR) = \sqrt{\mathrm{E}[(RR - \mu_R)^2]}$
SDSD	Standard deviation of differences between adjacent RR intervals	SDSD= $\sqrt{\mathrm{E}[(R_d - \mu(R_d))^2]}$, with $R_{d_i} = RR_{i+1} - RR_i$, and $\mu(R_d) = \mathrm{E}[R_{d_i}]$
RMSSD	Square root of the mean of the sum of the squares of differences between adjacent RR intervals	RMSSD= $\sqrt{\mathrm{E}[R_d]^2}$
r_k	Serial correlation coefficient	$r_k = \frac{\sum_{i=1}^{N}(RR_i - \mu_R)(RR_{i+k} - \mu_R)}{\sum_{i=1}^{N}(RR_i - \mu_R)^2}$ with $k = 1, \ldots, 5$, and N the total number of RR intervals
Frequency domain		
LF	Power in the LF band	$0.04\mathrm{Hz} \leq f \leq 0.15\mathrm{Hz}$
HF	Power in the HF band	$0.15\mathrm{Hz} \leq f \leq 0.4\mathrm{Hz}$
LF/HF	Sympathovagal balance	LF/HF

marker of mortality after myocardial infarction, and heart failure [118, 126, 219]. In addition, several authors in the literature have shown that parameters like the mean of the RR-interval time series can be used to detect apnea episodes [38, 54, 147], and the occurrence of epileptic seizures [76, 146, 224], especially those originated from the temporal lobe [33]. Table 4.3 provides an overview of the different time domain indices used in this research.

4.3.2 Frequency domain analysis

The analysis of heart rate changes in the frequency domain allows to find periodicities in the HRV, which can be associated with the different biological mechanisms that modulate changes in the heart rate [129, 153]. The action of the different biological mechanisms can be observed in different frequency bands. For example, variations below 0.04Hz are considered to be modulated by circadian rhythms, thermoregulatory cycles and regulatory mechanisms of blood pressure [130]. However, the origin of these variations is not yet fully understood [29], and no clear consensus has been reached on their interpretation [129]. Furthermore, very long-term ECG recordings (24h) are needed to be able to accurately quantify variations at such a low frequency band. This becomes a limitation for ambulatory ECG analysis, where changes in the physical activity of the subject contribute to the information content in this band. Therefore, any conclusion derived from these variations can be misleading.

Heart rate fluctuations in the low frequency band (LF; 0.04-0.15Hz) are thought to be mediated by both sympathetic and parasympathetic branches of the ANS [29]. On the other hand, high frequency (HF; 0.15-0.4Hz) changes have been observed to be mediated mainly by the parasympathetic activity through the respiratory sinus arrhythmia (RSA). In addition, the effect of the respiratory rate can be clearly observed as a peak in this frequency band, but special care must be taken in the interpretation when this rate goes below the lower limit of 0.15Hz [51].

The quantification of the power content in the LF and HF bands is regularly done using the FFT algorithm. This algorithm assumes stationarity, and it requires the time series to be sampled at a fixed rate. With this in mind, it is clear that the RR-interval time series, or tachogram, must be resampled in order to obtain interpretable results[5]. Resampling frequencies are normally between 2 and 10Hz [51], and cubic spline interpolation is preferred. The reason for this is the low overestimation of the power in the frequency bands that is achieved using cubic splines [52]. A different approach that can be implemented to quantify the power in LF and HF is the Lomb-Scargle periodogram [168], where the power spectral density of an uneven sampled time series can be computed.

Apart from fulfilling the requirement of a regular sampling time, the assumption of stationarity should be satisfied. This can be done by means of the short-time Fourier transform (STFT). The STFT analyses the frequency content of the signal by means of a moving window, where stationarity is then

[5]If the tachogram is not resampled, the results depend on the local heart rate, and it is no longer possible to make a comparison between different segments or subjects.

assumed. Another way to overcome this problem of stationarity is by means of the well-known *wavelet transform*. The quantification of the power in the "approximated" LF and HF bands can be also obtained from the different wavelet coefficients. However, special care must be taken in the selection of the decomposition levels, since they are determined by the selected sampling frequency. In this research the STFT algorithm was used to compute the frequency domain indices, and the frequency bands are indicated in Table 4.3.

4.4 ECG-Derived Respiration

Respiration plays an essential role in the diagnosis and monitoring of different conditions, such as respiratory diseases and sleep disorders. However, its recording is often associated with the use of invasive and intrusive sensors such as respiratory belts and thermistors. Despite the fact that these sensors are regularly used in a hospital setting and are unavoidable in different medical tests, it is very rare to find them in ambulatory systems. In fact, several monitoring systems avoid using these sensors, not only because of their interference with the natural breathing, but also because of the costs associated with their use. For these reasons, several studies have focused on the derivation of respiratory information from the ECG signal. Some have proposed algorithms to derive respiratory rates from the tachogram, and others have shown that an approximation of the respiratory signal can be computed from changes in the morphology of the ECG [38, 54, 108, 109, 193]. This approximated signal is called the ECG-derived respiration (EDR), and its computation is possible due to the mechanical interaction between the respiratory movements and the morphology of the ECG. More specifically, during each breathing cycle the electrical impedance of the thorax, and the relative position of the electrodes with respect to the heart, change due to variations of the amount of air in the lungs. Consequently, these variations can be detected from the changes in the amplitude of the different waves of the ECG.

In this research, three algorithms based on morphology changes are used to approximate the respiratory signal when no real respiration is available.

4.4.1 R-peak amplitude

This methodology, first described in [54], uses the amplitude of the R-peaks in a baseline corrected ECG signal. The baseline is computed by means of two median filters, namely one of 200ms that removes the QRS complexes and P-

waves, and one applied to the result of the previous filter, of 600ms that removes the T-waves. The resulting baseline is then subtracted from the original ECG, and the amplitudes of the R-peaks make up the EDR. In the remainder of this document, this EDR will be denoted by R_{amp}. It is worth noting that the time stamps of the EDR signal correspond to the position of the R-peaks, as it is illustrated in Figure 4.11.

4.4.2 Principal component analysis

The EDR signal computed using PCA takes into account not only the variations of the R-peak amplitude, but also the modulation of the whole QRS complex due to respiratory movements. Here, all QRS complexes are first detected and then segmented using a symmetric window of 120ms around the R-peaks. The length of this window was selected based on the width of a normal QRS complex: $100ms \pm 20ms$ [51]. Next, the mean is removed from the segments, and they are then aligned in a matrix $\mathbf{X}$ [119]. Since the amplitude of all points around the R-peaks are modulated by respiration, a mean variation is derived by means of PCA. This mean corresponds to the R_{pca}, which is the first principal component obtained from the eigendecomposition of the covariance matrix of $\mathbf{X}$: $\mathbf{\Sigma} = \mathbf{X}\mathbf{X}^T$. This methodology is summarized in Figure 4.11.

4.4.3 kernel principal component analysis

In [217] it was shown that a non-linear version of PCA, i.e., KPCA (Chapter 2), achieves better approximations of the real respiratory signal, when compared with the two previous methodologies. In this case, a non-linear transformation of the input space is performed by means of an RBF kernel function, which allows to take into account linear and non-linear interactions between respiration and the morphology of the QRS complexes. The input space corresponds to the 120ms windows that were segmented around the R-peaks, and the new, transformed space is called the feature space. It is in this new space that the traditional linear PCA is performed. After that, the EDR, R_{kpca}, is computed in the feature space and then mapped back to the input space by means of pre-image as discussed in Section 2.2.2.

As mentioned in Chapter 2, a problem with KPCA is the selection of the model parameters, σ^2 and n_c. For this particular application, the number of components is already known, namely $n_c = 1$, and only the kernel parameter σ^2 needs to be selected. This is because the first principal component accounts for most of the variance in the morphology of the ECG, which is caused by respiration [119]. Hence, the selection of σ^2 can be performed either by using

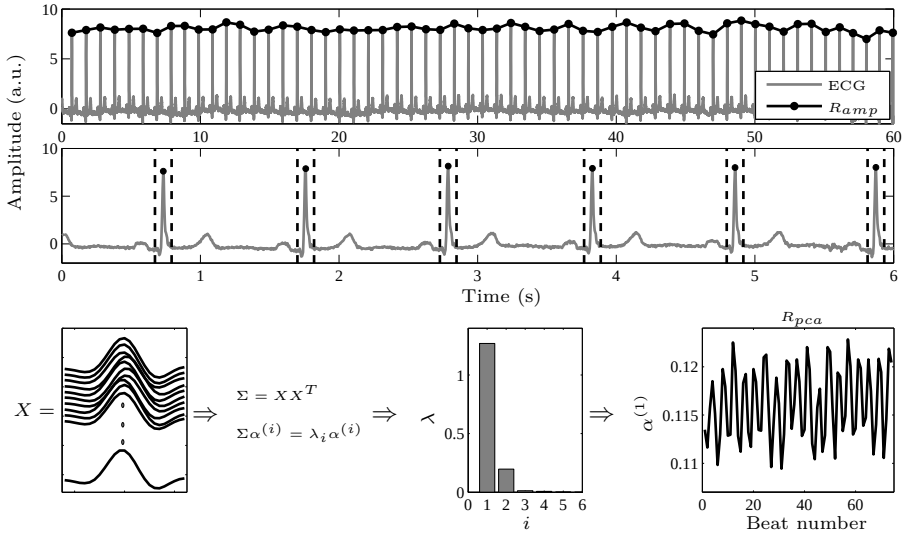

Figure 4.11: *Computation of the EDR signals using the R-peak amplitude, R_{amp} (top), and principal component analysis R_{pca}. To compute R_{pca} a symmetric window of 120ms around each R-peak is selected as shown in the second plot, from top to bottom. All the windows are organized in a matrix X (bottom left), and its first principal component corresponds to R_{pca}. Note that the first component describes most of the variance of the QRS complexes. For the computation of R_{kpca}, a non-linear transformation of X, instead of X is used.*

the methodology proposed in chapter 3, or through the maximization of the Shannon entropy of the kernel matrix, proposed in [217]. These two criteria will be compared next.

4.4.4 Evaluation of the EDR algorithms

Data

In order to evaluate the three EDR algorithms and the two model selection criteria for kPCA, the publicly available Physionet Fantasia database [96] was used. This dataset consists of simultaneous recordings of ECG and respiration. Single-lead ECG (lead II) and respiratory signals were recorded from 20 young (age between 21 and 34 years), and 20 elderly (age between 68 and 85 years) healthy subjects. The respiratory signals were measured using a respiratory belt around the thorax. All subjects were in supine position while watching

the movie Fantasia (Disney, 1940), and the signals were recorded for about 120 minutes with a sampling frequency of 250Hz. The heart beats were all automatically annotated and visually verified. In total, 4772 minutes were collected from this dataset and used in this evaluation.

Comparison

To evaluate the different EDR algorithms and model selection criteria for KPCA, the following procedure was implemented. First, all ECG and respiratory signals were segmented into epochs of one minute. Next, the position of the R-peaks was determined for each epoch, using Algorithm 7 (see Section 4.2). After that, 4 EDR signals were computed for each ECG segment. These signals correspond to R_{amp}, R_{pca}, R_{kpca1}, and R_{kpca2}, where the last two were obtained using MDD and the maximization of the Shannon entropy, respectively. As mentioned before, the time instances for the EDR signals correspond to the position of the R-peaks. Therefore, the signals were resampled at 2Hz using cubic spline interpolation. In order to evaluate their resemblance to the real respiratory signals, the latter were down-sampled at the same frequency, and the correlation coefficient and the mean magnitude squared coherence (MSC) were computed between each pair of signals. The correlation coefficient was determined as the maximum value of cross correlation over 10 lags [217], and the MSC was computed as

$$C_{xy}(f) = \frac{|P_{xy}(f)|^2}{P_{xx}(f)P_{yy}(f)}, \tag{4.4}$$

where $P_{xx}(f)$ and $P_{yy}(f)$ are the power spectral densities (PSD) of the signals x and y respectively, and $P_{xy}(f)$ is the cross-power spectral density of x and y. The PSD was computed using the Welch's algorithm, with a 1024 point fast Fourier transform (FFT), and a Hamming window of 30s with an overlap of 50%. The length of the window was selected in order to assure that at least 4 cycles at the lowest frequency limit of the respiratory band are included. This band corresponds to the HF band, which ranges from 0.15Hz to 0.4Hz [129].

For the computation of the mean MSC it is important to consider only information about the respiration. Therefore, the frequency band is defined around the fundamental frequency of the real respiratory signal, and the cutoff frequencies correspond to the half-peak amplitude values. This is illustrated in Figure 4.12.

For the calculation of the EDR using KPCA, the kernel parameter was tuned for each minute separately. This was done in order to consider baseline and dynamic changes in the respiration..

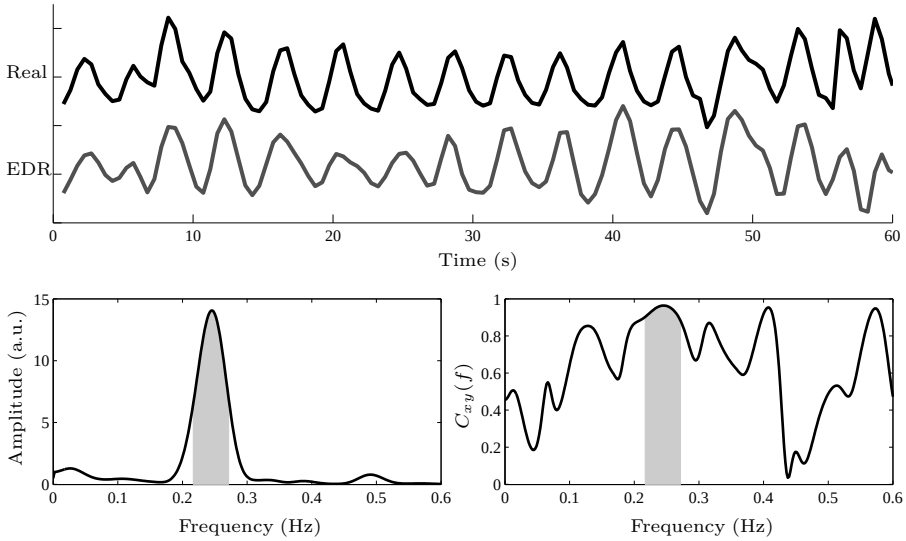

Figure 4.12: *Computation of the mean magnitude squared coherence (MSC). (top) Reference respiratory signal and EDR. (bottom left) PSD of the real respiratory signal, where the shaded area indicates the frequency range between the half-peak amplitude values of the fundamental respiratory frequency. (bottom right) Coherence between the respiration and the EDR signal. The mean value of the coherence $C_{xy}(f)$ inside the shaded area is used for the evaluation of the resemblance between the EDR and respiration. a.u. stands for arbitrary units.*

Statistical analysis

The differences between performances of the EDR algorithms and model selection criteria were evaluated with a confidence interval of 95%, using the Kruskal-Wallis test and the multi comparison test with Bonferroni correction. The numerical results and the box plots indicate the median and the 25th, and 75th percentiles.

Results

For the first part of the analysis, all 4772 minutes in the dataset were used for the comparison between the different techniques. The results of this comparison are presented in Figure 4.13(top), where the correlation coefficients and mean MSC between each EDR signal and the real respiration are indicated. Note that

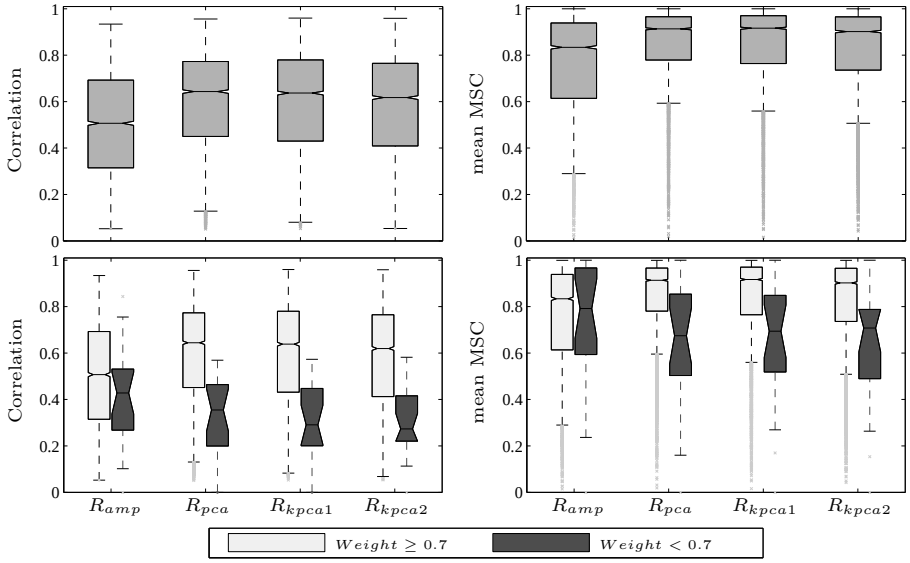

Figure 4.13: *Boxplots of correlation and mean magnitude squared coherence (MSC) between the reconstructed and the original respiratory signals for all the 4772 ECG minutes. Note that the approximation of the EDR based on the R-peak amplitude, namely R_{amp} is slightly worse than those computed using the other techniques. No significant differences were found either in the techniques or in the model selection criteria for kPCA.*

the values of correlation and mean MSC are slightly lower for R_{amp} than for the other EDR signals, nevertheless no significant differences were found between the different methods. In the second part of the analysis, the weights obtained using the ACF discussed in section 4.1.2 were computed for all segments. A threshold of 0.7 was then applied to split the data into two groups, one with clean segments and one with "contaminated" segments. The selection of this threshold was done at random, since the main goal was to determine how the different algorithms performed in the presence of noise or transients in the signals. As can be seen in Figure 4.13 (bottom), there are big differences in the correlation and mean MSC between the EDR signals obtained with PCA and kPCA, and the real respiration. Moreover, it is clear that the R-peak amplitude method appears to be less sensitive to noise. This is not a surprise, since it is well known that the performance of PCA is significantly compromised in the presence of noise [44]. Therefore, these findings can be considered in real life applications, where transients, artifacts, changes in baseline, and noise contaminate the ECG signals.

4.5 Conclusions

This chapter presented general algorithms that can be used for the pre-processing and analysis of the ECG signal in a wide variety of applications. The first problem tackled in this chapter was the detection of ECG artifacts. The main goal was to introduce two novel approaches, namely one based on the ACF, and one based on MDD. These two approaches were evaluated on a real dataset, and it was concluded that the ACF was the most appropriate and accurate method to find contaminated segments in the ECG. This due to its high accuracy and sensitivity to differentiate between normal and contaminated ECG segments, and the simplicity of its computation and interpretation. However, one disadvantage of the method is that it requires a training phase where ECG segments of relatively high signal-to-noise ratio are used.

In the second part of this chapter, the detection of the QRS complexes, and the construction of the tachogram were discussed. A modified Pan-Tompkins algorithm was proposed, and evaluated on a public dataset, where it was found that its performance was comparable with those obtained with algorithms proposed in the literature, and it outperformed other methods for the detection of ectopic beats. After testing the QRS detection algorithm, the construction of the tachogram was discussed. Furthermore, the different time and frequency domain parameters used for the analysis of heart rate variability were presented.

The last part of this chapter discussed the derivation of respiratory information from the ECG signals, and three different algorithms were described for this purpose, namely the R-peak amplitude, PCA, and kernel PCA. These algorithms were tested on a public dataset, where the approximated respiratory signals were compared against the real respiration. In addition, two model selection criteria for kPCA were tested, namely the maximization of the Shannon entropy and MDD. It was concluded that for real life applications the methodology based on the R-peak amplitude is less sensitive to noise than those based on PCA. In addition, no significant differences were found between the two model selection algorithms for kPCA, in this particular application. These findings suggest that the simplest method to extract respiratory information from the ECG offers reliable and robust performance, when compared to methods based on principal component analysis. However, some complex interactions between respiratory movements and the morphology of the ECG might be missing with this simple algorithm.

All methodologies presented and tested in this chapter can be easily implemented and used in several applications. Hence, a set of graphical user interfaces and Matlab functions were developed and are now available for future research on ECG (see Appendix B).

Chapter 5

Quantifying cardiorespiratory interactions

Physiological mechanisms like those responsible for the modulation of heart rate and respiration play an essential role in maintaining homeostasis, and regulating central nervous functions like arousal [79]. In fact, interactions between both mechanisms are strongly related to the capacity of the individuals to react to different stimuli. This is why the detection and quantification of directional interactions between heart rate and respiration are crucial for understanding the effect that different diseases and conditions have on the physiological control mechanisms of the autonomic nervous system. With this in mind, different linear approaches such as cross-correlation and coherence can be used to investigate the interactions in time and frequency domains of both heart rate and respiration. These methodologies allow to quantify linear interactions between time series, but they are not able to describe temporal changes in the interactions between the time series. For instance, the presence of non-stationarities poses a problem for this type of methodologies. This limitation can be overcome by phase rectified signal averaging (PRSA), which is an approach capable of quantifying quasi-periodicities in the dynamics of one process or in the interaction between two different systems. This approach will be described in section 5.1.

A different methodology that can be used to quantify directional interactions between heart rate and respiration is information dynamics, and it will be discussed in section 5.2. It allows to decompose the information contained in the heart rate into two different components. One component quantifies the information coming from respiration, and the other one describes all

interactions between heart rate and other vagal and sympathetic modulators. Even though the methodology used in this research to compute these components assumes stationarity and joint Gaussian distribution between the variables, it has many advantages that will be described further.

The last approach that will be described in this chapter is based on subspace projections. It leads to the computation of different components of heart rate, and it provides not only the quantification of their information content, but also their temporal variations.

5.1 Phase rectified signal averaging

Phase rectified signal averaging (PRSA) is a methodology originally proposed in [21, 22] to detect quasi-periodicities in non-stationary signals, in the presence of noise and artifacts. An example of such signals is the RR-interval time series, RR, where several dynamic changes and interruptions in the periodic oscillations of the heart rate can be observed. The reason for this is related to the fact that the RR signal is influenced and modulated by different physiological mechanisms at the same time. In order to deal with the non-stationary nature of RR, PRSA reduces the original signal into a compressed version, where quasi-periodicities can be unveiled.

5.1.1 Univariate PRSA

PRSA reduces a given time series $\mathbf{X} = \{x_i\}_{i=1}^{N}$ to a shorter sequence $\mathbf{X}' = \{x'_j\}_{j=-L}^{L}$, with $x, x' \in \mathbb{R}$, and $L < N$. This shorter sequence $\mathbf{x}'$ of length $2L + 1$ is computed following the three steps illustrated in Figure 5.1, and described as follows.

Anchor point (AP) selection

Given the time series $\mathbf{X} = \{x_i\}_{i=1}^{N}$, all points x_i that satisfy a given condition are selected as anchor points (AP). This condition can be selected depending on the type of information that needs to be analyzed. The conditions used in this research are:

1. Increasing points:

$$x_i > x_{i-1} \tag{5.1}$$

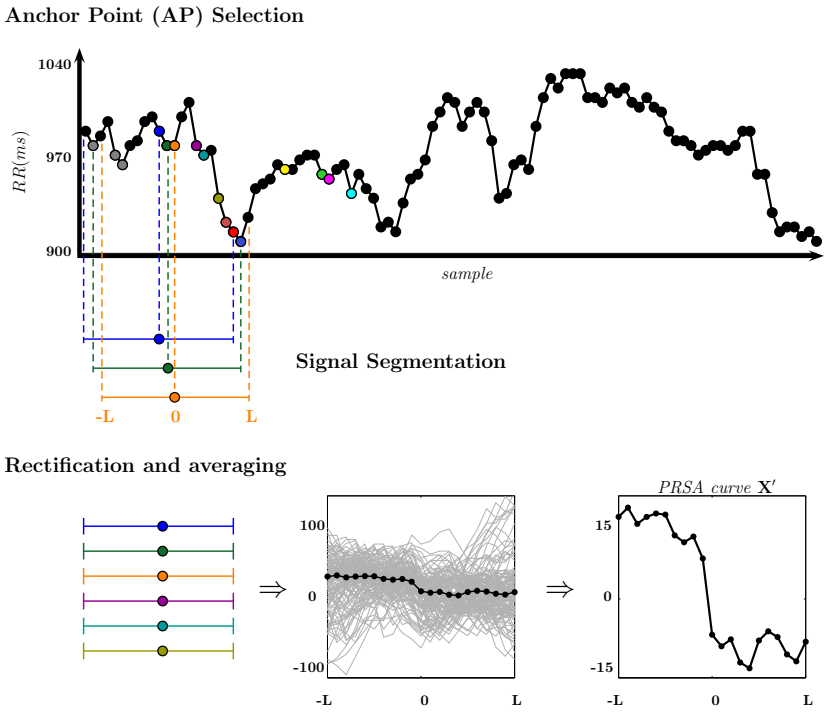

Figure 5.1: *Methodology to perform PRSA. The first step, illustrated at the top, selects all the decreasing points $x_i < x_{i-1}$ as AP. Next, the signal is segmented into windows of length $2L + 1$ around all AP. Note that the first 3 AP (in gray) are not used in the segmentation because they are not surrounded by the specified window. All the segments are then aligned around the AP and an average signal (PRSA curve) is calculated.*

2. Decreasing points:

$$x_i < x_{i-1} \tag{5.2}$$

3. Average increasing events:

$$\frac{1}{T}\sum_{j=0}^{T-1} x_{i+j} > \frac{1}{T}\sum_{j=1}^{T} x_{i-j} \tag{5.3}$$

4. Average decreasing events:

$$\frac{1}{T}\sum_{j=0}^{T-1} x_{i+j} < \frac{1}{T}\sum_{j=1}^{T} x_{i-j}, \tag{5.4}$$

where T is the number of samples in the time series taken for the average. This parameter can be seen as an upper limit of the frequency band that can be analyzed by PRSA. In [22] it was shown that T can be selected using $f \approx 1/2.5T$, with f the frequency under investigation.

Figure 5.1 shows the selection of the decreasing points, where it is clear that after satisfying condition 2, many points can be selected as AP. To avoid that artifacts are included in this selection, a threshold can be used to limit the inequality. For instance extreme values in the time series can be excluded from this selection if percentiles are set as thresholds.

Signal segmentation

A window of length $2L+1$ is defined around each AP, as it is indicated in Figure 5.1. Note that the first 3 AP are discarded because they are not surrounded by such a window. In this way a total of M anchor points and M windows remain. Since most of the points in the time series are selected as AP, it is clear that most of the windows will overlap. This becomes relevant when defining the parameter L, which should take into account the lowest frequency of interest. Since the presence of quasi-periodicities is normally unknown, it is crucial to find an optimal value for L that can reveal important information in the signal.

Rectification and averaging

In this step all the windows are aligned around the AP as indicated in Figure 5.1. Then, an average curve $\mathbf{X}'$, or so-called PRSA curve, is computed over all M windows. This curve now contains information of quasi-periodicities related to increases or decreases in the time series. All oscillations that are not phase synchronized with the AP are averaged out, as well as all the noise and artifacts.

From the PRSA curve illustrated in Figure 5.1 it is possible to identify very slow oscillations that occur before and after decreasing points in the time series. In addition, it is clear that a longer length L needs to be considered in order to observe the coherence time[1] of these oscillations. This time can be estimated from the decay of the oscillations around the AP. Other faster oscillations around the AP can be observed, and in order to better estimate their frequency it can be useful to look into the power spectrum of the PRSA curve. This procedure can be used to further analyze and quantify the information content

[1]Coherence time refers to the time during which the phase relationship between the oscillations remains nearly constant

in this curve. Other methods to study this signal include wavelet transform, linear regression among others.

5.1.2 Bivariate PRSA

The univariate version of PRSA allows to find quasi-periodicities in the presence of non-stationarity, noise and artifacts, by performing AP selection and segmentation on the same signal. In the bivariate version (BPRSA) however, these two steps can be done in separate signals, which leads to the investigation of the interrelations between two different processes [173]. In BPRSA the AP selection is done on the "driver" signal $\mathbf{X}$ (i.e. respiration), and the "target" signal $\mathbf{Y}$ (i.e. heart rate) is then used to perform the segmentation. Once the target is segmented, the BPRSA curve is computed as the average between the segments. Figure 5.2 illustrates the first two points of this procedure, namely the selection of the P and the segmentation of the target signal.

In BPRSA, the directionality of the causal relations between $\mathbf{X}$ and $\mathbf{Y}$ can be explored. This is one of the main advantages of BPRSA over other methodologies like cross-correlation and transfer function, where interactions between signals are also studied. Another advantage of BPRSA when compared with the cross-correlation function, is that it leads to flatter power spectra, where peaks and quasi-periodicities are enhanced. With this in mind, this methodology can be used here to investigate the interactions between respiration and the heart rate under different conditions. However, it is necessary to first understand how both PRSA and BPRSA can be related to the control mechanisms of heart rate and respiration.

5.1.3 Quantification and interpretation

Two different analyses can be performed using PRSA and its bivariate version. First, the quantification and interpretation of PRSA will be discussed. In this case, quasi-periodicities that precede an acceleration or deceleration of the heart can be analyzed. This is done by selecting the AP as the increasing or decreasing points in the heart rate signal, or RR interval time series. As a result, the "acceleration capacity" (AC) and "deceleration capacity" (DC) of the heart can be estimated from the PRSA curve [21] as

$$AC(DC) = \frac{1}{4}(x'_0 + x'_1 - x'_{-1} - x'_{-2}). \tag{5.5}$$

This characterization of the PRSA curve has been already investigated in the framework of risk stratification [106, 21], and of autonomic changes related

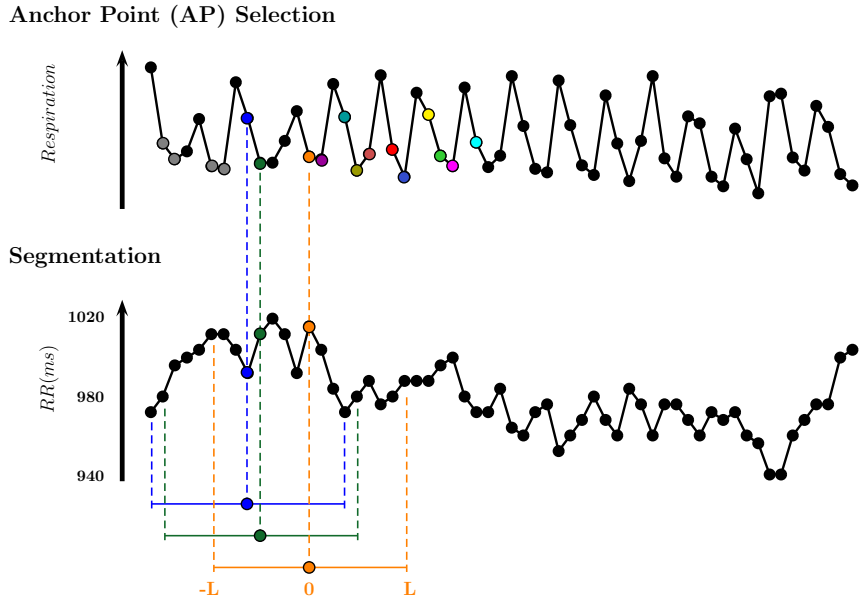

Figure 5.2: *Methodology to perform the bivariate version of PRSA. The first step, illustrated at the top, selects all the decreasing points $x_i < x_{i-1}$ as AP in the driver, respiratory signal. Next, the target signal RR is segmented into windows of length $2L+1$ around all the positions of the AP in the driver. The rectification and averaging steps are the same as in the univariate version.*

to age [127, 43]. For instance, in [43] it was found that the capacities of the heart to accelerate and decelerate are reduced with age in healthy volunteers. With this in mind and considering that DC and AC are strong indicators of the modulation of the ANS, the authors of [43] hypothesized that reductions in DC are related to a reduction in vagal tone. In a different study [106], the DC parameter was proven to be a good predictor for cardiovascular risk, where patients with low risk presented a higher DC value than their non-survivor counterparts. From these findings it is possible to assume that changes in the DC value correspond to changes in the parasympathetic activation of the heart. Unfortunately, no association between the AC parameter and the sympathetic activation can be made thus far. The reason for this relies on the fact that accelerations are thought to be also modulated by the parasympathetic system. Furthermore, some of the studies mentioned before did not manage to see any difference in AC between patient and control groups. Therefore, it is not possible to conclude that accelerations are mainly regulated by sympathetic activation.

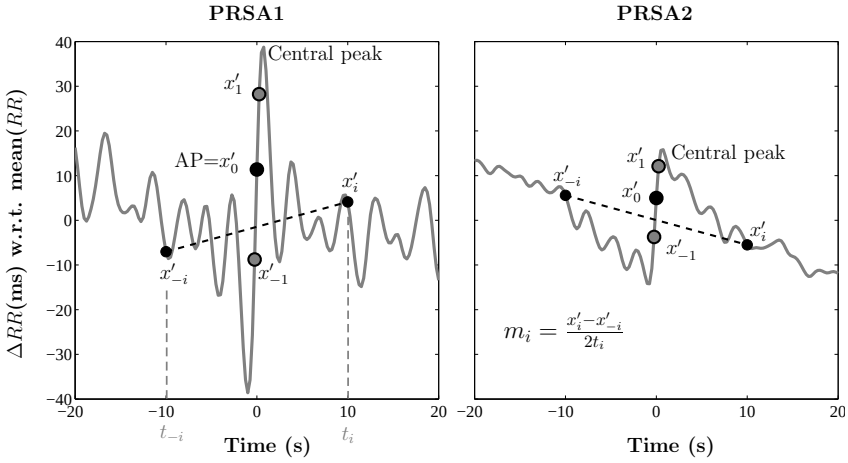

Figure 5.3: *Quantification of the PRSA curves. The anchor points x'_0, the central peak of the PRSA curves, and the segment connecting two points at a distance t_i from anchor point are indicated. The computation of the slope m_i of the segment is also defined.*

A different informative characteristic of the PRSA curve is the amplitude of its central peak. Figure 5.3 illustrates this parameter on two different PRSA curves, where it is clear that the amplitude of the central peak of PPRA1 is larger than for PRSA2. In [106] it was shown that this parameter can be associated with the strength of the regulation of the heart rate by the ANS. That is, the lower the cardiovascular risk, the more pronounced the central peak. As a result, a more enhanced regulation of the heart rate and a stronger reaction to external perturbations are observed. It is expected that control subjects have a more pronounced central peak than patient population with autonomic dysfunctions.

The last informative feature computed from the PRSA curve, which is used in this research is the slope of the line connecting points at specific distances from the AP [22]. In Figure 5.3 the computation of this slope is indicate as m_i. In this way, the time-reversal symmetry of the system, which cannot be seen in any power spectra, can be investigated. This can provide important information on the relation between what happens before and after the heart was accelerated or decelerated. From Figure 5.3 it is clear that oscillations at times t_i and t_{i-1} are not synchronized since the slopes of the dashed curves are clearly not equal to zero. In addition, the curves are clearly not symmetric with respect to the anchor point, which indicates different dynamics in the deceleration ($t < 0$) and stabilization of the heart ($t > 0$).

The features described above can be also derived from the curve obtained with the bivariate version of PRSA, namely BPRSA. However, their interpretation is different, since the BPRSA curve describes the interactions between two systems. Some studies have already used this methodology to explore the interactions between blood pressure and heart rate. For example, in [23, 144], BPRSA was used to assess the baroreflex sensitivity in heart failure patients, and it proves to be a good predictor of mortality. Other studies have highlighted the potential use of BPRSA to investigate the respiratory sinus arrhythmia (RSA), and it has also been proven to be a good predictor of mortality in post-infarction patients. In the last framework, a causal directional link can be assumed, where the respiration R triggers changes in the heart rate (i.e. RR-interval time series). In this way the AP are selected from the respiratory signal, and the segments are then obtained from the RR. Hence, oscillations in the heart rate that are in phase with the respiratory cycles are investigated. In the event that the BPRSA curve is flat, it can be concluded that no coupling between the heart rate and respiration exists. On the other hand, all the quasi-periodicities enhanced in the BPRSA curve correspond to periodicities triggered by inspiration or expiration periods. Therefore, it is clear that the information to the right of the AP in the BPRSA curve represents everything that follows the trigger event in the respiratory cycle. Conversely, oscillations observed before the AP can be related to everything that caused the respiratory event. As in the PRSA curve, the middle peak is strongly associated with the response of the heart rate to trigger events. One can expect that a dysfunctional cardiorespiratory coupling will produce a less pronounced peak. Furthermore, the AC and DC values described above can also be used here, but they reveal information on how fast the heart rate changes in phase with the respiratory cycle.

The selection of the parameters T and L strongly depend on the application. This also applies for the interpretation of the curves, hence, a deeper analysis on these curves will be discussed further in this book.

5.2 Information dynamics

Information dynamics is a framework derived from the field of dynamical information theory [124] that allows to measure the complexity of a system (i.e. heart) and its connectivity to other systems (i.e. respiration) in an integrated network. It leads to the estimation of the amount of information stored in the system, and the amount of information transfer from one system to the other. The *information storage* can be associated with how well the past of a system can be predicted from its own past [124], while *information transfer*

is the amount of information carried by a target system that can be predicted from its driver, and not from its own past [171].

An important advantage of this methodology over other more traditional time and frequency domain approaches (used to measure interactions between systems) is that it reveals causal links within the network. Furthermore, it allows the description of the spatial and temporal changes in the interactions, by introducing the probability of a system to arrive at a certain state given its past. This is the so-called *transition probability*, which provides the estimation of more dynamic information measures [64]. The formulation of these entropy-based measures will be described next.

5.2.1 Entropy decomposition

The time series method proposed in [64] is used here to formulate the different information measures.

Assume a stationary stochastic process $\mathbf{U} = [\mathbf{X}, \mathbf{Y}]$, with $\mathbf{X} = \{X_n\}_{n=1}^{N}$, $\mathbf{Y} = \{Y_n\}_{n=1}^{N}$, and N the length of the time series. X_n corresponds to the present of the driver system (i.e. respiration R), and Y_n the present of the target (i.e. RR) at time n. Moreover, all realizations of the processes $\mathbf{X}$ and $\mathbf{Y}$ before the time n are denoted by $\mathbf{X}_n^- = [X_{n-1}, X_{n-2}, \dots]$ and $\mathbf{Y}_n^- = [Y_{n-1}, Y_{n-2,\dots}]$ respectively. Figure 5.4(a) shows segments of respiration and RR, and it indicates the different measurements described above.

Given $\mathbf{U}$, the transition probabilities $p(Y_n|\mathbf{Y}_n^-)$ and $p(Y_n|\mathbf{X}_n^-)$ indicate the probability of process $\mathbf{Y}$ to be in the state defined by Y_n, knowing its own past and the past of the driver system. It must be noted that the respiratory sinus arrhythmia defines a causal relationship from $\mathbf{X}$ to $\mathbf{Y}$, where the heart rate is known to be driven by changes in respiration. This relationship implies that the uncertainty about Y_n can be resolved not only by knowing its own past, namely $\mathbf{Y}_n^-$, but also by taking into account the information transferred from $\mathbf{X}_n^-$. In fact, this resolution of uncertainty about Y_n can be quantified using the *predictive information* ($P_{\mathbf{Y}}$), defined as

$$P_{\mathbf{Y}} = H(Y_n) - H(Y_n|\mathbf{X}_n^-, \mathbf{Y}_n^-). \tag{5.6}$$

The first term corresponds to the *Shannon entropy* of Y_n, which refers to the uncertainty contained in the sample Y_n. The second term on the other hand, refers to the *conditional entropy* of Y_n, which quantifies the residual amount of information carried by Y_n that cannot be predicted by the past of the process $\mathbf{U}$. Figure 5.4(b) illustrates the interactions between the present of $\mathbf{Y}$ (i.e., Y_n), and the past of the system. $H(Y_n)$ is indicated by the red circle, the information

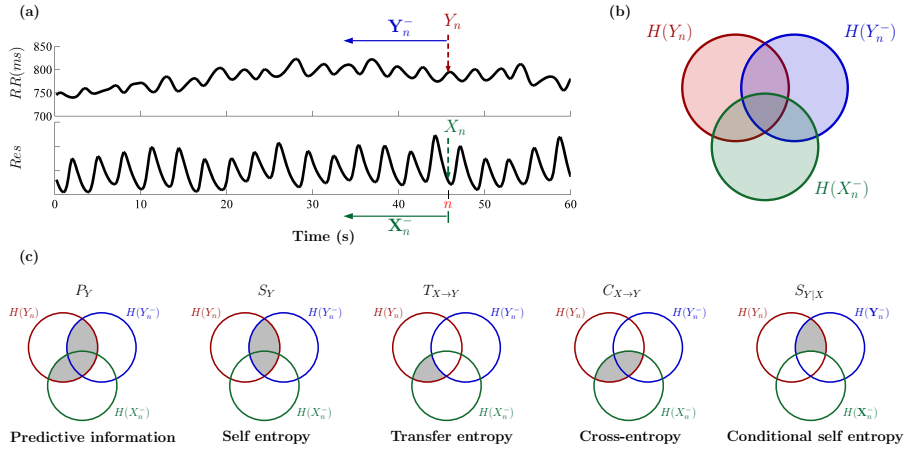

Figure 5.4: *Illustration of entropy decomposition using information dynamics. (a) One minute measurements of RR-interval time series (RR = $\mathbf{Y}$) and a respiratory signal (Res = $\mathbf{X}$). The horizontal blue and green arrows indicate the past of the heart rate and respiration, respectively. Y_n and X_n indicate the observation of heart rate and respiration, respectively, at the time n. (b) Illustration of the amount of information contained in the present of heart rate (red), and the past of heart rate (blue) and respiration (green). The intersections represent the amount of shared information between the different variables. (c) Interpretation of the different entropy terms. The shaded areas correspond to the information retrieved by each term.*

contained in the past of the system variables is indicated in blue and green, and the amount of information shared by the variables can be understood as the intersection region.

If the contribution to Y_n from the past of only one particular variable needs to be retrieved, the definition in (5.6) can be rewritten as

$$P_{\mathbf{Y}} = H(Y_n) - H(Y_n|\mathbf{Y}_n^-) + H(Y_n|\mathbf{Y}_n^-) - H(Y_n|\mathbf{X}_n^-, \mathbf{Y}_n^-), \tag{5.7}$$

where the term $H(Y_n) - H(Y_n|\mathbf{Y}_n^-)$ refers to the *self-entropy* $S_{\mathbf{Y}}$ [125], which corresponds to the amount of information carried by Y_n, that can be predicted by its own past. In other words, $S_{\mathbf{Y}}$ can be seen as the information storage of $\mathbf{Y}$. Large values of $S_{\mathbf{Y}}$ indicate that there is a high predictability of the heart rate. However, no distinction between the information coming solely from heart rate and the information transferred from respiration can be made. The second term quantifies the amount of information transferred from $\mathbf{X}_n^-$ to Y_n, and that cannot be predicted from $\mathbf{Y}_n^-$. This corresponds to the *transfer entropy*, $T_{X\to Y}$ [172].

The expression in (5.6) can alternatively be written as

$$P_{\mathbf{Y}} = H(Y_n) - H(Y_n|\mathbf{X}_n^-) + H(Y_n|\mathbf{X}_n^-) - H(Y_n|\mathbf{X}_n^-, \mathbf{Y}_n^-), \tag{5.8}$$

where the predictive information is defined as the sum of the *cross-entropy* $C_{X\to Y}$ and the *conditional self entropy* $S_{Y|X}$ [64]. $C_{X\to Y}$ quantifies the amount of information shared between Y_n and the past of $\mathbf{X}$. When there is a large amount of information transferred from the respiration to the heart rate, one expects this value to be also large. $S_{Y|X}$ corresponds to the residual amount of information that can be retrieved from the past of $\mathbf{Y}$. Figure 5.4(c) illustrates the meaning of these entropy terms.

5.2.2 Quantification of entropy measures

The computation of $S_{\mathbf{Y}}$, $T_{X\to Y}$, $C_{X\to Y}$ and $S_{Y|X}$ is done using the approach presented in [63] and [65], which makes use of the link between information theory and predictability. This procedure assumes that the process $\mathbf{U}$ has a joint Gaussian distribution. Under this assumption it is possible to describe the dynamics of $\mathbf{U}$ using a linear vector autoregressive (VAR) model of order p.[2] This description of $\mathbf{U}$ allows to relate the conditional entropy terms to the error probabilities of a regression model as

$$H(Y_n) = \frac{1}{2}\ln(\sigma^2_{Y_n}) + \frac{1}{2}\ln(2\pi e) \tag{5.9}$$

$$H(Y_n|\mathbf{Y}_n^-) = \frac{1}{2}\ln\left(\sigma^2_{Y_n|\mathbf{Y}_n^-}\right) + \frac{1}{2}\ln(2\pi e) \tag{5.10}$$

$$H(Y_n|\mathbf{X}_n^-, \mathbf{Y}_n^-) = \frac{1}{2}\ln\left(\sigma^2_{Y_n|\mathbf{X}_n^-\mathbf{Y}_n^-}\right) + \frac{1}{2}\ln(2\pi e), \tag{5.11}$$

where e is the Euler's number, $\sigma^2_{Y_n}$ is the variance of Y_n, $\sigma^2_{Y_n|\mathbf{Y}_n^-}$ and $\sigma^2_{Y_n|\mathbf{X}_n^-\mathbf{Y}_n^-}$ are the variances of the residuals of the linear regressions of Y_n on $\mathbf{Y}_n^-$ and on $[\mathbf{X}_n^-, \mathbf{Y}_n^-]$ respectively. These prediction error variances are expressed in terms of variances and covariances matrices as [63, 65]

$$\sigma(Y_n)|\mathbf{V} = \sigma(Y_n) - \Sigma(Y_n, \mathbf{V})\Sigma(\mathbf{V})^{-1}\Sigma(Y_n, \mathbf{V})^T, \tag{5.12}$$

where $\mathbf{V}$ can be either $\mathbf{Y}_n^-$ or $\mathbf{X}_n^-$, and $\Sigma(\cdot)$ and $\Sigma(\cdot,\cdot)$ correspond to the covariance and cross-covariance matrices respectively.

This procedure leads to an exact computation of the different entropy terms when the condition of Gaussianity is fulfilled. However, when this is not the

[2] In this study, the order p is selected using the Akaike Information Criterion (AIC).

case, the results are still valid, but they correspond to an approximation where only the linear part of the process is modeled.

5.2.3 Statistical significance

A consequence of using the parametric representation in (5.12) is that the statistical significance of the entropy terms can be computed. This is done by means of F-statistics, where two linear regressions are compared, namely the restricted linear regression and the unrestricted one. In the restricted linear regression Y_n is predicted from $\mathbf{Y}_n^-$ with the residual sum of squares indicated by RSS_r. Conversely, in the unrestricted regression Y_n is predicted from $[\mathbf{X}_n^- \mathbf{Y}_n^-]$ with RSS_u. The prediction errors produced by both regressions are compared and used to perform the test statistic defined as

$$F = \frac{(RSS_r - RSS_u)/p}{RSS_u/(N - nc)}, \tag{5.13}$$

where N is the length of the predicted time series and nc the coefficients for the unrestricted model. A given entropy term is considered statistically significant if its corresponding F is larger than the critical value of the Fisher distribution with $(p, N - nc)$ degrees of freedom, and a significant level $\alpha = 0.05$.

5.2.4 Interpretation on the cardiorespiratory interactions

In this approach, the time series RR and R are considered as observations of the underlying physiological mechanisms.

The link between each entropy term and the cardiorespiratory system can be made using the following scenarios. On the one hand, the values for $C_{X\to Y}$ and $T_{X\to Y}$ are expected to be larger in normal subjects than in subjects for whom the energy expenditure is not minimized by the RSA. In other words, differences in cardiorespiratory coupling between subjects can be observed when comparing these two entropy terms.

On the other hand, for patients for whom the heart rate is more fixed and predictable, the values of $S_{Y|X}$ and S_Y are expected to be larger than for normal subjects. $S_{Y|X}$ quantifies all variations in heart rate that could not be explained by respiration but are explained by its own past. As such, it reflects physiological mechanisms, different from respiration, producing predictable heart rate dynamics. The predictability of the heart rate is considered a limiting factor to react to acute situations like epileptic seizures [101]. In this case, a

lot of information is believed to be stored in the heart rate, which results in less flexibility and less adaptation to external factors.

Concerning the different entropy decompositions (5.7) and (5.8), it is clear that using (5.8) allows to split the information carried by heart rate into two components: respiration and others (e.g. sympathetic activation).

5.2.5 Why use information dynamics?

The main reason to select this approach based on information dynamics is that it allows to quantify the amount of information actively stored in the heart rate and the amount of information transferred from respiration to heart rate. This is one of the most important advantages of this method over more classical approaches such as coherence, since the concept of causality and directionality in the interactions can be exploited. However, frequency domain methods that extend the coherence towards an assessment of directional interactions (e.g., directed transfer function and partial directed coherence) can be seen as an alternative. Nonetheless, in this study the approach based on information dynamics is preferred for two reasons: first, it allows a nice decomposition of the overall predictability into the storage and transfer components, which is not straightforwardly guaranteed by frequency domain methods; second, it takes the whole causal influences from driver to target (i.e., from respiration to heart rate) into account, while a frequency domain method limiting the analysis to the conventional HF band may miss influences occurring outside of this band.

The entropy measures described above are extensions of the concepts of mutual information (MI) and conditional MI to the analysis of the “dynamic" information shared between the present of the target variable (i.e., heart rate) and its past and that of the other variable (in this case, respiration). Originally, measures of MI can quantify only “static" information shared between variables. However, the extension leading to the formulation of information dynamics, allows to use MI measures for assessing temporal flows of information, making them applicable in the context of time series analysis where the temporal information is essential to capture dynamic interactions.

5.3 Subspace projections

It is clear that the heart rate contains information about the activity of several physiological mechanisms at the same time. In fact, when the influence of one specific mechanism needs to be isolated, a decoupling methodology like

the one proposed in [40, 42], can be used. This methodology is based on Orthogonal Subspace Projections (OSP), and it allows to split the information contained in one signal into different components. In the case of the heart rate, OSP can be used to compute one component related to respiratory changes, and one related to physiological mechanisms different from respiration (e.g., sympathetic activation).

5.3.1 Formulation

Given are two physiological signals $\mathbf{X}$ and $\mathbf{Y}$. All the information contained in $\mathbf{Y}$ that is linearly related to $\mathbf{X}$, can be computed by projecting $\mathbf{Y}$ onto a subspace $\mathbb{V}$ defined by variations in $\mathbf{X}$. In order to define the basis for $\mathbb{V}$, three different approaches are implemented in this research and will be described later in this section.

After constructing a basis $\mathbf{V}$ for the subspace $\mathbb{V}$, any signal $\mathbf{Y}$ can be projected into the subspace, by means of

$$\mathbf{Y}_{\mathbf{X}} = \mathbf{P}\mathbf{Y}, \tag{5.14}$$

with $\mathbf{P}$ a projection matrix defined as

$$\mathbf{P} = \mathbf{V}(\mathbf{V}^T\mathbf{V})^{-1}\mathbf{V}. \tag{5.15}$$

Note that $\mathbf{Y}_{\mathbf{X}}$ describes the dynamics of $\mathbf{Y}$ related to $\mathbf{X}$, and the part of $\mathbf{Y}$ that is related to other mechanisms can be computed as

$$\mathbf{Y}_{\mathbf{X}^\perp} = \mathbf{Q}\mathbf{Y}, \tag{5.16}$$

with

$$\mathbf{Q} = I - \mathbf{P}. \tag{5.17}$$

This procedure is summarized in Figure 5.5, where $\mathbf{Y} = RR$ and $\mathbf{X} = R$. It is clear that the main result of this algorithm is the computation of a component of heart rate related to respiration denoted by $\mathbf{Y}_{\mathbf{X}}$, and a component related to mechanisms other than respiration denoted by $\mathbf{Y}_{\mathbf{X}^\perp}$.

5.3.2 Subspace definitions

Three different approaches are presented to compute the basis $\mathbf{V}$ for the subspace defined by variations in $\mathbf{X}$. These approaches are tuned for the computation of components of heart rate related to respiration, and components related to other mechanisms.

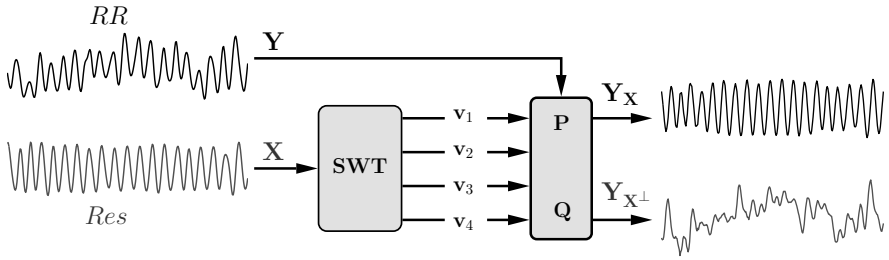

Figure 5.5: *Decoupling of heart rate (RR) and respiration (Res) using orthogonal subspace projections. The respiratory signal denoted by* $\mathbf{X}$ *is decomposed using stationary wavelet transform (SWT). This results in a respiratory space spanned by the basis* $\mathbf{V} = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4\}$. $\mathbf{Y}$ *is projected into this subspace by means of the matrix* $\mathbf{P}$ *and a component related to respiration* $\mathbf{Y}_\mathbf{X}$ *is computed. A component related to mechanisms other than respiration,* $\mathbf{Y}_{\mathbf{X}^\perp}$*, is computed using the matrix* $\mathbf{Q}$*, which is the orthogonal complement of* $\mathbf{P}$.

Wavelet decomposition

In [40], this subspace is defined by means of a 5 level, wavelet decomposition of the respiratory signal (i.e. $\mathbf{X}$) using *db4*[3] as a mother wavelet. From this decomposition, 6 different components of $\mathbf{X}$ are compute, namely 5 vectors $[\mathbf{d}_1, \ldots, \mathbf{d}_5]$ describing the detail coefficients at different levels, and 1 vector $\mathbf{a}_5$ corresponding to the approximation coefficients of level 5. Each one of these components describes dynamics at different frequency bands. For instance, if the sampling frequency of $\mathbf{X}$ is 5Hz, the corresponding approximated frequency bands are the following:

$$\begin{aligned}
&\mathbf{d}_1, \quad 1.25 \quad \leq f_1 \leq 2.5 \\
&\mathbf{d}_2, \quad 0.62 \quad \leq f_2 \leq 1.25 \\
&\mathbf{d}_3, \quad 0.31 \quad \leq f_3 \leq 0.62 \\
&\mathbf{d}_4, \quad 0.15 \quad \leq f_4 \leq 0.31 \\
&\mathbf{d}_5, \quad 0.07 \quad \leq f_5 \leq 0.15 \\
&\mathbf{a}_5, \quad 0 \quad < f_1 \leq 0.07.
\end{aligned} \tag{5.18}$$

[3]Daubechies wavelets 4

Consequently, if one wants to define a subspace $\mathbf{V}$ that describes respiratory dynamics at the HF band of HRV ($0.15 \leq f \leq 0.4$) [129], the detail coefficients of levels 3 and 4 can be used, namely $\mathbf{d}_3$ and $\mathbf{d}_4$. In addition, in order to impose a linear relationship between $\mathbf{Y}$ and $\mathbf{X}$, a delay m is introduced. In other words, the basis for the subspace $\mathbf{V}$ are $\{\mathbf{d}_3, \mathbf{d}_4, \mathbf{d}_3^-, \mathbf{d}_4^-\}$, where $\mathbf{d}_3^- = \{\mathbf{d}_3^{-1}, \ldots, \mathbf{d}_3^{-m}\}$ and $\mathbf{d}_4^- = \{\mathbf{d}_4^{-1}, \ldots, \mathbf{d}_4^{-m}\}$ are the delayed versions of $\mathbf{d}_3$ and $\mathbf{d}_4$ respectively up to a time m. In [216] this delay is set to 3 seconds. However, it has been shown in [51] and [194] that the phase difference between respiration and cardiac changes is not always constant. For this reason, this study proposes to compute the cross-correlation between the heart rate (i.e. $\mathbf{Y}$) and the respiratory signal $\mathbf{X}$ over 15 lags (i.e. 3s), and use the time with the maximum value of correlation as the delay parameter. This procedure is summarized in Algorithm 8.

Algorithm 8. Subspace definition using wavelet decomposition

Input: Heart rate signal $\mathbf{Y}$, and respiratory signal $\mathbf{X}$ sampled at 5Hz. Frequency band f to be described.
Output: Subspace $\mathbf{V}$.

1. Compute the 5 level wavelet decomposition of $\mathbf{X}$ and extract $\mathbf{W} = \{\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3, \mathbf{d}_4, \mathbf{d}_5, \mathbf{a}_5\}$.
2. Compute the correlation between $\mathbf{Y}$ and $\mathbf{X}$ over 15 lags (i.e. 3s), and set m as the time with the maximum value of correlation.
3. Find the vectors in $\mathbf{W}$ that intersect the specified frequency band f using (5.18). This new set of vectors will be denoted by $\mathbf{W}_f$.
4. Delay $\mathbf{W}_f$ from 1 to m times and obtain $\mathbf{W}_f^- = \{\mathbf{W}^{-1}, \ldots, \mathbf{W}^{-m}\}$.
5. Set $\mathbf{V} = \{\mathbf{W}_f, \mathbf{W}_f^-\}$.

Filtering

The reference respiratory signal $\mathbf{X}$ is first filtered using a band-pass, fifth order, Butterworth filter with cutoff frequencies at 0.04Hz and 0.15Hz. Next, the filtered signal and its delayed versions using a delay from 1 to m seconds as indicated in the previous approach, are used to construct the basis $\mathbf{V}$. Note that the range of frequencies used here corresponds to the low frequency (LF) band commonly used in HRV. In order to describe the high frequency (HF) band, $\mathbf{X}$ is filtered using the same filter, but now with cutoff frequencies at 0.15Hz and 0.4Hz. This approach is summarized in Algorithm 9.

Algorithm 9. Subspace definition using filtering

Input: Heart rate signal $\mathbf{Y}$, and respiratory signal $\mathbf{X}$. $0.04 \leq f_{LF} \leq 0.15$, and $0.15 \leq f_{HF} \leq 0.4$.
Output: Subspaces $\mathbf{V}_{LF}$ and $\mathbf{V}_{HF}$, describing respiratory changes in the low- and high-frequency bands commonly used in HRV.

1. Filter $\mathbf{X}$ using a band-pass, fifth order, Butterworth filter with cutoff frequencies at 0.04Hz and 0.15Hz. The resultant signal will be denoted by $\mathbf{X}_{LF}$.
2. Filter $\mathbf{X}$ using a band-pass, fifth order, Butterworth filter with cutoff frequencies at 0.15Hz and 0.4Hz. The resultant signal will be denoted by $\mathbf{X}_{HF}$.
3. Compute the correlation between $\mathbf{Y}$ and $\mathbf{X}$ over 15 lags (i.e. 3s), and set m as the time with the maximum value of correlation.
4. Set $\mathbf{V}_{LF} = \{\mathbf{X}_{LF}, \mathbf{X}_{LF}^{-}\}$ and $\mathbf{V}_{HF} = \{\mathbf{X}_{HF}, \mathbf{X}_{HF}^{-}\}$, where $\mathbf{X}_{LF}^{-} = \{\mathbf{X}_{LF}^{-1}, \ldots, \mathbf{X}_{LF}^{-m}\}$ and $\mathbf{X}_{HF}^{-} = \{\mathbf{X}_{HF}^{-1}, \ldots, \mathbf{X}_{HF}^{-m}\}$ are the delayed versions of $\mathbf{X}_{LF}$ and $\mathbf{X}_{HF}$ respectively, with a delay ranging from 1 to m seconds.

Driver signal

The subspace related to respiration is constructed using the original $\mathbf{X}$ and its delayed versions, using from 1 to m seconds. This procedure is described in Algorithm 10.

Algorithm 10. Subspace definition using the driver signal

Input: Heart rate signal $\mathbf{Y}$, and respiratory signal $\mathbf{X}$.
Output: Subspace $\mathbf{V}_{\mathbf{X}}$.

1. Compute the correlation between $\mathbf{Y}$ and $\mathbf{X}$ over 15 lags (i.e. 3s), and set m as the time with the maximum value of correlation.
4. Set $\mathbf{V}_{\mathbf{X}} = \{\mathbf{X}, \mathbf{X}^{-}\}$ with $\mathbf{X}^{-} = \{\mathbf{X}^{-1}, \ldots, \mathbf{X}^{-m}\}$ the delayed version of $\mathbf{X}$, with a delay ranging from 1 to m seconds.

The reason to use these three approaches to compute the basis for different subspaces $\mathbf{V}$, relies on the definition of the frequency bands. The first approach uses wavelet decomposition, which imposes frequency bands depending on the level of decomposition and the sampling frequency of the signal. In the second approach, the traditional LF and HF bands are explored. Finally, the last approach allows to describe "all" information of the heart rate that is related to respiration.

5.3.3 Quantification and interpretation

In order to quantify the different components $\mathbf{Y_X}$ and $\mathbf{Y_{X^\perp}}$, their relative power can be computed as [40]

$$F_{\mathbf{X}} = \frac{\mathbf{Y_X^T Y_X}}{\mathbf{Y^T Y}} \tag{5.19}$$

and

$$F_{\mathbf{X}^\perp} = \frac{\mathbf{Y_{X^\perp}^T Y_{X^\perp}}}{\mathbf{Y^T Y}}. \tag{5.20}$$

These ratios indicate the amount of power of each component relative to the total power of the signal. In this way, the interactions between heart rate and respiration can be quantified. For instance, in normal patients where a strong cardiorespiratory coupling takes place, the relative power of the component of heart rate related to respiration is expected to be higher than in patients with a cardiorespiratory dysfunction. On the other hand, the relative power of the orthogonal component quantifies the strength of other mechanisms affecting the heart rate.

An advantage of this methodology is that not only the relative power of each component can be obtained, but also a time series describing the temporal interactions between respiration and heart rate. Figure 5.6 illustrates the application of this approach on real data. From the top panel it is clear that there are changes in the heart rate that are modulated by the respiration. Hence, after filtering R using a band-pass filter with cutoff frequencies at 0.15Hz and 0.4Hz as indicated in Algorithm 9, step 2, the two components of the heart rate are derived. These components are indicated in Figure 5.6(middle panels), where the variations of heart rate related to respiration are clearly distinguished from those regulated by other mechanisms. Furthermore, when the power spectral densities (PSD) of these components are compared against the one of the real RR (Figure 5.6, lower panels) it is shown that other mechanisms different from respiration, are also acting on the HF band. These can be mainly related to other parasympathetic activity that takes place in the regulation of the heart rate [129].

Since the main interest here is to determine whether or not there is a strong interaction between respiration and heart rate, the relative power computed using (5.19) can be used. For the example presented in Figure 5.6, the power ratio between $\mathbf{Y_X}$ and $\mathbf{Y}$ in the HF band is equal to 0.51. In the case that respiration does not act as the main modulator, this value is expected to go down. However, it is important to keep in mind that respiration is not the only vagal activity influencing changes in the heart rate.

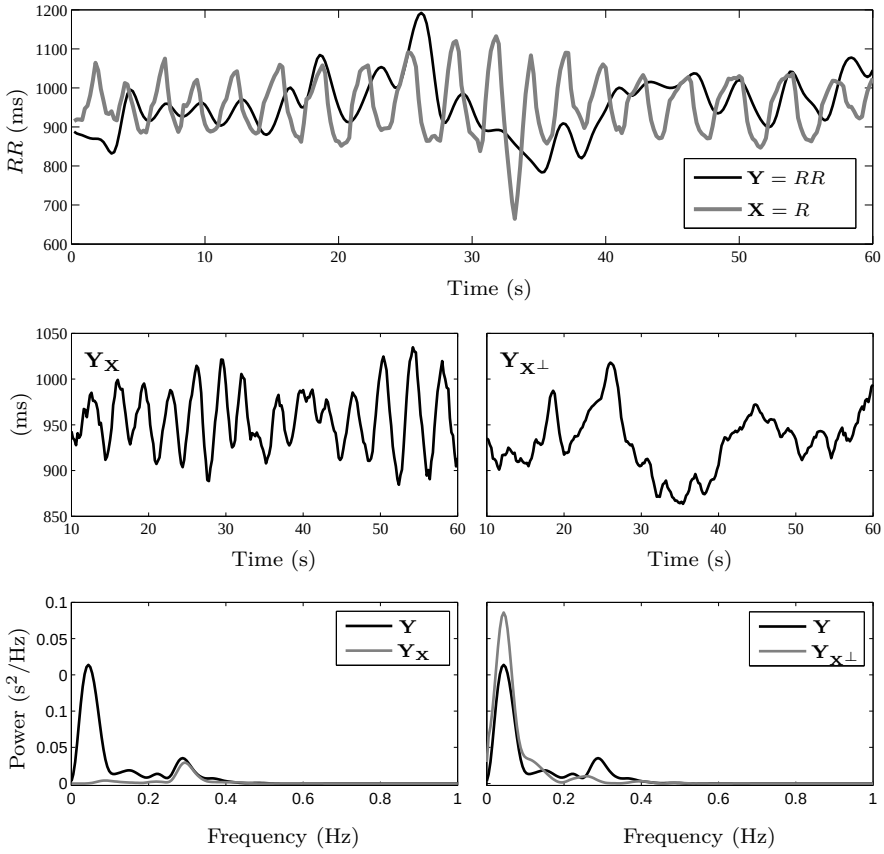

Figure 5.6: *Decoupling of the RR interval time series from the respiratory signal R using orthogonal subspace projections. The original signals are indicated on the top panel, where the y axis indicate the units of the RR. The respiration is superposed for comparison. The panels in the middle show the time series obtained using the filtering approach in the HF band, to compute the basis* $\mathbf{V}$*. These time series were obtained with a delay of 50 samples in the construction of the basis, hence, with a sampling frequency of 5Hz, they only take values from 10s. The power spectral densities of the dependent* $\mathbf{Y}_{\mathbf{X}}$ *and independent component* $\mathbf{Y}_{\mathbf{X}^\perp}$ *are shown in the lower panels. These PSD were computed using the Pwelch algorithms with a Hamming window of 30s and 50% overlap. Note that the power of* $\mathbf{Y}_{\mathbf{X}^\perp}$ *in the frequency band between 0.15Hz and 0.4Hz is not zero, which might indicate other mechanisms acting at this frequencies.*

The study of these time series and their relative powers will be included in the applications, where a deeper analysis will be provided.

5.4 Conclusions

This chapter introduced three different approaches that are used to describe and quantify interactions between two time series. In the first part of this chapter phase rectified signal averaging (PRSA) and its bivariate version BPRSA were discussed. These methodologies allow to find quasi-periodicities in time series in the presence of non-stationarities, noise and artifacts, which constitutes one of the main advantages of (B)PRSA. In particular, BPRSA also leads to the identification of causal relationships between two systems, which is of paramount importance in the study of cardiorespiratory interactions.

The second part of this chapter described an approach based on the emerging field of information dynamics, where the information contained in a process can be split into different components. These include the information stored in the system and the information transferred from one system to another. This can be exploited to analyze and separate the information contained in the heart rate, that is solely coming from the respiration. In addition, this methodology allows to quantify all information related to mechanisms different from respiration that take part in the regulation of the heart rate. These mechanisms include sympathetic activation and other parasympathetic modulations.

In the last part, a methodology based on subspace projections was presented. As described in information dynamics, subspace projection also allows to decompose the information contained in the heart rate into two components. However, the main difference is that the latter not only quantifies the amount of information in each component, but also provides two time series where temporal changes can be observed. This can become relevant when other parameters based on time and frequency representations are needed to characterize each component.

As can be seen, these three methodologies offer three different representations of the information shared between the respiration and the heart rate. Therefore, their interpretations might be complementary. This will be discussed in the next chapter, where the cardiorespiratory interactions in epilepsy will be investigated.

Part II

Applications

Chapter 6

Cardiorespiratory interactions in epilepsy

This chapter presents some of the main contributions of this research. The algorithms described in the first part of this document will be used to study the effect of epilepsy on the autonomic control mechanisms. In particular, the effects of interictal (i.e. in between seizures) activity on the heart and respiratory controls will be discussed. Many studies have focused on the analysis of HRV during seizures [175, 24, 102], especially those originated from the temporal lobe [33, 200, 123]. However, little is known about respiratory changes and cardiorespiratory interactions during interictal activity in childhood epilepsy. It is in this respect, that the results discussed in this chapter bring new insights into the understanding of the underlying mechanisms in epilepsy. The first part of this chapter will be devoted to introduce the main concepts in epilepsy that will be used along this research. In the second part, the cardiorespiratory interactions in children suffering from West syndrome will be investigated. Different methods to quantify these interactions will be used, and the effect that the continuous interictal activity, typical in West syndrome, has on the autonomic functions will be identified. Then, two other types of epilepsy will be studied, namely, temporal lobe and absence epilepsy. Again, the cardiorespiratory interactions will be quantified during in-between seizure activity, and differences on the cardiorespiratory control in these two types of epilepsy will be discussed. The fourth part of this chapter will present the analysis of pre-ictal (i.e. before the seizure) interactions between heart rate and respiration. Finally some conclusions and future research questions will be presented.

6.1 Introduction to epilepsy

Epilepsy is a brain disorder characterized by the occurrence of unprovoked seizures. These are synchronous or abnormal neuronal discharges that may affect different regions of the brain [70]. Many factors can be related to the origin of seizures, for instance, brain injury, brain tumor, stroke, structural or metabolic disorder of the brain, and genetic factors, amongst others. However, in most of the cases the causes are unknown [27].

According to the World Health Organization, epilepsy accounts for 0.5% of the global burden of disease [2]. Nearly 50 million people world wide suffer from epilepsy, and about 70% of them can be successfully treated with anti-epileptic drugs (AED). The incident annual proportion of epilepsy is between 40 to 70 new cases per 100 000, and most of these cases are children younger than 2 years or adults older than 65 years. Furthermore, epileptic patients suffer from consequences such as a higher risk of being involved in accidents and of suffering from sudden unexplained death (SUDEP). In fact, when compared to general population, their risk of premature death is 2 to 3 times higher [2].

6.1.1 Definitions

Some important terms that will be used along this chapter will be defined as follows:

- *Epilepsy:* neurological disease with at least two unprovoked seizures separated by more than 24 hours [70]. In addition, epilepsy can also be considered in patients who had one unprovoked seizure and after that a higher risk of occurrence of new seizures [69].
- *Seizure:* synchronous or abnormal discharges in a group of neurons in the brain.
- *Ictal activity:* time during an epileptic seizure
- *Interictal activity:* time between seizures
- *Seizure semiology:* symptomatology experienced by the patients before, during and after the seizure.
- *Status epilepticus:* prolonged seizure activity for more than 30 minutes, or many consecutive seizures in a period of time longer than 30 minutes. During this time the patient does not achieve full recovery, and cannot return to normal functioning.

- *Refractory epilepsy:* is a type of drug-resistant epilepsy, where the occurrence of seizures cannot be controlled.
- *SUDEP:* Sudden unexplained death in epilepsy
- *Epileptogenic zone:* is the "necessary and sufficient" region of the cortex for the initialization of seizures [128].
- *Vagus Nerve Stimulation (VNS):* Intermittent stimulation of the left vagus nerve suppresses the occurrence of seizures [86].

6.1.2 Diagnosis and treatment

Different methods are used to diagnose epilepsy:

- Neurological history of the patient
- Physical examination
- Electroencephalography (EEG)
- Neuroimaging tests like CT scan and MRI.

In order to consider that epilepsy is resolved in an individual, he or she must be seizure-free for at least 10 years, and not have taken anti-epileptic drugs (AED) for the last 5 years. In addition, if the individual was suffering from an age-dependent epilepsy, and he or she is older than the upper limit of the syndrome, then resolved epilepsy is also determined.

The main goal of treatment strategies in epilepsy is to control or stop the occurrence of seizures. This is done in order to improve the quality of life of patients, and avoid that they experience life threatening episodes. This usually can be achieved using AED, however it is known that in 30% of cases, epileptic seizures cannot be controlled [2]. In those cases surgical therapy can be seen as an alternative depending on the region of the brain where the *epileptogenic zone* is located. When a patient is not selected as candidate for surgery, VNS can be an option.

6.1.3 Seizure types

A broad classification of seizures is presented by the International League Against Epilepsy (ILAE)[27], and it consists of two main groups: partial and generalized seizures. A different groups called "unclassified epileptic seizures"

is also identified as a type of seizures. It includes all seizure activity that cannot be associated with any of the known classes, and it will not be discussed in this work.

This classification of seizures follows the one presented in [15] and is described as follows.

I. Partial seizures

These seizures are originated from a particular region in the brain and they may or may not spread throughout the entire brain. Partial seizures are subdivided into three main categories: simple partial, complex partial, and secondarily generalized seizures. The difference between simple and complex partial is the involvement of consciousness, which is only compromised in the latter category.

A. *Simple partial seizures* affect one particular region of the brain, and depending on that, different seizure manifestations can take place. For instance, when seizures involve the motor cortex, manifestations like abnormal movements, weakness, or coordinated actions in a group of muscles can take place. Other manifestations of simple partial seizures might include somatosensory, autonomic and physical distortions.

B. *Complex partial seizures* involve loss of consciousness and often do not have any extreme abnormal representation. This is to say, subjects experiencing this type of seizures can act like day dreaming or perform automatisms without being conscious of doing so. Complex partials are typically originated from the frontal or temporal lobes and rapidly spread to areas of the brain in charge of awareness and consciousness.

C. *Secondary generalized seizures* start in one part of the brain and rapidly spread to both hemispheres. These seizures involve loss of consciousness and typically abnormal, excessive motor activity as in tonic-clonic seizures. Due to their similarity with tonic-clonic, secondarily generalized seizures are difficult to differentiate. In addition, some times the partial onset is so short that no local manifestations can be distinguished. These seizures may last up to 5 minutes, and after that they are considered a medical emergency.

II. Generalized seizures

These seizures spread throughout a big part, or the entire brain. These types of seizures are typically characterized by loss of consciousness and in some cases by uncontrolled motor activity.

This group can be subdivided into six main categories: Absence, myoclonic, clonic, tonic, tonic-clonic, and atonic seizures.

A. In *absence seizures* the subject experiences loss of consciousness for up to 30 seconds, and in some cases the seizures are accompanied by slight movements like blinking, and chewing. Absence seizures are most common in childhood epilepsy, and they can occur several times during the day. Furthermore, they are typically unnoticed which makes them difficult to track, unless an EEG monitoring system is used.

B. *Myoclonic seizures* are very short episodes of contraction of a muscle or a group of muscles. These seizures last for no longer than 5 seconds, hence, they are very difficult to detect.

C. *Clonic seizures* consists of rapid and repetitive involuntary contraction/relaxation of a muscle (i.e. jerking). These seizures are very rare, and they are most of the time accompanied by a tonic phase.

D. *Tonic seizures* are characterized by stiffening of the muscles, and they are rarely accompanied by loss of consciousness. These seizures occur typically during sleep and they may last for few seconds (up to 20s). Patients performing regular activities (e.g. walking, standing) before the seizure onset will probably fall, or be involved in an accident, due to the change in muscular tone characteristic in this type of seizures.

E. *Tonic-clonic seizures* are a convulsive type of seizures that involve a tonic phase and a clonic phase. Moreover, subjects that experience this type of seizures loose consciousness and are at higher risk of being involve in accidents or life threatening episodes if they are not attended during the occurrence of the seizure. The reason for this is that these seizures might be accompanied by tongue biting and absence or shallow breathing that can cause the patient to asphyxiate. Usually these seizures last for minutes, often between 1 to 3 minutes, and their recovery period is typically longer when compared with other types of seizures.

F. *Atonic seizures* also known as drop attacks, are characterized by loss of muscular tone, which causes the subject to fall. These seizures rarely involve loss of consciousness, and they last no more than 15 seconds.

6.2 Cardiorespiratory interactions in West syndrome

West syndrome (infantile spasms-IS) was first described in 1841 by W.J. West, who reported that his son was presenting clusters of tonic-flexor brief spasms accompanied by a long-term regression in psychomotor development [179]. Since then, children suffering from spams with an unknown cause are identified as suffering from the "West syndrome".

West syndrome is an age-dependent epilepsy that typically begins during the first year of life, and is considered a catastrophic form of epilepsy [155] due to its deep effect on the developmental outcome. Children suffering from West syndrome typically present poor response to anti-epileptic drugs, and once the spasms resolve over time, these children often develop other types of refractory seizures in their late childhood [179].

Besides developmental delay, this syndrome is also accompanied by hypsarrhythmia in the EEG, which is characterized by high-amplitude waves in a chaotic EEG background. An example of a typical EEG activity recorded from a child suffering from West syndrome is shown in Figure 6.1, and compared with a normal EEG segment. As can be seen in the figure, the high-voltage slow waves are clearly visible in all EEG channels, and sometimes the spike-discharges can spread throughout the whole brain and become generalized [155]. Furthermore, the hypsarrhythmia in the EEG often occurs during non-rapid eye movement (REM) sleep, and it is significantly reduced in REM sleep [215].

Keeping in mind that a continuous interictal activity is observed in the EEG of children suffering from this disorder it is possible to think that the cardiac and respiratory control mechanisms as well as the cardiorespiratory interactions are deeply compromised in this type of epileptic activity. This hypothesis is supported by different studies, where cardiac autonomic dysfunctions have been observed in chronic epilepsy [87, 66, 49, 99]. Additionally, others have shown that the severity of the autonomic changes is correlated with the duration of epilepsy [167], but more importantly, these autonomic dysfunctions are thought to be strongly related to sudden unexpected death in epilepsy [121]. Therefore, it is of paramount importance to monitor the cardiac and respiratory control in this type of epilepsy.

The aim of this study was to identify cardiac and respiratory changes during continuous interictal activity, typical of early epileptic encephalopathy. Cardiac changes have already been studied in [99], but here, these changes are investigated together with those in the respiratory signal. Furthermore, interactions between heart rate and respiration in children suffering from West syndrome are quantified and compared against those in control subjects.

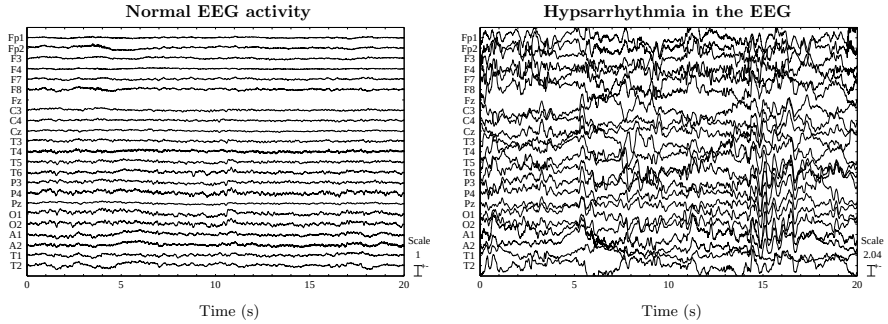

Figure 6.1: *Comparison between a normal EEG activity and a chaotic EEG behavior with high-amplitude waves characteristic in West syndrome.*

6.2.1 Data

Single-lead ECG signals were extracted from 24h video-EEG recordings of 10 children suffering from West syndrome (mean age 6 months, range 4-9 months), and 14 control subjects (mean age 8 months, range 5-11 months). No significant differences in age were found between both populations. All subjects in the study were referred to the epilepsy clinic of the university hospital Leuven (UZ Leuven, in Belgium). On the one hand, control subjects were investigated for non-epileptic sleep myoclonus and they all presented normal EEG activity. On the other hand, "West" patients presented epileptic spams, developmental arrest, and hypsarrhythmia in the EEG. The signals were recorded before treatment with adrenocorticotropic hormone (ACTH), which has been found to significantly reduce or completely eliminate infantile spasms [179]. However, some of them were receiving other treatment at the moment the video-EEG was performed. These treatments are indicated in Table 6.1. Patients did not have any anatomical lesion or brain stem abnormalities that could cause respiratory dysfunction [183], and this was confirmed by means of magnetic resonance imaging (MRI).

Simultaneous EEG, ECG (lead II) and video were analyzed by two EEG specialists, who paid special attention to clinical subtle signs of infantile spasms (IS). The ECG signals were extracted from the 24h video-EEG recordings, and they were sampled at 250Hz. Unfortunately no respiratory information was available in these recordings, hence, the three methodologies described in section 4.4 were used to derive this information from the ECG signals. The mean duration of the ECG signals is 14h (range 10h-22h), where both sleep and awake periods were included. In total about 320 hours of recordings were collected.

Table 6.1: *Characteristics of West patients*

no.	Age at the onset of IS (months)	Gender	Treatment at the time of EEG	Etiology	Age at EEG (months)
1	5	Female	Vigabatrin	Symptomatic ACM infarct	7
2	4	Male	Clonazepam	Unknown	11
3	6	Male	Valproate	Unknown	6
4	5	Female	Vigabatrin	Unknown	6
5	7	Female	Vigabatrin, valproate	Unknown	9
6	9	Male	None	Trisomy 21	10
7	4	Male	Valproate, topiramate	Tuberous sclerosis	10
8	9	Male	None	Unknown	9
9	5	Female	None	Unknown	5
10	7	Female	Vigabatrin	Mowat-Wilson	7

6.2.2 Methods

Each ECG signal was analyzed in non-overlapping segments of 5 minutes. This was done in order to get a good estimation of dynamics taking place at the low frequency band of the HRV [129]. The results obtained with this segmentation will be compared with those reported in [99, 100].

The R-peaks were first detected using Algorithm 7 (see Section 4.2), and then used to derive the respiratory information from the ECG signals using the three EDR algorithms discussed in Section 4.4. The derived EDR signals will be denoted by R_{amp}, R_{pca}, and R_{kpca}. After that, the RR interval time series were computed for each segment, and the HRV parameters discussed in Section 4.3 were derived. Before doing that, the segments were detrended by removing the linear baseline.

All segments of 5 minutes were then characterized by a set of time and frequency domain attributes, one RR time series, and 3 EDR signals. The RR and the EDRs were used to compute the cardiorespiratory interactions by means of (B)PRSA, information dynamics, and subspace projections as presented in Chapter 5.

In order to avoid using highly contaminated segments, the noise level estimation

(i.e. weight) described in Algorithm 4, steps from 1-4, proposed in Section 4.1.2 and based on the autocorrelation function similarity was computed for each segment. The normalized weights obtained in Algorithm 4, step 4 were used to discard all segments with weights lower than 0.9. Note that at this point no 95th percentile threshold was used. Instead, a threshold of 0.9 was experimentally defined.

Differences between segments of the normal and West groups were evaluated using the Kruskal-Wallis test, with a 95% confidence. In addition, in order to avoid computing the different HRV metrics under the presence of non-stationarities, the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test was used to select all RR stationary segments. By doing this, changes in the autonomic tone within the segment were not considered for the analysis [51]. The KPSS test assesses the null hypothesis that a time series is stationary around a deterministic trend.

6.2.3 Results in West syndrome

In total, 3843 segments of 5 minutes were extracted from the dataset, and only 365 segments had a weight lower than 0.9. Of these 365, 294 belong to normal subjects and 71 to West syndrome patients. Consequently, the final data included in the analysis consisted of 1868 control, and 1610 West syndrome segments.

HRV analysis

After performing the KPSS test, 279 and 438 segments were selected as trend stationary, corresponding to the control and West syndrome groups respectively. However, since significant differences in HRV parameters have been observed between day and night [132], only segments collected between 9h and 17h were included. This means that 75 controls, and 130 West syndrome segments were analyzed. The time and frequency domain parameters for the selected segments are indicated in Figures 6.2 and 6.3. Only the time domain features that are significantly different are indicated. Note that the standard deviation of all RR intervals is significantly lower in West syndrome cases, which is an indication of a more fixed heart rate. Furthermore, the RMSSD values indicate a lower vagus-mediated control for West syndrome cases than for control subjects.

To compute the frequency domain parameters, all the RR signals were sampled at 5Hz using cubic spline interpolation, and a 512-point Welch periodogram with a Hamming window of 40s and 50% overlap was used to compute

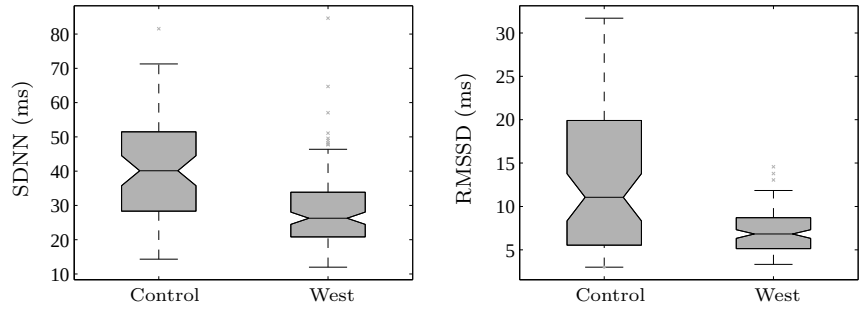

Figure 6.2: *Time domain HRV parameters with significant differences ($p < 0.05$).*

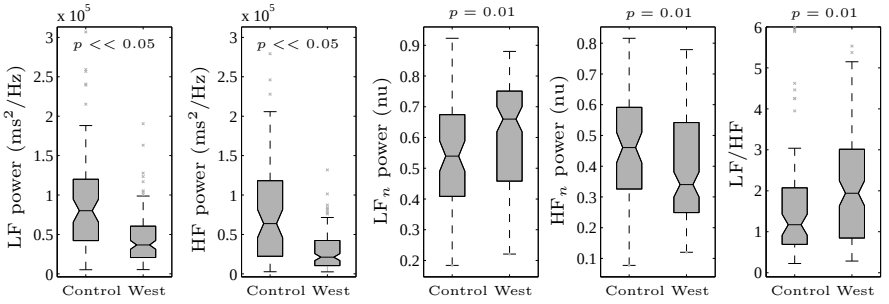

Figure 6.3: *Frequency domain HRV parameters. Note the reduced individual and total power in the West cases. The normalized values of LF (LF_n) and HF (HF_n) indicate a shift towards a more pronounced sympathovagal balance in West syndrome, which can be seen in the LF/HF values.*

the PSD. When looking at the frequency domain features (Figure 6.3), it is clear from the LF and HF powers that autonomic control for the West syndrome cases is significantly attenuated when compared to control subjects. However, when observing the normalized components: LF_n=LF/(LF+HF) and HF_n =HF/(LF+HF), an increased LF_n and a reduced HF_n is observed in West syndrome. This can be related to a reduced vagal tone, which results in an increased sympatho-vagal balance as can be seen in the LF/HF values of Figure 6.3.

Cardiorespiratory analysis

In order to analyze the cardiorespiratory interactions, three methods were used, namely PRSA, information dynamics, and subspace projections. Furthermore,

three EDR signals were used to quantify these interactions, R_{amp}, R_{pca}, and R_{kpca}. These respiratory signals together with the RR series were resampled at 5Hz.

The first analysis was performed using PRSA, and since this method is able to work with non-stationary signals, all the segments in the dataset were used at this point. The univariate version of PRSA was computed using all the deceleration points as anchor points (AP). In other words, all the increasing points in the RR were used as anchor points[1]. On the other hand, to compute the bivariate version of PRSA, all increasing points in the respiratory signal were used as AP. The reason to use deceleration points in the heart rate relies on the fact that it can be directly linked to vagal activity, which plays a key role in risk stratification [21].

A window of 30s around the anchor points was used to segment the RR. Figure 6.4 shows the PRSA and BPRSA curves for one West and one Control segment. These two examples show that the central peak of the PRSA curve is less pronounced in the West syndrome patient, thus it is possible to expect that these patients have a weaker regulation of the ANS [106], which is confirmed with the HRV studies discussed before. Moreover, the central peaks of the BPRSA curves computed using all three EDRs, show a less pronounced cardiorespiratory coupling in West syndrome. As mentioned before, these (B)PRSA curves only show two particular examples, but a better quantification of the differences between West syndrome and controls can be observed in Figure 6.5, which shows the (B)PRSA parameters derived from all the segments. It is important to remember that in this case, the slope (SB) of the line connecting the point before the AP and the AP, and the slope (SA) of the line connecting the AP with the point after the AP, indicate how fast the heart responds to decelerations, and it is clear that for West syndrome cases these slopes are significantly ($p < 0.05$) lower. This can be an indication of a less pronounced vagal regulation in these patients, which is in agreement with the findings described above. Another important parameter derived from the PRSA curve is the deceleration capacity (DC), which is significantly lower ($p < 0.05$) in patients than in controls.

From the BPRSA curves it is possible to see how fast the heart reacts to changes in the inspiration period. This is the case, because as in the PRSA curves, the AP were defined as the increasing points in the respiratory signal. In Figure 6.5 is is shown that the slopes before the AP in the BPRSA curves are lower for West syndrome cases. However, only significant differences are observed when the R_{amp} and R_{kpca} signals are used as drivers.

[1]The RR series is inversely proportional to the instantaneous heart rate. This means that increases in heart rate correspond to reductions of the RR intervals.

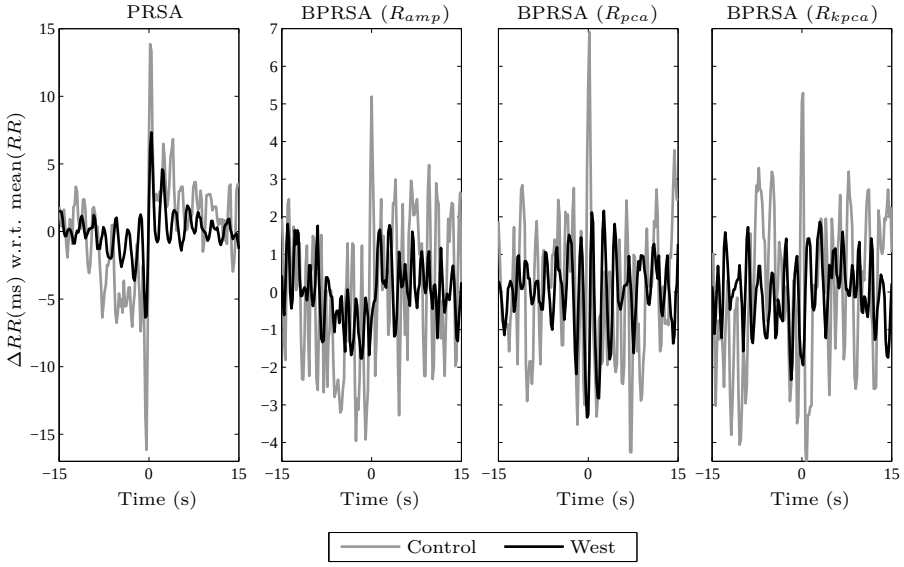

Figure 6.4: *PRSA and BPRSA curves. Note that the central peaks (time=0) are less pronounced for West syndrome cases than for controls. This is an indication of a reduced vagal-tone, and a suppressed cardiorespiratory coupling [106].*

Apart from analyzing the RR and EDR signals using PRSA, the different entropy measures computed using information dynamics, were also computed on the stationary segments. Nevertheless, no significant difference was found in this analysis. Furthermore, the subspace projection method described in Section 5.3 was used to decompose the $\mathbf{Y} = RR$ time series into two components: one related to respiration ($\mathbf{Y}_{\mathbf{X}}$) and one containing sympathetic modulations and possibly other vagal activations ($\mathbf{Y}_{\mathbf{X}^\perp}$). At this point, three methodologies were used to build the respiratory subspace, namely, wavelet (Algorithm 8), filtering (Algorithm 9), and the complete EDR signal (Algorithm 10)[2]. As a result, all group comparisons of the power contained in the different components were significantly different. Therefore, only the results obtained using the entire respiratory signal R_{amp} (Algorithm 10) are indicated in Figure 6.6. Note that for control cases, most of the variations in the heart rate are coupled with variations in the respiration, as can be seen in the power values of $\mathbf{Y}_{\mathbf{X}}$. This is not the case for West syndrome patients, where it is clear that a cardiorespiratory dysfunction might be taking place. Remember that for this decomposition, the entire frequency band was taken into account. In other

[2]These algorithms were presented in Section 5.3

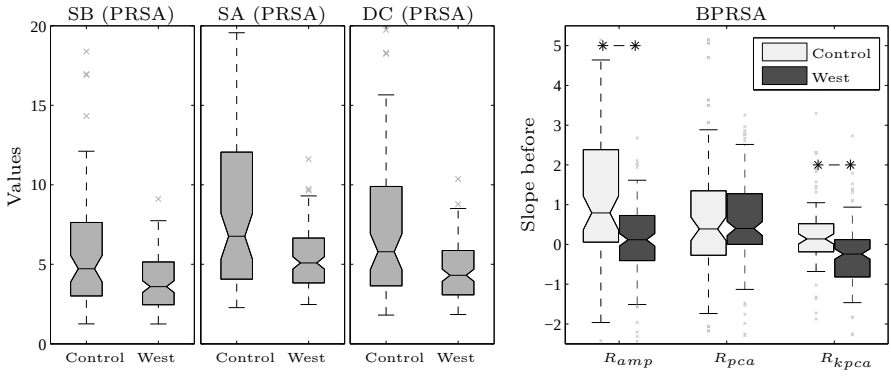

Figure 6.5: *Parameters derived from the (B)PRSA curves. The slope before (SB) the AP, slope after (SA) the AP, and deceleration capacity (DC) are all significantly lower for West cases ($p < 0.05$). Concerning the BPRSA curves, only the slopes before the AP in the curves computed using R_{amp} and R_{kpca} are significantly different for both groups ($p < 0.05$), and are indicated by $* - *$.*

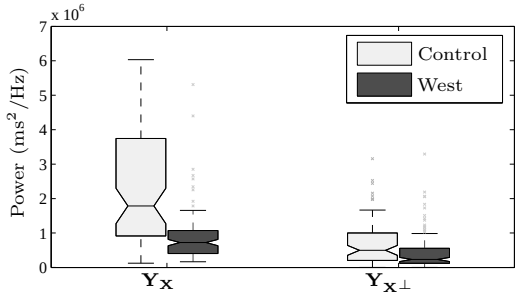

Figure 6.6: *Power ratios of the heart rate component related to respiration ($\mathbf{Y_X}$), and the component related to mechanisms different from respiration ($\mathbf{Y_{X^\perp}}$). A comparison between the power ratios for control subjects and West syndrome patients is indicated, and both group comparisons are significant ($p < 0.05$).*

words, no sub-bands, LF or HF, were considered in this analysis. This means that all effects affecting both respiration and HR are in a way quantified by the relative power of $\mathbf{Y_X}$. On the other hand, other variations in the HR are quantified by the power of $\mathbf{Y_{X^\perp}}$. From these results it is also possible to see an attenuated autonomic effect, as it was observed from the LF and HF values in the HRV analysis.

6.2.4 Discussion

From the HRV and cardiorespiratory analyses it is clear that patients suffering from West syndrome have a significantly reduced autonomic function. This is observed from the reduced SDNN, which is correlated with all the variability of the heart rate [129]. From these values, and from the powers in the LF and HF band it is clear that West syndrome patients suffer from autonomic dysfunction that can significantly compromise their defense mechanisms against hypoxia and asphyxia [72, 79]. In other words, their heart rate is more fixed, and less adaptable in acute situations that can lead to sudden unexplained death in epilepsy (SUDEP). Another important indication of this is the RMSSD parameter, which is an indication of vagal/parasympathetic control. It was found that when compared with control subjects, West syndrome patients have a significantly lower RMSSD. The latter has been correlated with SUDEP [57, 56], and it was determined that lower RMSSD values are an indication of a more dysfunctional control of the cardiac rhythm, which can increase the risk of developing lethal arrhythmias that may lead to SUDEP.

The fast adaptation of the heart rate can also be observed in the PRSA curves, where again, it was shown that it is significantly reduced in West syndrome patients. Moreover, the coupling between heart rate and respiration appears to be compromised as well. This is clear from the bivariate PRSA curves, and from the two different components of the heart rate, namely $\mathbf{Y}_{\mathbf{X}}$, and $\mathbf{Y}_{\mathbf{X}^{\perp}}$. The involvement of cardiorespiratory interactions indicate that the gas exchange in West syndrome patients is not optimal, which can be again interpreted as a dysfunction of the defense mechanisms against hypoxia and asphyxia. With this in mind, and taking into account the results reported in [100], it is possible to conclude that West syndrome deeply compromises the autonomic control mechanisms of heart rate and respiration.

Apart from quantifying the effect of epilepsy on the ANS, this study showed that PRSA and subspace projections can be used together with the well-known HRV parameters, to quantify variations in the regulators of the heart rate, especially vagal activation.

Keep in mind that the data used in this study consisted of continuous interictal activity, hence, autonomic dysfunction is not only typical during seizures. However, in order to determine if this is particularly the case in West syndrome, other seizures will need to be considered. In this respect, the next study was performed on interictal activity in temporal lobe and absence epilepsy in childhood.

6.3 Cardiorespiratory Variability in TLE and Absence Epilepsy

As mentioned before, seizures can be classified into two main groups, namely, focal and generalized seizures [27]. Focal seizures originate from one particular region and may or may not spread throughout the whole brain causing a secondary generalization. Generalized seizures on the other hand, involve spike-and-wave discharges in the whole brain, and most likely result in loss of consciousness [32]. The manifestations of the seizures closely depend on the location of the epileptic discharges. For example, seizures affecting the motor cortex can manifest as convulsions, trembling, jerking or rapid contraction/extension of the muscles. These are the types of seizures that are most commonly associated with epilepsy. However, there are other types of seizures with many other manifestations that are not "clearly distinguishable" and can remain unnoticed by a third person (e.g., care giver). This is often the case in absence epilepsy (AE), which is characterized by generalized seizures that manifest themselves as impairment of consciousness and sometimes as staring. Absence seizures are more common during childhood, they last only a few seconds, and they can occur several times per hour. Another important characteristic of this type of seizures is that high amplitude spike-waves discharges at frequencies larger than 2.5Hz can be observed in all channels of the EEG, as shown in Fig. 6.7. The fact that these seizures are often unnoticeable by a third person makes it difficult to track the frequency of the seizures and to assist the patients when seizures occur.

A different type of seizures with *mild* manifestations corresponds to the complex-partial seizures in temporal lobe epilepsy (TLE). These seizures receive the name of complex-partial because they also cause impairment of consciousness and most of the times they do not result in secondary generalization [62]. Furthermore, partial seizures in TLE can be accompanied by lip smacking, chewing, involuntary but coordinated movements, and they can last between 30s up to 2min.

It is well known that seizures are responsible for cardiovascular, cardiorespiratory and gastrointestinal changes [59]. For example, in TLE, seizures are associated with episodes of apnea [33], during which the heart rate (HR) is significantly increased while the respiratory rate decreases. In [182] it was concluded that the younger the patient, the higher the probability to have apnea episodes. However, this can also be influenced by the amount of anti-epileptic drugs (AED) the patient was receiving. Moreover, in [176] it was shown that patients with TLE do not always develop apnea episodes, but they can still reach low values of oxygen saturation. This desaturation can be due

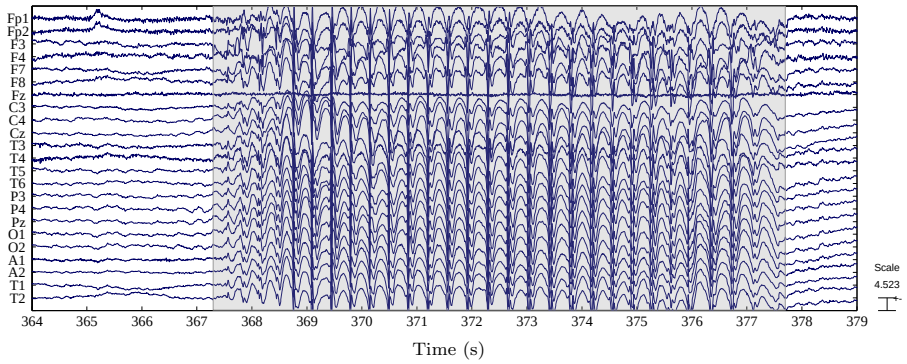

Figure 6.7: *EEG of a typical absence seizure. The high amplitude spike-and-wave complexes are indicated inside the shaded area. These oscillations have a frequency larger than 2.5Hz and they propagate throughout the whole brain. Note that the channel Fz located in the frontal area is not showing the same pattern. This is not related to the propagation of the epileptic discharges but to the malfunctioning of the channel during the whole monitoring period.*

to hypoventilation and result in sudden unexpected death in epilepsy (SUDEP) due to their inability to recover after an epileptic seizure. On the other hand, absence seizures do not seem to have any particular effect on the cardiac and respiratory functions. However, a more interesting effect is observed during interictal (seizure-free) activity. The heart and respiratory rate appear to be distorted even without the occurrence of an epileptic seizure. This could possibly be explained by a different connectivity in the brain when compared to normal subjects [101]. Furthermore, these distortions might play a key role in the pathophysiology of SUDEP.

Several studies have been conducted on autonomic dysfunction during seizures, however, limited attention has been given to seizure-free periods in epilepsy. For example, in [101] and [208] this was studied, but the heart rate and the respiration were analysed separately. As mentioned before, it was found that children suffering from AE showed deviations, both in the heart rate and the respiration, that were not related to seizure activity. This raises the interesting question whether or not the cardiorespiratory interactions are also affected. If this is the case with patients suffering from epilepsy, it could indicate that the mechanisms responsible for regulating central nervous functions, like arousal, are deeply affected. And this on its turn is closely associated with SUDEP.

This study investigates the differences in the cardiorespiratory interactions between patients with epilepsy and control subjects. Here, two types of epilepsy

are studied, namely, AE and TLE in childhood. The heart rate and respiratory signals are derived from ECG recordings, and the cardiorespiratory interactions are assessed by means of information dynamics [64]. The dataset used in this study will be described below.

6.3.1 Data

Single-lead ECG (lead II) signals were extracted from 24h video-EEG-ECG recordings, from 10 children with absence epilepsy (mean age 10.0±1.9 years), 10 children with temporal lobe epilepsy (mean age 10.3±2.7 years), and 10 control subjects (mean age 10.8±4.3) denoted respectively by AE, TLE and CO. These patients were referred to the epilepsy clinic of UZ Leuven in Belgium. Epilepsy patients needed a 24h video-EEG to assess the effect of therapy, and control subjects were suspected to suffer from epilepsy, but they were all found to be normal. None of the subjects was known to suffer from a cardiac nor a respiratory problem, and all patients were suffering from refractory (drug-resistant) epilepsy.

The ECG signals were recorded with a sampling frequency of 250Hz and they were segmented using the annotations provided by two independent EEG specialists, who indicated the start of each epileptic seizure. According to the guidelines presented in [129], the recommended range for heart rate variability analysis is defined from 250Hz to 500Hz. However, in order to increase the resolution of the R-peak locations the ECG was upsampled to 500Hz using cubic spline interpolation.

In order to capture the interictal activity, ECG epochs of 5 minutes were selected satisfying two conditions:

1. No epileptic discharges are observed in the EEG
2. The onset of any seizure is at least 30 minutes away.

In total 30 segments (one per subject) of interictal activity, were selected. These segments were selected during the same time of the day, namely between 13h and 15h, and during this time the children were awake and at ease on their beds.

6.3.2 Data processing

From each ECG epoch of 5 minutes, the RR interval time series and the ECG-derived respiration (EDR) were computed. The R-peaks were detected using

Algorithm 7 presented in Section 4.2, and verified by manual inspection. Due to the absence of recorded respiratory signals, breathing activity was derived from the ECG exploiting the well known respiration-induced modulation of QRS amplitude [143]. While this approach yields only an approximation of respiratory activity, techniques for the computation of EDR are widely used and have been validated thoroughly [103, 142, 169]. In this study, the respiration obtained with kPCA was used. Moreover, since it has been proven that differences in QRS location and RR time series measured from different ECG leads are strongly linked to breathing [78], all computations presented in this study are obtained using the same ECG lead, namely lead II.

The mean of the RR intervals, and the normalized power for each segment were computed. The power spectrum was computed using the Lomb-Scargle periodogram, and the frequency bands for the power computation were set as:

- Low frequency band (LF): $0.04\text{Hz} \leq \text{LF} < 0.15\text{Hz}$
- High frequency band (HF): $0.15\text{Hz} \leq \text{HF} < 0.4\text{Hz}$.

The power values were normalized to the sum of the total power in LF and HF. After that, the respiratory signals were filtered using a high-pass Butterworth filter with cutoff frequency at 0.05Hz, in order to remove baseline wander. Finally, the tachogram and the respiratory signals were normalized to zero mean and unit variance.

6.3.3 Statistical analysis

Differences between the groups (AE, TLE and CO) were evaluated using the Kruskal-Wallis test, which is a non-parametric version of the one-way ANOVA, with 95% of confidence. Since three groups were analyzed, a multi-comparison test was used with Bonferroni correction equal to $\alpha/3$, with $\alpha = 0.05$.

6.3.4 Results in TLE and AE

Once the RR interval time series RR, and the respiratory signal R_{kpca} were derived from each ECG segment of 5 minutes, time and frequency domain features were computed. Fig. 6.8 shows the comparison of the mean RR and the normalized power in the low and high frequency bands, between the three groups under study.

It is clear from Fig. 6.8 that the heart rate of patients suffering from AE is significantly higher than the one of control subjects. Remember that the figure

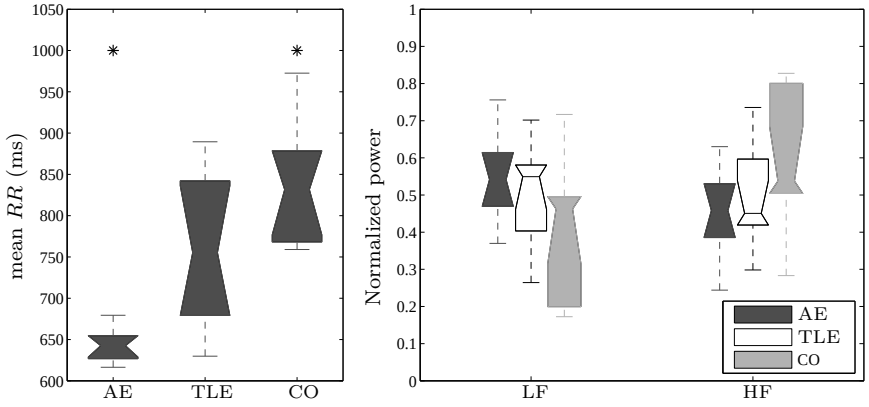

Figure 6.8: *Parameters derived from the RR interval time series (RR). Significant differences ($p < 0.05$) are indicated by *. Note that the mean RR is significantly lower (higher heart rate) for patients suffering from AE than for the control subjects. The frequency domain parameters were normalized to the sum of the total power in the low (LF) and high (HF) frequency bands. The LF band was taken from $0.04Hz$ to $0.15Hz$, and the HF band from $0.15Hz$ to $0.4Hz$.*

shows the RR intervals which is the inverse to the heart rate. These results were reported in [208] and [101], where several segments of one minute were used to make the comparison between AE and CO. Here, however, segments of 5 minutes were studied, and variations due to circadian rhythm were reduced due to the fact that the selected segments correspond to the same time of the day (from 13h to 15h). In addition, the patients were awake and at ease on their beds, which also reduces the influences of the body position on the heart rate and respiration. Fig. 6.8 also shows the frequency domain parameters of RR, where it is possible to observe that the power in the low frequency band is higher, although not significantly, in AE than in CO. This might be an indication that a higher sympathetic modulation is taking place in AE patients. This is also observed when looking at Fig. 6.9, which shows the mean and the standard deviation of the respiratory rates. It is clear from the figure that the variability of the respiratory rate is significantly higher in AE when compared to control subjects.

The main point made so far is that the differences in RR and R_{kpca} between CO and patients suffering from AE, were also found in the segments under investigation. The next step is to look at the interactions between both physiological signals, and determine if these types of epilepsy also have a profound effect on how these mechanisms interact with each other.

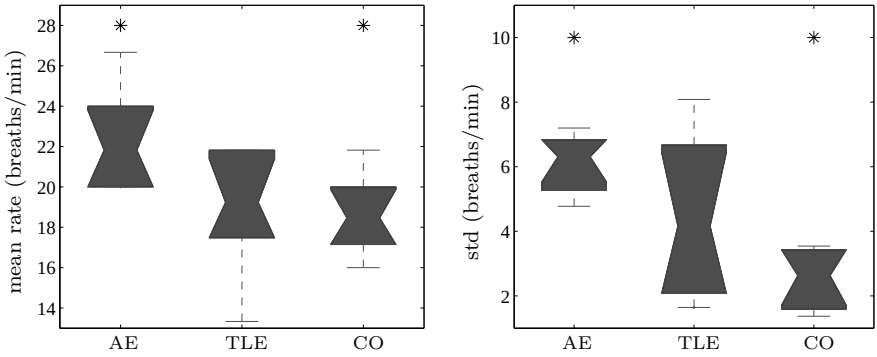

Figure 6.9: *Mean and standard deviation of the respiratory rate computed from the EDR signals (Res). Significant differences ($p < 0.05$) are indicated by *. Note that the mean respiratory rate and the variability for patients with AE is significantly higher than for CO patients, which suggests an increased sympathetic modulation.*

Fig. 6.10 shows entropy values for each group (AE, TLE and CO). Using the F-test to compute the statistical significance of the entropy terms[3], it was found that for all patients, the values of $S_{\mathbf{Y}}$, $C_{X\to Y}$ and $S_{Y|X}$ were significant, while values of $T_{X\to Y}$ were significant for 29 patients, with $9/10$ control subjects. Remember that two decompositions of $P_{\mathbf{Y}}$ were used: $S_{\mathbf{Y}}+T_{X\to Y}$ and $C_{X\to Y}+S_{Y|X}$. As can be seen from the figure, the only significant differences were obtained between AE and CO concerning the values for cross-entropy ($C_{X\to Y}$) and conditional self entropy ($S_{Y|X}$).

When looking at the predictive information, it is observed that heart rate variability is predicted slightly better in normal subjects than in patients with epilepsy. However, no significant differences were found. Therefore, it is very interesting to look at the different entropy terms in both decompositions: $P_{\mathbf{Y}} = S_{\mathbf{Y}} + T_{X\to Y}$ and $P_{\mathbf{Y}} = C_{X\to Y} + S_{Y|X}$. The first term corresponds to the self entropy $S_{\mathbf{Y}}$, for which only a slight trend towards higher information storage of RR in the control group is observed. Before drawing conclusions about the predictability of the heart rate in normal subjects, it is useful to look at the conditional self entropy $S_{Y|X}$. The latter quantifies the information stored in the heart rate, without taking into account information transferred from the respiration. In other words, $S_{Y|X}$ looks at mechanisms, different from respiration, that affect the heart rate (e.g. sympathetic activation). For this entropy term, it is observed that in control cases, there is less information shared between the heart rate and its past, which implies that the heart rate is

[3]Section 5.2.3 described how the statistical significance can be obtained for each entropy.

less predictable and more adaptable, thereby allowing a faster reaction on acute changes. This effect is more pronounced in the AE group, which also confirms that there is an interictal autonomic alteration in this kind of patients.

Turning now to the amount of information transferred from respiration to heart rate, it is necessary to look at the transfer entropy $T_{X\to Y}$, and the cross-entropy $C_{X\to Y}$. Fig. 6.10 also indicates that more information is shared between RR and Res in CO than in AE. TLE is in between both groups. This can be seen from both entropy terms, but the difference is statistically significant only for the cross-entropy. The higher sensitivity of the cross-entropy compared to the transfer entropy can be explained by the fact that the latter, in the presence of unidirectional interactions like those presumably existing from respiration to heart rate, tends to underestimate the information transfer ([65]). The significantly smaller information transfer measured by the cross-entropy in AE compared with CO suggests that the influence of respiration on the HR is also changed during AE, which confirms that not only respiration itself is altered in this type of epilepsy.

Additionally, it can be seen that the TLE group is more similar to the control group, which is an indication that TLE affects the cardiorespiratory interaction less than AE. This is in agreement with previous findings reported in [101] and [149].

6.3.5 Discussion

Prior studies have shown that when looking at interictal data of children suffering from epilepsy, physiological parameters like heart rate, respiration [101, 208] and heart rate variability [87] are significantly affected. In children with refractory epilepsy, the HRV is significantly lower, which can be explained by a reduced vagal activity. These observations were also clear in this study, where it was shown that patients with absence epilepsy (AE) present a reduced/increased parasympathetic/sympathetic modulation during seizure-free activity, in comparison to the control group. This was clear from the values of conditional self entropy that indicate a less flexible heart rate in AE, which can be seen as a predisposition of these patients to experience life-threatening episodes that can lead to sudden death. In [31, 87], it was discussed that an increased sympathetic activity may cause a reduction in the thresholds for life-threatening tachycardias or bradycardias. This needs to be taken into account when diagnosing children suffering from epilepsy, especially those suffering from AE.

Knowing that the heart rate and respiration are both affected in epilepsy, it is now crucial to also take into account their interactions, especially in AE. Two

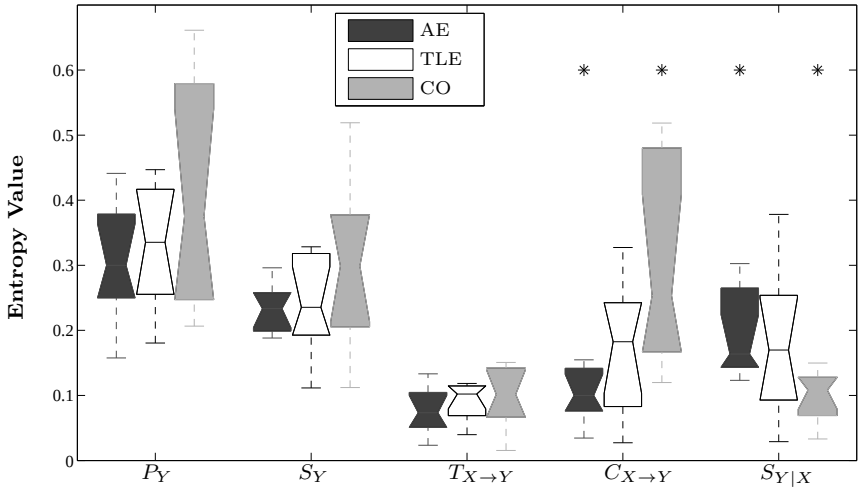

Figure 6.10: *Boxplots of predictive information P_Y, self-entropy S_Y, transfer entropy $T_{X\to Y}$, cross-entropy $C_{X\to Y}$, and conditional self entropy $S_{Y|X}$ for the three groups: absence epilepsy (AE), temporal lobe epilepsy (TLE) and control group (CO). Note that the second decomposition corresponding to $C_{X\to Y}+S_{Y|X}$, allows to find significant differences ($p < 0.05$, indicated by *) between AE and CO. This is an indication of an increased autonomic dysfunction in AE. TLE displays a less severe alteration to the autonomic controls, and it resembles better the properties of the control group.*

main points were observed here:

1. the information going from respiration to heart rate is significantly lower in AE than in CO,
2. and the heart rate in AE is highly predictable, which might indicate an increased sympathetic, and decreased parasympathetic activation.

On the other hand, no significant differences were found between TLE and CO. However, it is well known that these seizures are accompanied by bradycardia or tachycardia [182], and most of the times by apnea episodes [33]. Then again, this seems to occur during ictal activity in this type of epilepsy. Hence, the focus is now on interictal changes observed in AE.

It is known that in AE the thalamocortical network, which connects the thalamus and the cerebral cortex, has an important role in the generation of epileptic discharges [80]. Studies in rats have shown that the effectiveness of the synaptic responses in this network is lower in rats with AE. Furthermore,

there exist differences in the connectivity of this network when compared to controls. These two observations are very important, because the respiratory centers of the brain get input from brain systems like the thalamus, prefrontal cortex, amygdale and insula [154]. Hence, it is possible to think that the thalamocortical network could be more sensitive to changes in AE, causing a long term alteration to the respiratory control mechanism. In this way, the interactions between respiration and heart rate might also be affected, which can result in a dysfunctional regulation of central nervous functions, like arousal. This is of course a hypothesis that needs to be confirmed in further studies.

This study shows that there are some autonomic differences between AE and CO during seizure-free activity. Patients suffering from AE display a less adaptive heart rate and a modified cardiorespiratory interaction. This new insight might indicate a deeper underlying reason that is maybe not yet fully understood, such as for example the brain connectivity in absence epilepsy. Further research is needed to unravel the fundamental cause.

It was also shown that the approximate method based on information dynamics is a convenient and fast way to describe the linear interactions between respiration and heart rate. Using the exact formula for the entropy would take into account non-linear interactions, but this latter method is more demanding and prone to overfitting for short and noisy time series. However, non-linearities may need to be taken into account, since the dynamics of the cardio respiratory system might be more complex [79].

Future studies should investigate if these results change with the circadian rhythm. In other words, segments during periods of time outside the range between 13h and 15 should be analyzed.

6.4 Cardiorespiratory interactions towards seizure onset

Epilepsy is closely related to dysfunction of the cardiac and respiratory control mechanisms of the autonomic nervous system, as it was shown in the previous two studies (section 6.2 and 6.3). Furthermore, it is well known that epileptic seizures are not only accompanied by motor activity and stress responses, but also by different autonomic changes, such as tachycardia, bradycardia and hyperventilation [175, 146, 149]. Cardiac changes are thought to be caused by the propagation of epileptic discharges to the central autonomic network [175], which on its turn influences the output of the respiratory control mechanisms [102]. The presence of ictal cardiac disturbances has been reported in several

studies [175, 24]. In fact, these disturbances can also occur seconds before (pre-ictal) or after (post-ictal) the EEG changes marking the onset or offset of the seizure [102, 123]. In this study, the focus is on pre-ictal changes, and they can be seen as autonomic symptoms, or the first manifestations of seizures. Therefore, they constitute an important factor for the development of early detection systems.

As mentioned before, respiratory changes have also been observed during different types of seizures. For example, in [149] it was reported that seizures generated from the temporal lobe are typically accompanied by episodes of apnea. However, it is not yet understood whether apnea is secondary to the seizure, or if cardiac changes result as a response to oxygen desaturation caused by apnea [175]. It is certain, in any case, that there is a direct and profound effect of epilepsy on the cardiorespiratory control. Furthermore, the coupling between respiration and heart rate is thought to be affected due to deep involvement of autonomic centers during seizures. With this in mind, the question is whether cardiorespiratory interactions change "before" the EEG onset. And if this is the case, whether these changes show a transient behavior towards the seizure onset. In this study, these two hypotheses are addressed by means of information dynamics.

6.4.1 Data

The dataset used in this study consists of single-lead ECG (lead II) recordings extracted from continuous 24-hour video-EEG monitoring of 37 children admitted to the epilepsy clinic of UZ Leuven, Belgium. The mean age of the patients was 9.2 years (range 3-16 years), they were all suffering from refractory epilepsy, and none of them was receiving any drug with chronotropic action[4]. In total, 87 seizures were recorded, of which 48 were of focal nature, with 28 originated from the frontal lobe (FLE) and 20 from the temporal lobes (TLE), and 39 were generalized seizures. The latter group includes 10 generalized absence seizures (GAB), and 29 tonic or tonic-clonic seizures (GTN). The onsets of the seizures were annotated by two different EEG specialists based on video and EEG.

In addition to the ECG recordings, the respiratory efforts measured around the abdomen (R_a) and thorax (R_t) were also available for 18 seizures (8 FLE, 10 GAB). The sampling frequency for both the ECG and respiration is 250Hz.

[4]No changes in heart rate are caused by the medication

6.4.2 Pre-processing

Each ECG recording was used to compute the RR interval time series, and the EDR signals: R_{amp}, R_{pca}, and R_{kpca}, using the same procedure implemented in the previous studies. When the real respiratory signals were available, they were first low-pass filtered at 2Hz, and then sampled at the position of the R-peaks. These signals will be denoted by R_a and R_t, corresponding to the respiratory effort measured in the abdomen and thorax respectively.

6.4.3 Dynamic computation of cardiorespiratory interactions

In order to determine whether there is a transient behavior in the cardiorespiratory coupling before/during the seizures, it is important to consider the non-stationary nature of the signals, especially around seizure activity. Keeping this in mind, a moving window of 100 heart beats (HB) with an overlap of 99 HB is used to analyze the data. For each window, a set of entropy values $\mathcal{S}_{ri} = [C_{X\to Y}, S_{Y|X}, T_{XY\to Y}, S_Y]$ is computed, with $r = [R_{amp}, R_{pca}, R_{kpca}, R_a, R_t]$, $i = 1, \ldots, N$, and N the total number of windows in the signal.

6.4.4 Results and discussion

Figure 6.11 shows the RR and all the respiratory signals during an absence seizure. Note that after the seizure onset, the similarity between the real respiratory signals and the estimated ones is reduced. Remember that absence seizures are accompanied by loss of consciousness and staring, which means that no strong movement is expected to affect the ECG and real respiratory recordings. Therefore, it is possible to think that the direct relationship between ECG amplitude and respiratory movements is affected by this type of seizures. Furthermore, the respiration pattern is clearly affected during the seizure, together with the cardiorespiratory interactions, as discussed in the previous study,.

For the analysis, segments of 5 minutes were extracted from the dataset, starting 3 minutes before the seizure onset. Following the procedure described above, N sets of entropy parameters $\{\mathcal{S}_{ri}\}_{i=1}^{N}$ were computed for each pair $[RR, R_{amp}]$, $[RR, R_{pca}]$, $[RR, R_{kpca}]$, $[RR, R_a]$, and $[RR, R_t]$ considering that the last two pairs were only available for 18 seizures.

In Figure 6.12 it is shown that there are indeed transient changes in the amount of information transferred from respiration to heart rate. Furthermore, it can be seen that these changes start occurring before the seizure onset in the EEG,

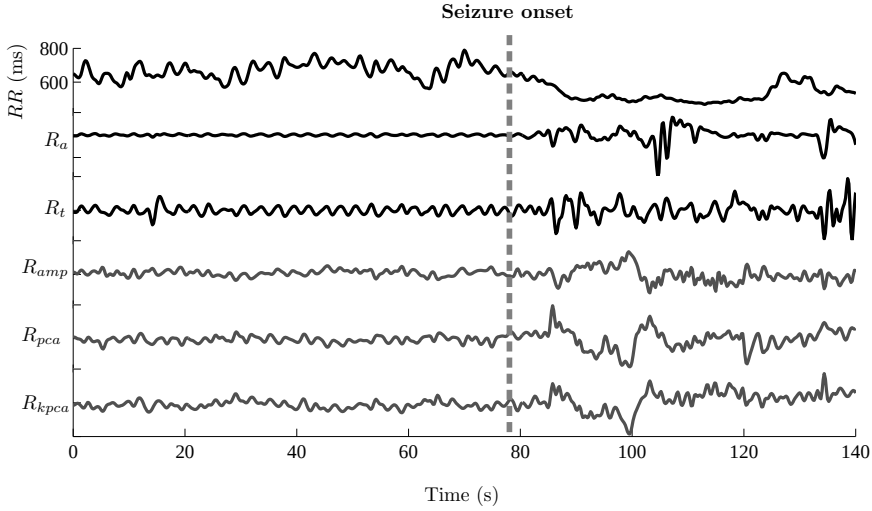

Figure 6.11: *Cardiorespiratory signals during an absence seizure. Note the visual similarity between the estimated and real respiratory signals before the EEG onset. This similarity becomes less clear with the seizure development.*

which means that these can be used for early detection of epileptic seizures, even when no "obvious" changes in heart rate are present. In order to quantify the changes in cardiorespiratory interaction around the seizures, the mean value for each entropy parameter was computed in a window of 60s around the seizure onset. This value was then compared with the one computed from a window of 60s in a seizure-free region. Differences between these values were evaluated with a confidence interval of 95% using the Kruskal-Wallis test.

It turns out that the interactions between respiration and heart rate depend on the type of seizure. Hence, the findings for the two types of seizures will be described separately in more detail below.

Focal seizures

Figure 6.13 shows the self entropy S_Y (left panel) and the conditional self entropy $S_{Y|X}$ (right panel) for four different respiratory signals (R_{pca}, R_{kpca}, R_a, and R_t). The values of S_Y and $S_{Y|X}$ for segments around the seizures (light grey boxplots) are compared with the values of S_Y and $S_{Y|X}$ within a normal segment (dark grey). Note that the information shared between heart rate and its own past (S_Y) is significantly different around the onset, which means that a more predictable heart rate is present around seizures. This has already

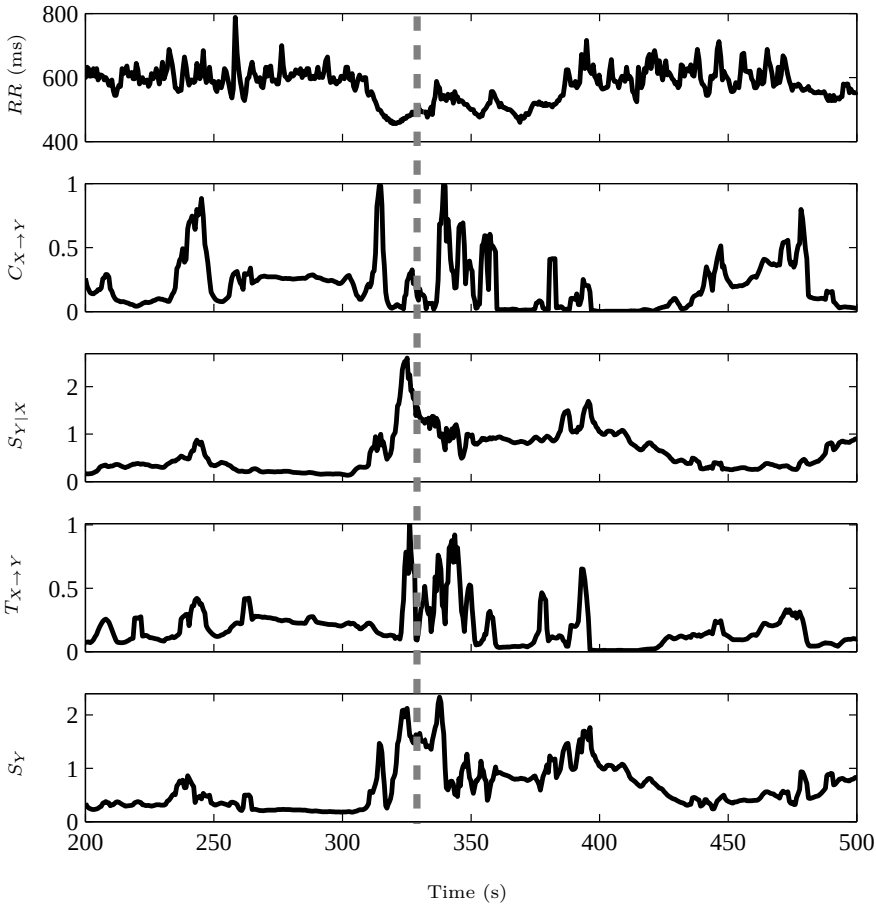

Figure 6.12: *Information dynamics analysis of an RR segment around a generalized seizure. The red dashed line indicates the onset of the seizure in the EEG. Note that the heart rate changes even before the EEG onset, and interactions between respiration and heart rate, are also affected by the seizure.*

been investigated in previous studies [123, 149], where it was shown that this effect is more significant for TLE. However, here it is observed that the seizures originated from the frontal lobe, are also accompanied by highly predictable heart rates, and can be easily classified using these values of entropy.

An important observation is the predictive role of respiration. This can be seen in Figure 6.13 when S_Y (left panel) is compared with $S_{Y|X}$ (right panel), as the differences are not as significant anymore when influences of respiration are removed (i.e. $S_{Y|X}$). Furthermore, it is remarkable that clear separations

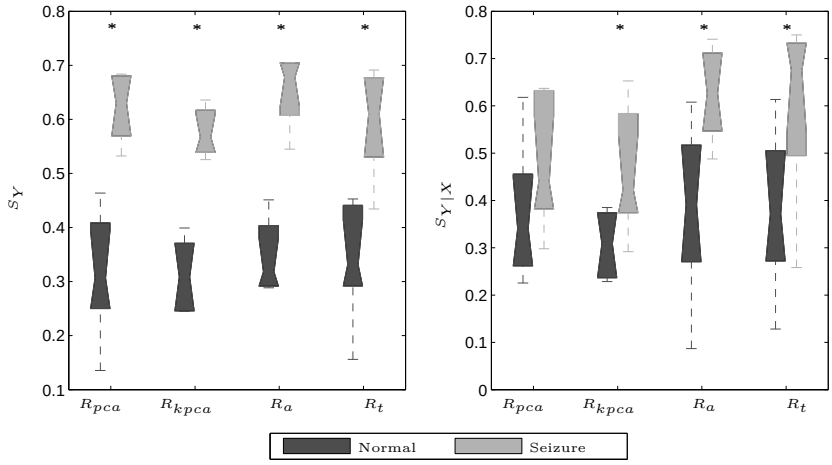

Figure 6.13: *Self entropy S_Y and self conditional entropy $S_{Y|X}$ values around the seizure onset, and during seizure-free activity. Four different respiratory signals were used, and significant differences are indicated by (*). Note that the discrimination of epileptic seizures is enhanced when the amount of information shared between the present of the heart rate, and its own past without removing respiratory influences (i.e. S_Y), is used.*

can be obtained even with an ECG-derived respiratory signal, which can be important in applications where a limited amount of sensors are available.

An explanation for these changes in cardiorespiratory information dynamics in TLE and FLE, can be found in the anatomic localization, relative to the central autonomic network, of the temporal and frontal lobe structures [102]. This proximity allows for the propagation of epileptic discharges from one of these lobes to the central autonomic network, which can result in cardiac and respiratory disturbances.

Generalized seizures

A different situation is observed for generalized seizures. Here a distinction is made between tonic/tonic-clonic, and absence seizures. Figure 6.14 shows that the amount of information shared between the past of respiration and the present of RR (i.e. $C_{X\to Y}$ and $T_{X\to Y}$) is significantly reduced in absence seizures. This is an important finding, since this type of seizures do not show a consistent effect on the heart rate [102]. The later is confirmed by the measurements of $S_{Y|X}$, during absence seizures. A slight increase in the residual

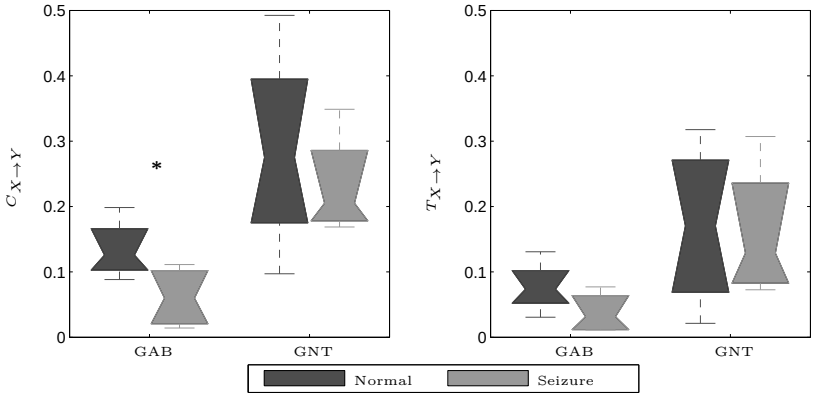

Figure 6.14: *Cross entropy $C_{X\to Y}$ and transfer entropy $T_{X\to Y}$ for regions around absence seizures and generalized tonic/tonic-clonic seizures, compared with values for seizure-free regions. The respiratory signal used in this case corresponds to R_{kpca}. However, the results do not differ significantly when a different respiratory signal is used. The significant difference is indicated by (*).*

information explained by physiological parameters, other than respiration, is observed. However, this increase is not significant, which is in agreement with previous findings reported in the literature [149]. The reason why the respiration plays a significant role in this type of seizures can be explained by the involvement of the thalamocortical network in absence epilepsy [102]. Even though these results seem promising, the example presented in Figure 6.11 needs to be considered. It was clear that the relationship between the real respiratory efforts measured in the abdomen and thorax, and the ECG-derived respiratory signals are "visually" different after the seizure onset. Therefore, future studies should pay special attention to the analysis of respiratory flow, SpO_2, CO_2, in order to get a better view on the effect that absence seizures have on respiration.

Concerning tonic/tonic-clonic seizures, no significant differences were found. This is not a surprise, since this type of seizures are characterized by motor activity, which needs to be considered when recording signals like the ECG. This is to say, in presence of movements the quantification of cardiorespiratory changes becomes unreliable. For this reason, other modalities need to be taken into account to be able to find physiological changes that may lead to this type of seizures.

6.5 Conclusions

This chapter presented the results of the cardiorespiratory analysis on children suffering from epilepsy. It was shown that epilepsy deeply compromises the control mechanisms of the heart rate and respiration, which might be associated with an increased risk of sudden unexplained death in epilepsy. On the one hand, it was observed that children suffering from West syndrome and absence epilepsy have a reduced vagal tone during interictal activity. This is however, not the case in temporal lobe epilepsy, where it seems that patients only experience these autonomic changes during ictal activity. Nevertheless, it was shown that the epileptic children studied in this work, have a more fixed heart rate and a reduced cardiorespiratory coupling, which can compromise their defense mechanisms against asphyxia and hypoxia. These findings might be an indication that a deeper underlying mechanism is taking place in epilepsy, and can be responsible for autonomic dysfunctions leading to SUDEP. However, further research is needed in order to unravel the fundamental cause of cardiac and respiratory dysfunctions in epilepsy.

Apart from quantifying the effect that epilepsy has on autonomic functions, this chapter showed that PRSA, information dynamics, and subspace projections, can be used to quantify interactions between heart rate and respiration. Furthermore, it was shown that they offer complementary information to HRV analysis, in the sense that additional information about the reaction of one of the branches of the ANS, in particular of the parasympathetic branch can be obtained. Despite the promising results presented in this chapter, further investigations on the non-linear interactions between heart rate and respiration need to be performed. This is of paramount importance since these interactions are known to be highly non-linear. However, the linear approximations obtained in this study, already provided new insights into the autonomic effects, and hopefully they will be taken into account for the design of new treatment strategies in epilepsy.

Chapter 7

Detection of epileptic seizures from single-lead ECG

This chapter proposes two novel algorithms to detect epileptic seizures from single-lead ECG. The first algorithm quantifies changes in the morphology of the ECG signal using principal component analysis (PCA), and it detects the seizures using kernel spectral clustering (KSC). The second algorithm, on the other hand, uses (B)PRSA and it extracts different parameters from the curves, which are then clustered using KSC. The development of these two algorithms was inspired by the findings discussed in the previous chapter, and in particular, by the fact that heart rate and respiration display a typical behavior towards the seizure onset. Furthermore, the techniques described in Part I are implemented in such a way that an "online" detection system can be envisioned.

The first part of this chapter gives an introduction to detection systems based on ECG. Next, the two datasets used in this study are described, and the methodology of the detection algorithms is presented. The results obtained with both proposed algorithms are compared and discussed. Finally, conclusions and future work in this direction will be presented.

7.1 Introduction to seizure detection

The control mechanisms of the autonomic nervous system are known to be deeply involved during epileptic seizures [33]. For example, changes in heart rate can occur during and even a few seconds before and after the EEG onset of the seizures [102, 123, 185, 200, 224]. These changes can be caused by either the involvement of the central autonomic centers during seizures, or by motor activity and/or stress responses to the seizures.

Different studies have shown that not only the heart rate, but also the morphology of the ECG changes during the seizure. In [185] it was shown that the QT interval was shortened during the early post-ictal phase in patients suffering from refractory temporal lobe epilepsy. Additionally, ECG abnormalities such as T wave inversion and ST elevation/depression, have been observed in the pre-ictal period of partial and generalized seizures [224]. These studies focused on the presence of ictal tachycardia in adults with refractory epilepsy. In [102] the focus was however, on childhood epilepsy and it was reported that changes in heart rate can be observed in temporal-lobe and frontal-lobe seizures. Generalized seizures present a bigger challenge due to deep motor involvement and/or a lack of significant changes in heart rate. These findings were obtained by assuming that seizures are characterized by a linear and a quadratic behavior in the HR or RR interval time series [200, 123, 102]. An illustration of this can be observed in Figure 7.1. This is of course valid in the presence of tachycardia or bradycardia, and when no ECG artifacts affect the detection of the R-peaks.

A different effect of epilepsy was observed in [149], where it was reported that epileptic seizures can be accompanied by apnea episodes, especially when the onset is localized in one of the temporal lobes. Moreover, as it was shown in the previous chapter (section 6.4), the cardiorespiratory interactions have a particular behavior towards the onset of partial and generalized seizures. Under these circumstances, it is possible to envision a detection system based on single-lead ECG that can improve the quality of life of patients suffering from epileptic seizures. This improvement in quality of life can be achieved due to the fact that the ECG has an inherent advantage over EEG in the sense that it is easier to record in a long-term ambulatory setting, hence, more portable systems can be developed. In addition, changes can be observed in the ECG even before the EEG onset, making it possible to early-detect the onset of a seizure. This last point is very important for the development of "closed-loop" systems like vagus nerve stimulation, where the ultimate goal is to stop seizures from occurring and prevent SUDEP [199].

In order to contribute to the development of seizure detection systems based

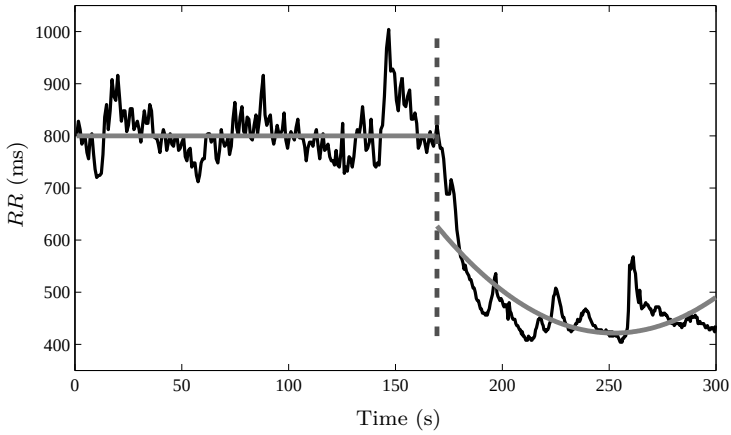

Figure 7.1: *Pattern of heart rate changes caused by a temporal lobe seizure. Note that a linear, and a quadratic curve are fit on the RR interval time series. This curve fitting was proposed in [123] to detect the breaking point indicated by the dashed line, which corresponds to the onset of the seizure.*

on more portable sensors, this chapter presents a methodology that uses only the ECG signal to detect epileptic seizures. This methodology proposes a novel feature that describes the changes of the ECG morphology by means of principal component analysis (PCA). This new feature allows to identify changes of the morphology of the ECG, in particular of the QRS-complex, that can be caused by seizures. Furthermore, PRSA and BPRSA are used to quantify changes in the cardiorespiratory interactions.

Different types of seizures are considered, namely, focal seizures with onset in the frontal lobe or in one of the temporal lobes, and generalized seizures. In addition, both childhood and adult epilepsy are considered. The biggest challenge is to detect generalized seizures, especially those with deep motor involvement, and with no obvious effect in the heart rate.

7.2 Datasets

In this study, two ECG datasets were used. The first one was used as a training set, and it is the same one used in the study presented in section 6.4. It consists of single-lead ECG recordings from children suffering from refractory epilepsy. These children were admitted to the epilepsy clinic of UZ Leuven in Belgium, and were monitored using 24-hour video-EEG. ECG and EEG

signals, amongst others, were measured in parallel. In total, 37 patients were monitored and 98 seizures were identified. The mean age of the patients was 9.2 years (range 3-16 years). The onset of the seizures was annotated based on EEG and video, by two different EEG specialists. Two types of seizures were recorded: 48 with focal onset and 50 with generalized onset. Twenty of the focal seizures originated from the temporal lobe (TLE) and the other 28 from the frontal lobe (FLE). The generalized seizures were 29 tonic or tonic-clonic (GTN), 11 myoclonic (GMY), and 10 generalized absence seizures (GAB). Once the seizures were identified, a 5 min ECG segment was extracted, starting 3 minutes before the EEG onset of the seizure. The ECG signals were recorded at a sampling frequency of 250Hz.

The second dataset used in this study was the *Post-Ictal Heart Rate Oscillations in Partial Epilepsy*[5], which contains 10 partial seizures and is publicly available in the Physionet *databank* [82]. This set consists of 7 ECG recordings, extracted from video-EEG monitoring of 5 women whose age ranged from 31 to 48 years, and who were suffering from epilepsy. All the seizures contained in this dataset were of focal nature with or without secondary generalization. The ECG signals were sampled at 200 Hz, and their duration ranged from 1h30m up to 3h45m. This dataset was used in order to test the algorithm on long term ECG recordings.

7.3 Methodology

The methodology implemented in this study consists of three different phases. Firstly, the ECG signals were pre-processed and the R-peaks were identified. Secondly, different signals and features were extracted from the ECG. Lastly, supervised and unsupervised learning techniques were implemented in order to detect the onset of the seizures.

7.3.1 Preprocessing and R-peak detection

Initially, the Physionet dataset was up-sampled at 250Hz in order to comply with the lower limit of the recommended frequency range in HRV [129], and to have both datasets sampled at the same frequency. After that, the power line interference in the ECG was removed using a notch filter at 50Hz for the data collected in the UZ Leuven, and at 60Hz for the Physionet dataset. Next, the R-peaks were detected using using Algorithm 7, described in Section 4.2. Finally, visual verification of the position of the R-peaks was performed on the 5-minute segments. This manual verification was done in order guarantee the

proof of concept, and that the results obtained at this point were not corrupted by artifacts. The online version of this approach should include the automatic detection and removal of segments with artifacts using Algorithm 4, proposed in Section 4.1.2.

Once all R-peaks were detected, and the RR time series derived from the time differences between consecutive R-peaks, the three EDR signals, namely, R_{amp}, R_{pca}, and R_{kpca} were computed using the methodologies described in section 4.4.

7.3.2 Continuous quantification of ECG morphology changes

After the R-peak detection, 5 signals were computed using the procedure described below, and outlined in Figure 7.2. This procedure was inspired by the derivation of the respiratory signal from the ECG using PCA [119], and is one of the contributions of this research.

a) A window of 60ms (m samples) before and after each R-peak was selected in order to capture the QRS-complexes.

b) Five ($n = 5$) consecutive QRS complexes were assembled into one QRS-matrix $X \in \mathbb{R}^{n \times m}$, and aligned with respect to the R-peaks. The definition of n was done experimentally, as the lowest amount of heart beats that lead to a better differentiation between normal and pre-ictal/ictal activity.

c) Principal component analysis (PCA) was performed on X, and the five non-zero eigenvalues were selected as the new set of parameters.

d) Shift one heart beat and repeat the algorithm until the end of the ECG segment.

After this procedure, the evolution of the five non-zero eigenvalues was obtained with a resolution of 5 heart beats, and a vector of 5 features was used to characterized each window shift. This feature vector for the ith window corresponds to $\{\lambda_{i,j}\}_{j=1}^{5}$.

7.3.3 Continuous measurement of cardiorespiratory interactions

The second algorithm proposed in this research to characterize the ECG, and perform seizure detection, quantifies cardiorespiratory interactions using

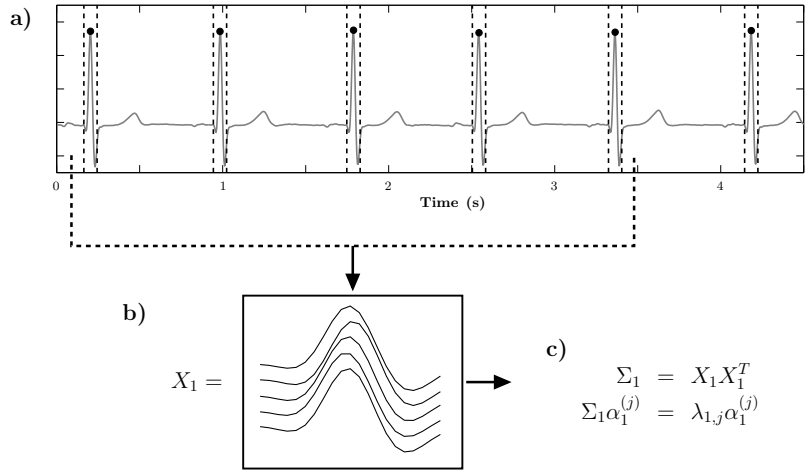

Figure 7.2: *Procedure to construct the QRS-matrix, and derive the five non-zero eigenvalues $\lambda_{1,j}$ of the matrix X_1, where $j = 1, \ldots, 5$. The window of 5 heart beats is shifted throughout the whole ECG segment and the evolution of the eigenvalues $\lambda_{i,j}$ is computed, where $i = 1, \ldots, N$ and $N+4$ the total number of heart beats in the segment.*

phase rectified signal averaging (PRSA)[1]. In order to simulate a "continuous" detection of epileptic seizures, PRSA was implemented using a moving window of 80 heart beats. The length of the window was determined experimentally, as the length that allowed to better differentiate between normal and pre-ictal/ictal activity. Each pair of signals was then used to extract two different PRSA curves, one where the RR was used as trigger and target simultaneously, and one where the EDR triggers the RR. This procedure was repeated for every window shift, in steps of one heart beat with an overlap of 79 beats (see Figure 7.3).

Phase rectified signal averaging (PRSA)

Phase rectified signal averaging (PRSA) [22] is able to detect quasi-periodicities in non-stationary signals like the RR, in the presence of noise and artefacts. This is done after reducing a given time series $\mathbf{Y} = \{Y_i\}_{i=1}^{N}$ to a shorter sequence $\mathbf{Y'} = \{Y'_j\}_{j=-L}^{L}$, with $Y, Y' \in \mathbb{R}$, and $L < N$. The shorter sequence $\mathbf{Y'}$ of length

[1] See chapter 5 for details on this methodology

$2L + 1$ is computed following the three steps represented in Figure 7.3[2], and described as follows.

1. The anchor points (AP) were selected from the driver signal $\mathbf{X}$ based on the conditions $X_i < X_{i-1}$ and $X_i > X_{i-1}$. In the univariate version of PRSA, $\mathbf{X} = \mathbf{Y}$, with $\mathbf{Y} = RR$, and in BPRSA $\mathbf{X}$ corresponds to the respiratory signal.

2. The target signal, RR, was segmented using a window of length $2L + 1$ around each AP. In total, M windows were segmented.

3. An average curve $\mathbf{Y}'$ was computed over all M windows.

Different parameters were used to quantify the (B)PRSA curves [93]. These parameters correspond to the slopes of the lines connecting the point before and after the AP in the (B)PRSA curves. These features will be denoted by S_{RR} and S_R for the PRSA derived using only the RR and using both the EDR and the RR, respectively. These slopes will give an indication on how fast the heart reacts to acceleration or deceleration, and how the inspiration/expiration cycle is related to changes in the HRV. In this study, it is hypothesized that during epileptic seizures, these reaction times will be affected due to sympathetic activations, or due to the involvement of the cardiorespiratory controls during the seizure. Besides, it is expected that the coherence times within the signal change due to the significant changes in HRV that might take place in the pre-ictal/ictal period. As discussed in section 5.1 the coherence time can be estimated from the slope of the line connecting points at a certain distance from the AP in the PRSA and BPRSA curves[3]. These slopes will be denoted by Δ_{RR} and Δ_R, respectively. Summarizing, for a moving window of 80 heart beats with an overlap of 79 beats and steps of 1 heart beat, the ith window was characterized by the following feature vector: $[S_{i,RR}, S_{i,R}, \Delta_{i,RR}, \Delta_{i,R}]$, with $i + 80$ the total number of heart beats in the ECG segment.

In total, 11 features were extracted from the ECG, RR and R signals, and they are summarized in Table 7.1.

7.3.4 Seizure detection

Once each window was characterized by the set of 11 features (see Table 7.1), two different unsupervised algorithms were implemented to split the data into

[2]These steps were described in Section 5.1, Figures 5.1 and 5.2
[3]For instance the slope connecting the points Y'_{-k} and Y'_k

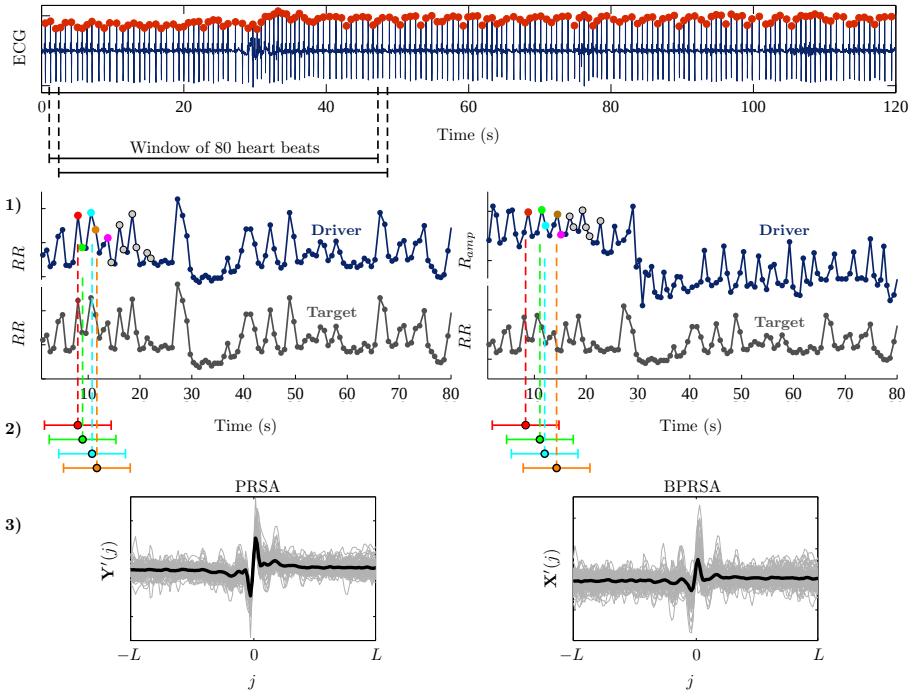

Figure 7.3: *Computation of the (B)PRSA curves. A window of 80 heart beats in the ECG is used for the illustration. From each window, two signals are computed namely RR and R_{amp}. 1) Selection of anchor points in the driver signal RR (left) and R_{amp} (right). The first 12 AP are indicated. 2) Segmentation of the target signal RR, using a window of length $2L$ around each anchor point. 3) Phase rectification and signal averaging of all extracted segments around AP. The average (B)PRSA curves are indicated in black.*

different groups. These algorithms were k-means clustering, and kernel spectral clustering (KSC) [10]. The reason to use unsupervised learning techniques relies on the fact that the labels (i.e. seizure onsets) are only known for the EEG signal, and early or late responses in the ECG are expected.

For k-means the number of clusters was selected using the average Silhouette value [107], and for KSC using the balanced line fit (BLF) proposed in [10].

Additionally, a classifier based on thresholding was used. These thresholds were defined as the 25th and 75th percentiles of each feature during the first 2 minutes, which were considered to be free of seizures.

Table 7.1: *Set of features derived from the ECG, RR interval time series, and respiratory signal R corresponding to any of the EDR signals.*

Method	Feature ID	Derived from		
		ECG	RR	EDR
mean	$\bar{RR}$		X	
SDNN	$\sigma(RR)$		X	
Morphology	$[\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5]$	X		
PRSA	$[S_{RR}, \Delta_{RR}]$		X	
BPRSA	$[S_R, \Delta_R]$		X	X

7.4 Results and discussion

After preprocessing the 5 minute ECG segments and detecting the R-peaks, 11 segments were identified as containing artifacts around the seizure onset (i.e. 180s). Of these 11 segments, 9 correspond to the generalized group with motor involvement, where artifacts caused by muscle movements are observed. Manual correction of these segments was performed, but for most of the cases it was not possible to assess the location of the R-peaks[4]. Hence, movement artifacts must be removed from the ECG before any seizure analysis.

7.4.1 Detection using morphology information

After finding the position of the R-peaks, the eigenvalue signals and the RR interval time series were computed for each ECG segment separately. Remember that each point of the eigenvalue signals was obtained using the previous 5 heart beats. Figure 7.4 shows an example of these 6 signals for a given segment. Note that the onset of the seizure, around 180s is clearly visible in the RR series and in the different eigenvalues. However, changes around the onset of the seizure (dashed line) are more visible in RR and λ_{3-5}. The reason for this is that the first components (λ_1, λ_2) are more affected by the respiratory modulation, as discussed in [119], and the other components become more relevant when the morphology of the QRS complex becomes more heterogeneous. In other words, most of the variance of the QRS complexes cannot be described by only the first two components.

[4]All experiments were carried out in MATLAB R2012a on an Intel® Core™ i7, 3.4GHz, 7.8 GB RAM, running Ubuntu 12.04 LTS.

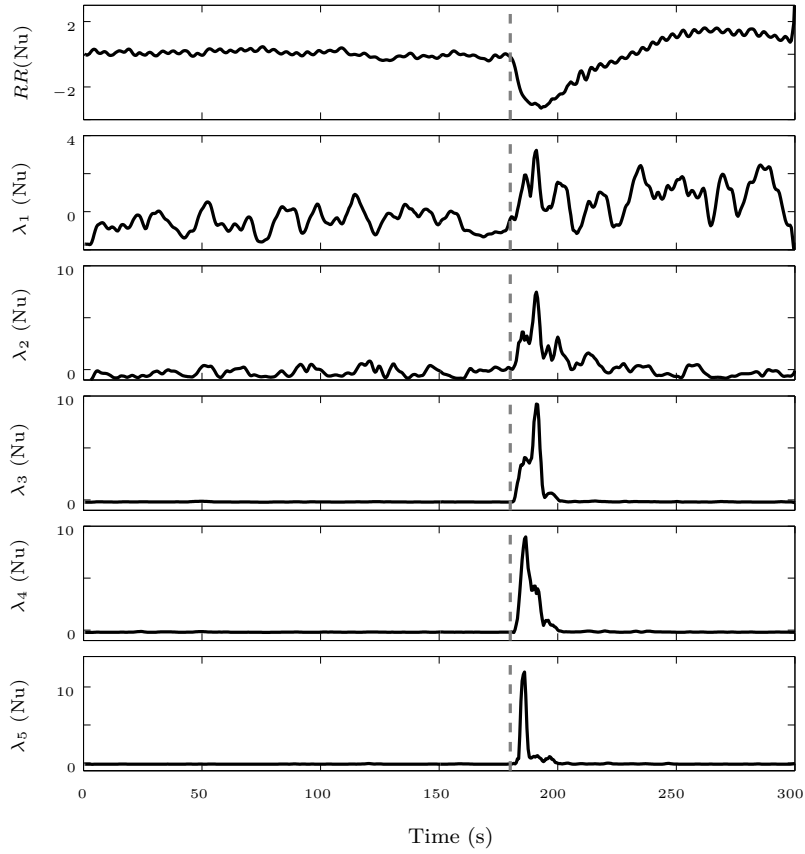

Figure 7.4: *Signals derived from the ECG after R-peak detection. The first one corresponds to the RR interval time series, and the other 5 indicate the evolution of the eigenvalues throughout the ECG segment. The dashed lines indicate the EEG onset of the seizure at 180s.*

The extracted signals were then segmented into epochs of 5 seconds, and each epoch i was characterized by the average values of all features, namely, $[\bar{RR}_i; \sigma(\bar{RR})_i; \bar{\lambda}_{i1}; \bar{\lambda}_{i2}; \bar{\lambda}_{i3}; \bar{\lambda}_{i4}; \bar{\lambda}_{i5}]$. These epochs were first classified as being *normal* or containing a *seizure* by means of the threshold defined as the 75th percentile of the first 2 minutes in all training segments. This is illustrated in Figure 7.5, where the dashed line indicates the 75th percentile of all training segments from segment 1 to segment 36. The training set consisted of segments of 40 patients, 20 suffering from focal seizures and 20 suffering from generalized seizures, selected using FS2 summarized in Algorithm 2. Figure 7.5 shows this

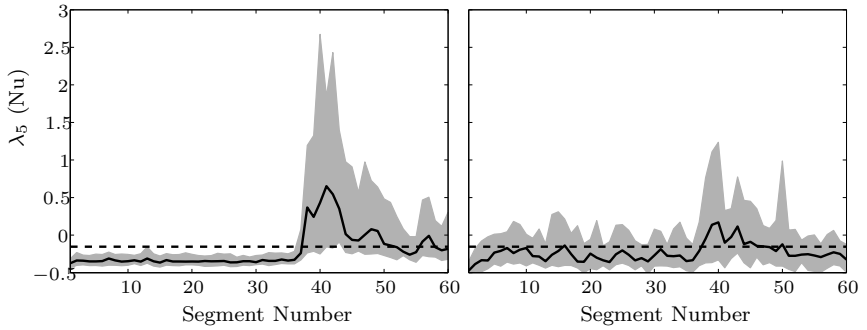

Figure 7.5: *Threshold for λ_5. Nu stands for normalized units. The shaded area indicates the 75th and 25th percentile per segment of 5 seconds. The median values are indicated by the solid line, and the dashed line corresponds to the threshold used for classification. A distinction between partial (left) and generalized (right) seizures is done. Note that in partial seizures, most of the segments around 200s (segment 40) lie above the threshold.*

threshold for λ_5, together with the 25th, 50th, and 75th percentiles of all epochs. A distinction is made between different types of seizures, where it is clear that the partial seizures are easier to detect by means of a simple threshold.

In addition to the classifier, k-means and KSC were implemented, and the optimal number of clusters was found to be 2. Each cluster was then manually characterized as containing seizure onsets or normal segments. Figure 7.6 indicates an example of cluster identifiers for each approach. It is observed that cluster 2 can be identified as containing seizure activity. Practically, once the clusters are characterized, new data points can be associated to one of the clusters by means of a distance measure from the point to the centroids of the clusters in the case of k-means. In this research, the Euclidean distances between the test segments and the centroids obtained with k-means were used to associate the segments to the cluster with the shorter distance. On the other hand, the out-of-sample extensions of KSC [10] were used to associate the test segments to either cluster. Table 7.2 indicates the results for the entire dataset, for each approach and each type of seizure.

As mentioned before, not all the eigenvalue signals indicate a clear change around the seizures. Therefore, the three different seizure detection approaches, namely thresholding, k-means and KSC, were implemented on different sets of features, and the best results (i.e. larger PPV) were obtained with the following parameters: $[\bar{RR}; \bar{\lambda_3}; \bar{\lambda_4}; \bar{\lambda_5}]$. Since the seizure onset annotations were done on the EEG signal, and it is known that cardiac changes can be observed before or after these annotations, a window of 60s around the annotated seizure onset

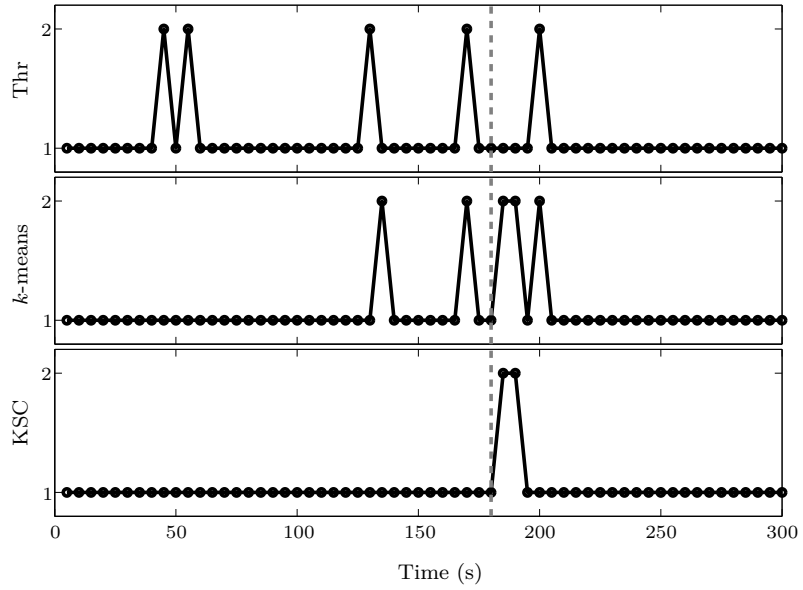

Figure 7.6: *Labels for each segment of 5 seconds after classification using threshold (top), clustering using k-means (middle), and KSC (bottom). The dashed line indicates the EEG onset. The cluster identifiers are indicated in the ordinate, where it is clear that cluster number 2 can be associated to seizure onsets.*

was used to calculate the number of false positives. In other words, if the clustering algorithm produced a "seizure" result within 30s before or 30s after the annotated seizure onset, the result was considered a true positive. If on the other hand, the result was outside this 60s window, the result was labeled as false positive. This is illustrated in Figure 7.7, where the entire procedure to detect seizures is indicated.

From Table 7.2 it is clear that kernel spectral clustering (KSC) detects 85% of the partial seizures, and only 57% of the generalized ones. If all the segments corresponding to generalized seizures that contain EMG artifacts are removed (24 out of 50), the PPV increases to 87%. Keep in mind that some of the generalized seizures used in this study involve motor activity, which implies the presence of EMG artifacts on the ECG. These artifacts reduce the reliability of the R-peak detection, and seizure detection algorithms. For this reason, the ECG should be cleaned from EMG artifacts, and different features derived from signals like EMG and/or accelerometers should be used to complement this study.

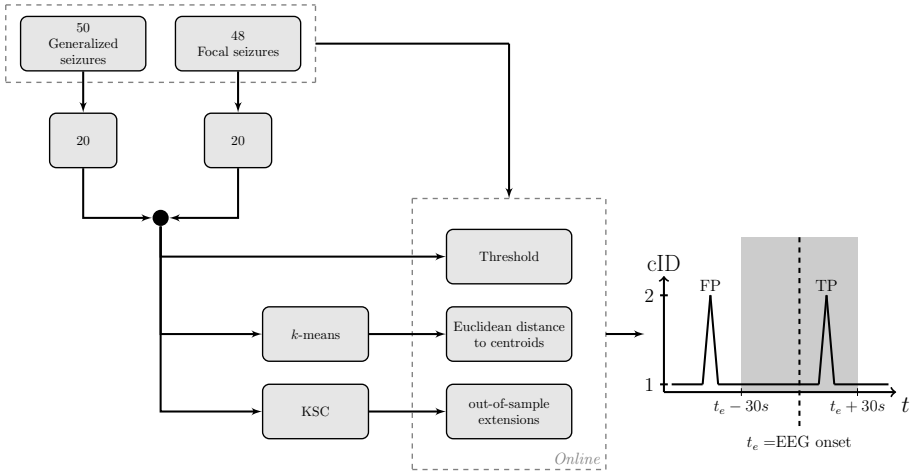

Figure 7.7: *General procedure to detect epileptic seizures. The dataset is first split into a training set of 40 seizures selected using FS2. Then, the threshold is defined, and the k-means and KSC algorithms are trained. Next, the whole dataset is associated to one of the clusters defined by the three different approachees, and the clusters are characterized as containing normal or seizure activity. The last step is to compute false positives and false negatives using a window of 60s around the EEG onset.*

Taking the results of KSC, for both types of seizures, about 75% of the changes in the ECG were lagging the onset in the EEG, while the rest was leading or simultaneous. The detection window for partial seizures was 7.4±14.3 seconds, and for generalized 9.6±20.5 seconds. These results refer to the fact that cardiac changes can be observed before EEG changes, and as mentioned before, a window of 60s around the EEG annotation was used to calculate the true positives. Hence, these latency times were obtained for each type of seizure separately.

These results are in accordance with previous studies [123] where the PPV for partial seizures was larger than 80%. The methodologies used in [123, 102] studied the full ECG segment in order to find the best curve fitting, and the exact seizure onset. However, the methodology proposed in this chapter uses a window of 5 seconds, hence, the detection of the seizure will be done with at least 5s delay.

The previous findings were obtained from the UZ Leuven dataset. Hence, the Physionet data was used to test these algorithms on long-term ECG signals. The performances achieved for this dataset are also included in Table

Table 7.2: *Seizure detection results using morphology changes in the ECG.*

Dataset	Seizures	Algorithm	TP	FP	PPV	Sens
UZ Leuven	Partial	Threshold	44	13	77%	91.6%
		k-means	43	18	70.5%	89.5%
		KSC	**43**	**7**	**85.7%**	**89.5%**
	Generalized	Threshold	40	39	50.6%	80%
		k-means	47	52	47.5%	94%
		KSC	43	32	57.3%	86%
Physionet	Partial	Threshold	10	14	41.6%	100%
		k-means	9	20	31%	90%
		KSC	10	9	52.63%	100%

7.2. Remember that this set only included data from adults, suffering from partial seizures. Therefore, it is hypothesized that the results are strongly biased by the typical morphological behavior observed in children. For this reason, the algorithm performs worse in this kind of dataset. Unfortunately, no additional dataset of adults was available at the moment that this study was performed. Future investigations must evaluate these morphological differences, and determine the age dependencies for the features proposed in this chapter. Nevertheless, it is remarkable to see that in about 17h of recordings only 9 false positives were detected when using KSC. These results can be improved after post-processing the cluster indicator, and after using a more appropriate training set.

7.4.2 Detection using (B)PRSA

The dataset collected in UZ Leuven, was first used to train the learning algorithm, and to determine the features that best differentiate between normal and pre-ictal/ictal activity. This was also done for all the EDR signals, namely $[R_{amp}, R_{pca}, R_{kpca}]$, and it was found that no significant differences exist between the results produced by either signal. For this reason, and taking into account the complexity of the algorithms, the R_{amp} signal was preferred.

Each ECG segment of 5 minutes was analyzed using a moving window of 80 heart beats as described above, and two PRSA curves were obtained for each heart beat. Both acceleration or decreasing points in RR, and deceleration or increasing points in RR were used, but the best results were obtained using acceleration points. This can be due to the fact that the most pronounced

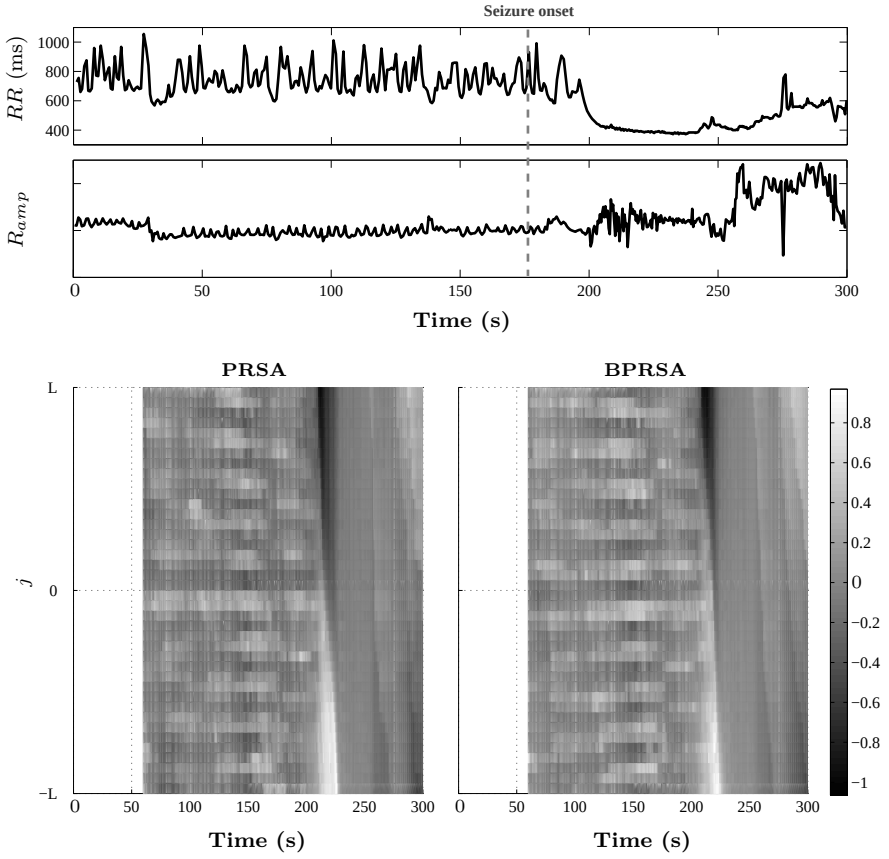

Figure 7.8: *(top) RR and R_{amp} signals. Note that the heart rate and respiration are significantly perturbed around the seizure onset. The (B)PRSA curves are indicated as vertical lines in the lower panels. It is clear that the perturbations of the heart rate and respiration are captured by the dynamics of the (B)PRSA curves around the EEG onset. Nevertheless, they are detected about 20s after the seizure onset.*

reaction of the heart rate towards the seizure is pre-ictal tachycardia. The parameter L was set to 20. This value was selected experimentally, as the length that allows to see changes in coherence times within the signal. Fig. 7.8 shows an example of a focal seizure originated from the temporal lobe, where clear changes in heart rate can be observed. These changes in temporal lobe seizures have already been reported in the literature [149][101]. However, in this study changes in cardiorespiratory coupling around the seizures, were also observed and quantified.

One PRSA curve was computed for each new heart beat, and as can be seen from the bottom panels of Figure 7.8, the shape of the curves changes around the seizure onset. Note that the time reversal symmetry of the curves around the anchor points ($j = 0$) is completely broken as the variability of the heart rate (HRV) is reduced. These changes in symmetry cannot be detected by any power spectral analysis [106], which makes PRSA a very powerful technique in this application. An important fact here is that the coherence time of the signal also changes. For instance, when the RR starts to decrease dramatically, and the analyzed segment includes both pre-ictal and ictal periods, the predictive power of the signal is reduced. These changes are linked to the reduced standard deviation of the HR during seizures [101, 149], which on its turn can be explained by the propagation of the seizures to the central autonomic network [101]. These changes in coherence time and symmetry, characteristic during seizures, can be quantified using the slope of the line connecting the furthest points to the AP of both PRSA curves, $\Delta_{RR} = (Y'_L - Y'_{-L})/(2L+1)$, and $\Delta_R = (X'_L - X'_{-L})/(2L+1)$, for PRSA and BPRSA respectively. The panels in the 3rd and 4th rows of Figure 7.9 show these changes. In addition, this can also be quantified using the first point of the PRSA curves, namely, Y'^k_{-L}, with $k = 1, 2, \ldots, M$, and M the total number of windows of 80 heart beats. This parameter indicates if the quasi-periodicities have already decayed or not when looking L heart beats in the past of the averaged AP. For normal activity, the values of Y'^k_{-L} should be close to zero as seen in Figure 7.9, 5th row.

Another feature that varies around the seizure onset is the slope S_{RR} of the line connecting the point before Y'_{-1} and the point after Y'_1 the AP of the PRSA derived using RR as target and trigger signals. Figure 7.9, 2nd row, shows how this slope becomes steeper during the occurrence of the seizure. The reason for this is, again, the reduced variability of the heart rate during this period, and a probable increased sympathetic activation. However, the mechanisms involved in the acceleration of the heart are not yet fully understood, hence, this effect cannot be associated solely to sympathetic activation [21].

The four features described above, namely $[S_{RR}, \Delta_{RR}, \Delta_R, Y'_{-1}]$ were used to characterize each heart beat, and k-means and KSC were then implemented to group them into clusters containing similar activity. The algorithms were trained using 40 seizures, 20 focal and 20 generalized (no myoclonic included) that were selected using the modified fixed-size algorithm (FS2) discussed in section 2.3.2. Keep in mind that KSC is an unsupervised learning algorithm, so no labels were used in training and validation. The optimal number of clusters and the optimal kernel parameter were selected using the balance line fit criterion described in [10]. The optimal number of clusters was found to be 3 using the BLF, and after analyzing each cluster separately, it was found

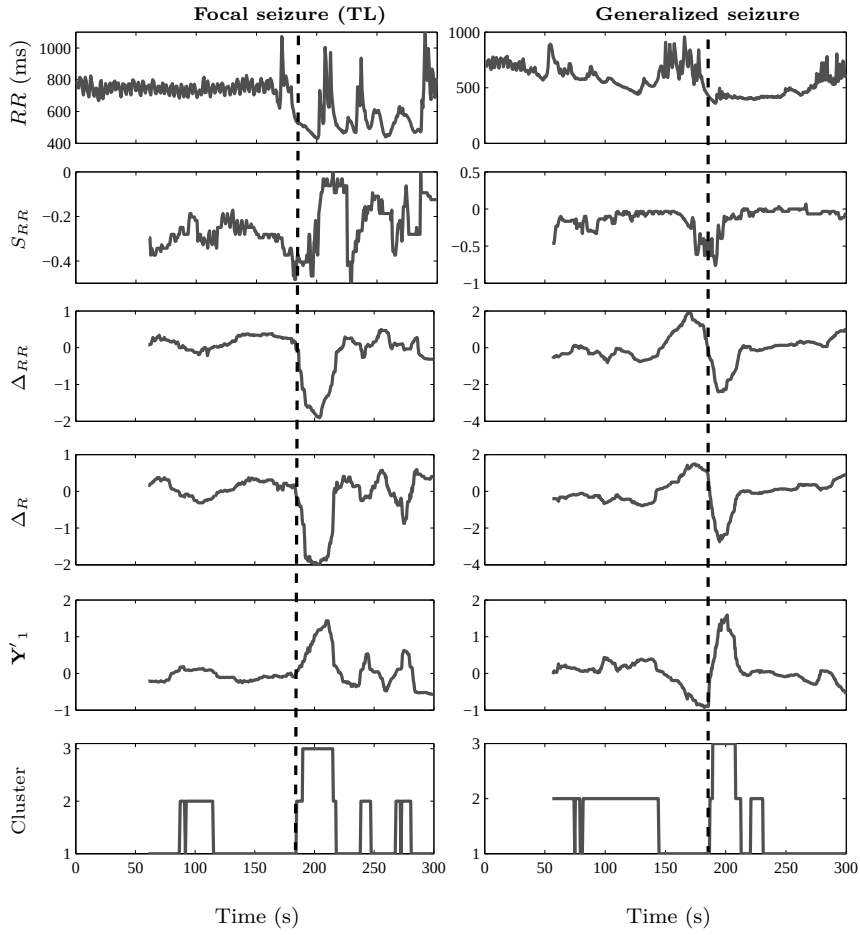

Figure 7.9: *Features extracted from the (B)PRSA curves of a focal seizure originated in the left temporal lobe (left), and a generalized seizure (right). From top to bottom: RR interval time series; slope S_{RR} of the line connecting the point before Y'_{-1} and the point after Y'_1 the AP of $\mathbf{Y}'$; Difference between the furthest points to the AP of both (B)PRSA curves, Δ_{RR}, and Δ_R; Variation of $Y_1'^k$ as a function of time; and cluster memberships indicated in the ordinate, where cluster 3 contains the epileptic seizures.*

that cluster number 3 contained most of the seizures of the whole set. In k-means the optimal number of clusters was found to be 2 using the silhouette

Table 7.3: *Seizure detection results using (B)PRSA.*

Dataset	Seizures	Algorithm	TP	FP	PPV	Sens
UZ Leuven	Partial	k-means	45	10	81.8%	93.7%
		KSC	48	5	**90.5%**	**100%**
	Generalized[a]	k-means	37	23	61.6%	94.8%
		KSC	36	7	**83.7%**	**92.3%**
Physionet	Partial	k-means	10	7	58.8%	100%
		KSC	10	4	**71.4%**	**100%**

[a]Myoclonic seizures are not included

value. The panels of the last row in Fig. 7.9 show the cluster membership with KSC as a function of time, and Table 7.3 indicates the positive predictive value (PPV) and sensitivity for each type of seizure and each algorithm. In order to determine the amount of TP, a window of 30s before and 30s after the EEG onset was used. This was the same procedure used in the previous approach, and indicated in Figure 7.7.

It is not a surprise that the results for generalized seizures are affected by myoclonic seizures, since they are brief, normally lasting no more than 2 seconds, and more importantly, they manifest themselves as rapid contraction/relaxation of a muscle. This implies motor involvement and very short changes (<2s). For this reason, the detection of these seizures remains a challenge. However, 9 of the 11 myoclonic seizures were detected, but 6 false positives were also included (PPV=60%). If these seizures are included in the generalized results, the total PPV was reduced to 76.5%. On the other hand, as mentioned in the previous approach, 24 seizures were contaminated by EMG artifacts. Hence, if these seizures are removed from the set, the PPV increases to 89% and *all* the generalized seizures are detected. Again, as stated before, other modalities like EMG or accelerometer data should be included in the analysis of motor seizures. The results obtained using (B)PRSA were compared with those reported using the morphology of the ECG (see Table 7.2). In this case, an improved PPV was achieved for both types of seizures, but what is more remarkable is that the algorithm that uses KSC, reduces the amount of false positives by more than 50% for the generalized seizures.

After training and validating the clustering algorithms using the UZ Leuven, the resulting models were applied to the Physionet dataset. The PPV and the sensitivity obtained for these partial seizures in adult epilepsy are indicated in Table 7.3. Note that these models were trained on a completely different

dataset, and still the KSC algorithm is capable to detect all partial seizures, and produce a relatively high PPV. A total PPV for partial seizures, including both the UZ Leuven and Physionet datasets is 86.56%, with sensitivity of 100%.

In order to estimate the latency of the detection, the results obtained using KSC were used. It was found that most of the seizures were detected more than 20 seconds after the seizure onset in the EEG. These delays are caused by the fact that a window of 80 heart beats is used in this approach. In addition, when compared with the previous study, where a window of only 5 seconds was used, it is clear that a compromise must be made between number of FP and latency of the detection. This will always depend on the application. For instance, less FP will be preferred for seizure diaries, follow-up of treatment strategies, amongst others. On the other hand, closed-loop systems like VNS could benefit from methodologies that offer low-latency detection.

7.4.3 Unifying seizure detection algorithm

At this point it is clear that the features that produce the best results in the detection of epileptic seizures are: $[\bar{RR}; \bar{\lambda_3}; \bar{\lambda_4}; \bar{\lambda_5}; S_{RR}; \Delta_{RR}; \Delta_R; \mathbf{X'}_1]$. However, when all these features were combined and used to cluster the data using KSC, no improvement was achieved in the detection of seizures.

7.5 Conclusions

This chapter presented two different algorithms for the detection of epileptic seizures from a single-lead ECG. The first one captured changes in the morphology of the QRS complexes by means of principal component analysis, and the second one quantified by means of (B)PRSA, the acceleration of the heart caused by an epileptic seizure. Both algorithms detected partial epileptic seizures with a PPV larger than 80%, in a window of 30s around the EEG onset. For generalized seizures, however, the performance went down, in particular when myoclonic seizures were included. Nevertheless, the approach based on PRSA allowed to detect generalized seizures with a PPV of 83%, which until now was not achieved by any other algorithm based solely on ECG analysis. Consequently, it is necessary to keep in mind that the ECG needs to be cleaned from EMG artefacts in order to be able to detect seizures with motor involvement, as it is the case for the generalized seizures used in this work and some partial seizures involving the frontal lobe.

A clear advantage of these algorithms is that they are easy to implement and can be adapted to home monitoring or online detection systems. These findings represent an important contribution to the development of closed-loop systems, where the aim is to abort epileptic seizures. However, training and validating the algorithm must be considered for each type of seizures, and to include different sources of data, to improve the performance for seizures with motor involvement. In other words, better performances can be expected if these algorithms are used to complement other approaches such as those based on accelerometers and EMG.

Another important consideration to be made is the fact that these algorithms can be used in two completely different approaches. For instance, for seizure detection, where it is required to detect seizures as soon as possible, the algorithm based on morphology can be considered. On the other hand, for diagnosis or follow-up of therapies an algorithm like the one based on (B)PRSA that has a higher latency of detection but a lower false positive rate can be considered, keeping in mind the considerations on motor seizures described above. In fact, they can both be considered for the development of an automatic report, where only the "medical" events are highlighted for the doctors to analyze. This will reduce the amount of time that is currently spent on manually annotating the occurrence of epileptic seizures.

Chapter 8

Sleep apnea detection from single-lead ECG

This chapter presents the last contribution of this research, which is a methodology based for the detection of sleep apnea. This methodology is based on the algorithms discussed in the first part of this book. In addition, it consists of two novel features to capture information from single-lead ECG. One feature builds upon the EDR algorithm based on PCA discussed in Section 4.4, and one uses the subspace projections approach discussed in Section 5.3. The latter is used to quantify the information shared between respiration and heart rate during apnea episodes. This chapter also reports on the added value of each of the three EDR methodologies described in Section 4.4 in the detection of sleep apnea. Furthermore, the evaluation of the different supervised classification techniques discussed in Section 2.1 will be performed.

Section 8.1 will give a brief introduction to sleep apnea detection from single-lead ECG signals. In Section 8.2 the proposed methodology will be described, and its performance will be evaluated on two different datasets. Results of this evaluation will be reported in Section 8.3, and discussed in Section 8.4. Concluding remarks will be presented in Section 8.5.

8.1 Introduction to sleep apnea detection

Sleep apnea is a sleep-related breathing disorder that affects about 10% of middle aged adults [159]. In addition, it is considered a risk factor for morbidity and mortality due to its long-term effect on the cardiovascular system [45]. This effect is related to different physiological mechanisms like systemic hypertension and increased sympathetic modulation that in a long term compromises the well-functioning of the heart.

Apnea and hypopnea are two types of respiratory events characterized by a complete absence or reduction of airflow, respectively, during at least 10s [147]. During these events the respiratory effort can differ, and this on its turn can be used to distinguish between different types of apnea. For instance, in obstructive sleep apnea (OSA) the patients stop breathing, despite the respiratory effort, due to an obstruction in the upper airways that avoids air entering the lungs. In central sleep apnea (CEN), patients also stop breathing, but in this case due to absence of respiratory effort. The combination of these two apnea types is called mixed apnea (MIX), and it consists of a reduction of respiratory effort that leads to an upper airway obstruction. Hypopneas are characterized by a drop in oxygen saturation caused by a reduction in the airflow. In addition, no clear ECG changes can be observed in the majority of the cases. As a result, hypopneas are the most challenging type of respiratory events to be detected solely from the ECG signal. In these cases other signals like oxygen saturation and airflow can complement the ECG analysis. Apneas and hypopneas can occur up to hundreds of times per night, and its recurrence can cause hypersomnolence, excessive daytime sleepiness, insomnia, nocturia, memory loss, attention deficit, and depression among others.

Currently, sleep-related breathing disorders are diagnosed using poly(somno)-graphy, which is a sleep test performed in a hospital or at home, and where the evaluation and supervision of a clinician is essential. Even though poly(somno)graphy is the most important diagnostic tool in sleep medicine, its high costs and low comfort significantly reduce its diagnostic power. Therefore, the development of more portable and less intrusive systems is of paramount importance.

During poly(somno)graphy several physiological parameters are recorded such as heart rate and respiration, which are deeply affected during episodes of apnea. Hence, several studies have focused on the detection of sleep apnea using as few signals as possible. For example, [38, 54, 108, 157] proposed to derive several features from the electrocardiographic (ECG) signal to detect sleep apnea. Moreover, instead of measuring the respiratory signal, those algorithms first estimate the respiration from the ECG, and then derive a set of features

that can later be used for the discrimination of apneas.

In this context, this study proposes a simple algorithm for the automatic detection of sleep apnea from the ECG signal. In addition, two different features that capture valuable information related to apneas are proposed. One describes changes in the morphology of the ECG by means of principal component analysis, and one computes the information shared between respiration and heart rate, by means of orthogonal subspace projections. Changes in the morphology of the ECG associated with obstructive sleep apnea events were first reported in [156], where it was concluded that algorithms that take into account changes in morphology perform better for the classification of apnea. Later in [34], changes in the P-waves were used to discriminate between normal ECG segments, and segments containing apneas. In this work, however, a novel methodology based on principal component analysis applied to the QRS complexes is proposed. Apart from these two features, this study compares three different algorithms to compute the ECG derived respiration (EDR). The first EDR algorithm is based on the R-peak amplitude [54], which has been widely used for apnea detection, the second one uses principal component analysis [119], and the last one uses its kernel version kPCA [217]. The last two algorithms have not been applied yet to the detection of sleep apnea, and this study aims to determine their added value in this type of applications. All the features derived from the heart rate and from the EDR signals are used as input to different classifiers, namely, linear discriminant analysis (LDA), support vector machines (SVM), and least-squares support vector machines (LS-SVM). Variations of these three classifiers are implemented and compared in this work. A final algorithm is then proposed and tested on two different datasets. An important addition of the proposed methodology is that a level of confidence is provided for each classified segment. This level is associated with the level of contamination of the ECG, which is used to derive all features.

8.2 Methods

8.2.1 Datasets

Two different datasets were used in this study. The first one is the Apnea-ECG database [157], publicly available[1] on Physionet [82]. This dataset contains 70 single-lead ECG signals sampled at 100Hz, of lengths ranging between 7h and 10h. Each ECG signal was manually annotated by an expert, on a minute-by-minute basis, for sleep apnea and hypoapnea events using additional signals

[1] http://www.physionet.org/physiobank/database/apnea-ecg/

Table 8.1: *Characteristics of patients in the Leuven dataset.*

Patient	Gender	BMI	AHI
1	M	40.8	72.6
2	F	26.3	3.2
3	M	21.7	0.6
4	M	30.5	42.7
5	M	27.8	45.6
6	M	29.4	26.5
7	M	31.3	29
8	M	30.9	29.6
9	M	26.6	36.2
10	M	36.6	44.2

such as respiration and oxygen saturation. Unfortunately, no differentiation between hypoapneas and apneas was indicated in the annotations. Eight of the 70 ECG recordings are accompanied by concomitant respiratory effort measured with respiratory belts in the abdomen and thorax. In total, 34 324 annotated minutes were extracted and used in this study.

The second dataset was recorded in the Sleep Laboratory of the University Hospital Leuven (UZ Leuven, in Belgium), and consists of 10 single-lead ECG (lead II) signals extracted from polysomnographic recordings of 10 different patients. The mean age of the patients was 48.4 ± 11.2, and they presented habitual snoring, witnessed apneas and hypersomnolence. The body mass index (BMI), gender, and apnea-hypopnea index for each patient are indicated in Table 8.1. Information about respiratory effort was acquired using two respiratory belts placed around the thorax and the abdomen. All signals were sampled at 200Hz and their duration ranges from 320 to 469 minutes (384 ± 47). In total 5956 minutes were extracted and annotated according to the *Chicago Criteria* [147], by a technician and a physician experienced on interpreting polysomnographic signals. These annotations correspond to the relative and absolute time of the event, duration, and type of respiratory event: obstructive apnea (OSA); central apnea (CEN); obstructive hypopnea (OSH); hypopnea (HPA); or mixed apnea (MIX). The events annotated as OSA and HPA are characterized by a reduced or complete absence of airflow respectively, during at least 10 seconds. As mentioned in Section 8.1, the detection of airflow reductions can be made from signals like the oxygen saturation and airflow. A central apnea (CEN) consists of a cessation of respiratory effort, and mixed apneas (MIX) are events with initial absence of respiratory effort that are accompanied by upper airway obstructions. A hypopnea was classified

as obstructive if there was presence of snoring during the event or increased inspiratory flattening of the nasal pressure flow signal compared to baseline breathing or associated thoracoabdominal paradox during the event but not during pre-event breathing.

8.2.2 ECG pre-processing

Each ECG signal is segmented into epochs of 60s. This is done based on the results reported in [55], where different epoch lengths were analyzed. It was concluded that segments of 60s allow to obtain the best performance when discriminating between respiratory events.

After segmenting the ECG, the power line interference at 50Hz is filtered out using a notch filter, and the mean of the signals is removed. Next, all ECG signals are up-sampled at 250Hz [129] by means of cubic spline interpolation. This is done in order to fall within the frequency range recommended in [129]: 250Hz-500Hz. These R-peaks are found using Algorithm 7 described in Section 4.2.

8.2.3 ECG contamination level

It is well known that different types of artefacts contaminate the ECG signals. Hence, when deriving signals like the tachogram it is necessary to determine whether an ECG segment is contaminated or not. By doing so, it is possible to estimate the reliability of the results obtained from that particular segment. To do this, the methodology proposed in Section 4.1.2 was used, where a weight ω is assigned to each ECG segment depending on how similar its autocorrelation function (ACF) is to the ACFs of the past segments. The similarity value for each segment is calculated using the cosine pair-wise similarity between ACFs, and its weight is then computed as the normalized sum of its pair-wise similarities. The closer ω is to one, the cleaner the segment is from artefacts. In [213], all segments with weights below a certain threshold (95th percentile) were removed from the dataset. Here, however, no threshold is applied. Instead, ω is used to provide a contamination level for each particular segment, and only segments with a certain value of ω are included in the training phase.

8.2.4 ECG-derived respiration

It has been shown by several authors [38, 54, 108, 109, 193] that respiratory events can be detected using an ECG derived respiration (EDR). The derivation

of this approximated respiratory signal is possible due to the mechanical interaction between the respiratory movements and the morphology of the ECG. This is to say that during each breathing cycle the electrical impedance of the thorax, and the relative position of the electrodes with respect to the heart change due to variations of the amount of air in the lungs. These variations can be detected from the changes in the amplitude of the different waves of the ECG.

Here, the three different methodologies described in Section 4.4 were used to compute the EDR signal: R_{amp}, R_{pca}, and R_{kpca}. The performances of the algorithm using each EDR will be assessed and compared against the one obtained using the real, reference respiratory signals. This is done in order to determine the added value of each methodology on the detection of sleep apnea from portable ECG-based monitoring systems.

8.2.5 Feature based on the QRS morphology

A new feature derived from the ECG is proposed and used here for the detection of sleep apnea. This feature is based on the computation of the EDR using PCA, which was presented in Section 4.4. The variation here is that the algorithm is applied to the matrix X^T rather than to X. The reason for this relies on the fact that when considering X^T, a “local" interpretation of changes in morphology around the R-peaks can be made. Furthermore, when an episode of apnea occurs, the changes in QRS morphology within one segment are more heterogeneous, an effect that can be observed from the eigenvalues of X^T. This is illustrated in Figure 8.1, where the eigenvalues obtained from two different ECG segments are indicated. The percentage of the variance explained by each component is also indicated, and it is defined as $(\lambda_i / \sum \lambda_i)$, where λ_i is the ith eigenvalue of the covariance matrix of X^T, $i = 1, \ldots, m$, and m the window size. Note that during a normal episode (i.e., non apneic) the variance explained by the first component describes more than 90% of the modulation of the QRS complexes. However, when apnea occurs, it is necessary to also consider the information contained in the second component, which explains more than 40% of the variance. Following these observations, the percentage explained by the second component is computed for each ECG segment of 60s, and it is used as a “morphology" feature during the classification phase.

The procedure discussed above is strongly connected to the singular value decomposition of the matrix $\mathbf{X}$, which is defined as $\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T$, where $\mathbf{U}$ are the eigenvectors of $\mathbf{X}\mathbf{X}^T$, $\mathbf{V}$ the eigenvectors of $\mathbf{X}^T\mathbf{X}$ and the non-zero elements of the matrix $\mathbf{S}$ correspond to the squared root of the eigenvalues

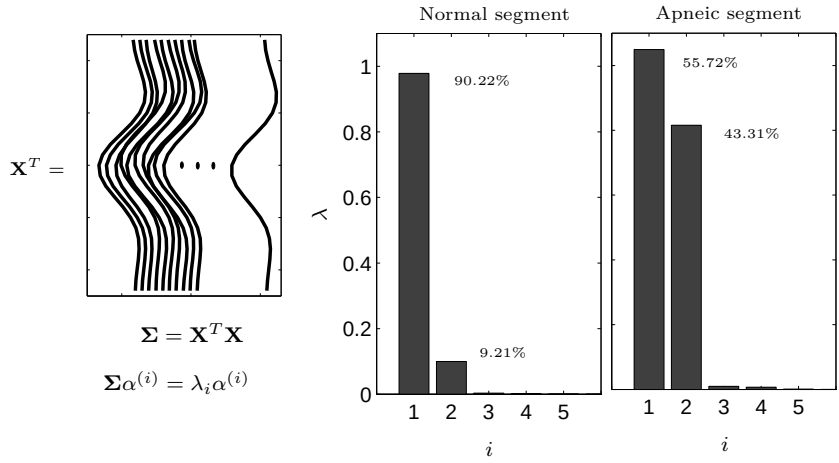

Figure 8.1: *Eigenvalues of the covariance matrix of* $\mathbf{X}^T$. *The percentage of variance explained by each component is indicated next to the first two eigenvalues for both cases: apnea and normal. Note that in order to describe most of the variance contained in the morphology of the QRS complexes for an apneic segment, it is necessary to also consider the second principal component.*

of $\mathbf{X}\mathbf{X}^T$ or $\mathbf{X}^T\mathbf{X}$. However, the PCA approach was adopted here in order to include a direct link with the EDR methodology.

8.2.6 Feature derived from subspace projections

It is well known that variations in heart rate are modulated by variations in respiration through the physiological mechanism called respiratory sinus arrhythmia (RSA) [28]. Therefore, the information contained in the heart rate signal could in principle be split into two different components, one related to respiratory changes, and one related to physiological mechanisms different from respiration (e.g., sympathetic activation). This decomposition is obtained using the methodology presented in section 5.3.

Given are two physiological signals $\mathbf{X}$ and $\mathbf{Y}$, where $\mathbf{Y}$ is the RR interval time series RR (resampled at 5Hz), and $\mathbf{X}$ corresponds to one of the following respiratory signals (resampled at 5Hz): R_{amp}, R_{pca}, R_{kpca}, R_a or R_t, where the last two correspond to the respiratory effort measured on the abdomen and on the thorax, respectively. All the information contained in $\mathbf{X}$ that is linearly related to $\mathbf{Y}$, can be computed by projecting $\mathbf{X}$ onto a subspace $\mathbb{V}$ defined by

variations in $\mathbf{X}$. In order to define a basis $\mathbf{V}$ for the subspace $\mathbb{V}$, Algorithms 8, 9, and 10 described in section 5.3, were used and they are briefly summarized below.

1. A 5 level wavelet decomposition of $\mathbf{X}$ using *db4* is used to define the basis $\mathbf{V}_1$, describing the low frequencies of respiration, namely frequencies below 0.07Hz. The detail coefficients of level 3, 4 and 5, and their delayed versions, are used to construct the basis $\mathbf{V}_2$, which describe respiratory variations between 0.07Hz and 0.6Hz. See Algorithm 8.

2. The reference signal $\mathbf{X}$ is first filtered using a band-pass filter with cutoff frequencies at 0.04Hz and 0.15Hz [129]. The filtered signal and its delayed version are used to construct the basis $\mathbf{V}_3$. In order to describe the high frequency (HF) band, $\mathbf{X}$ is filtered using the same filter, but now with cutoff frequencies at 0.15Hz and 0.4Hz. This new filtered signal and its delayed version are used to construct $\mathbf{V}_4$. See Algorithm 9.

3. The subspace related to respiration is spanned by $\mathbf{V}_5$, which is constructed using the original $\mathbf{X}$ and its delayed version. See Algorithm 10.

After constructing a basis $\mathbf{V}_i$, with $i = \{1, \ldots, 5\}$, the corresponding component $\mathbf{Y_{Xi}}$ can be computed. It is important to mention, that when computing $\mathbf{Y_{X5}}$, it is possible to describe all information contained in the heart rate that is *not* related to respiration. This will be denoted by $\mathbf{Y_{X6}}$, and can be computed as $\mathbf{Y_{X6}} = \mathbf{QY}$, with $\mathbf{P} = \mathbf{V}(\mathbf{V}^T\mathbf{V})^{-1}\mathbf{V}$, and $\mathbf{Q} = \mathbf{I} - \mathbf{P}_5$[2].

Once the component related to respiration $\mathbf{Y_{Xi}}$ is obtained, its relative power is computed as

$$F_i = \frac{\mathbf{Y}_{\mathbf{Xi}}^T \mathbf{Y_{Xi}}}{\mathbf{Y}^T\mathbf{Y}}, \tag{8.1}$$

and it corresponds to a feature that will be used as input to the classifier. In total, 6 features are computed for each ECG minute using orthogonal subspace projections, and they are summarized in Table 8.2, which also includes all the features used in this study. These other features correspond to the standard deviation of the RR interval time series (S_1), standard deviation of the respiratory signal R (S_2), and the serial correlation coefficients of the RR (r_k). For the last one, $RR = \{rr_i\}_{i=1}^{N}$, with N the total amount of RR intervals per segment. μ corresponds to the mean of the RR intervals per segment, and $k = 1, \ldots, 5$.

[2]For details on this computation see section 5.3

Table 8.2: *Set of features. Resp can be interpreted as any real or estimated respiratory signal*

Feature ID	Derived from			Frequency Range	Computation
	ECG	RR	Resp		
S_1		x		-	$std(RR)$
S_2			x	-	$std(R)$
r_k		x		-	$\frac{\sum_{i=1}^{m}(rr_i-\mu)(rr_{i+k}-\mu)}{\sum_{i=1}^{m}(rr_i-\mu)^2}$
PC	x			-	$\frac{\lambda_2}{\Sigma\lambda_i}$
F_1		x	x	0-0.07Hz	$F_i = \frac{\mathbf{Y_{Xi}^T Y_{Xi}}}{\mathbf{Y^T Y}}$, with $\mathbf{Y_{Xi}} = \mathbf{P}_i\mathbf{Y}$, $i = \{1,\ldots,5\}$, and $X_6 = \mathbf{QY}$
F_2		x	x	0.07-0.6Hz	
F_3		x	x	0.04-0.15Hz	
F_4		x	x	0.15-0.4Hz	
F_5		x	x	All	
F_6		x	x	All	

As mentioned before, the respiration modulates changes in the heart rate through the RSA, where the *RR* intervals are shortened during inspiration, and elongated during expiration. This mechanism allows for an improved gas exchange that leads to an optimal energy expenditure when respiration and heart rate are highly synchronized [28]. This synchronization is in some cases reduced during episodes of apnea [158]. However, a more important effect of apnea is related to an increased sympathetic modulation, which causes changes in heart rate that are not modulated by RSA [201]. The increased sympathetic activity has been already studied by means of the quantification of the sympathovagal balance, which can be computed as the ratio between the power of the heart rate in the low frequency band (0.04-0.15Hz) and the power in the high frequency band (0.15-0.4Hz). This ratio however, is defined in order to take into account *all* vagally and sympathetically mediated activities. Therefore, methodologies like the one described above, allow to determine whether or not respiration is the strongest modulator of the heart rate during a given minute epoch. This is to say, with this methodology, it is possible to extract the information of the heart rate that is modulated by mechanisms

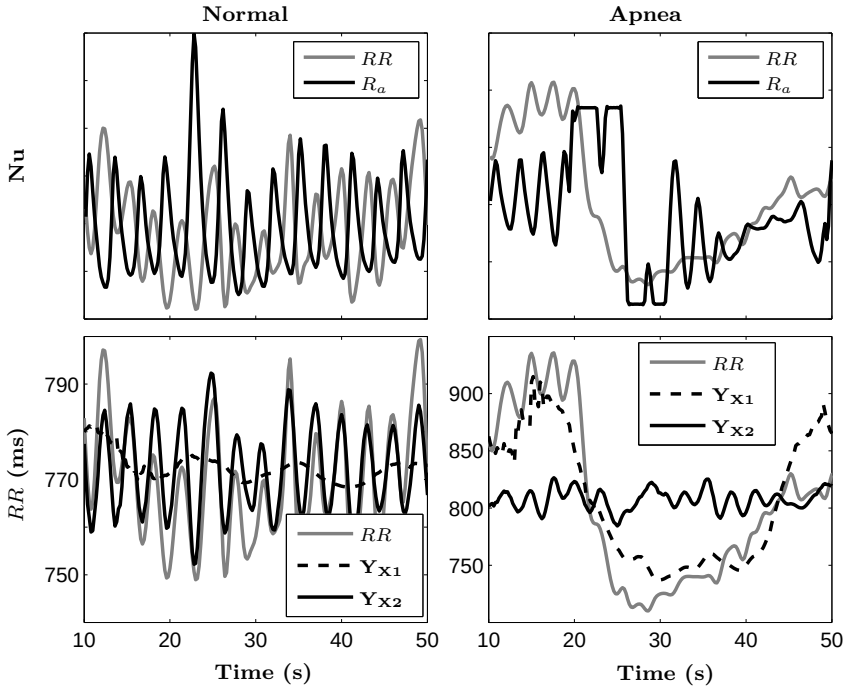

Figure 8.2: *(top) RR interval time series and respiratory signal measured in the abdomen for a normal (left) and an apneic (right) segment. Nu stands for normalized units. (bottom) Heart rate components derived using orthogonal subspace projections.* $\mathbf{Y_{X1}}$ *corresponds to the component related to respiratory dynamics in the range between 0-0.07Hz, and* $\mathbf{Y_{X2}}$ *describes variations in the range between 0.07-0.6Hz. Note that the components* X_1 *and* X_2 *get values from* $t = 10s$ *because the order of the model was 50, which corresponds to 10s for a sampling frequency of 5Hz.*

different from respiration. For instance, Figure 8.2 shows a normal minute of respiration and heart rate, and a minute containing obstructive sleep apnea. The two different heart rate components are also indicated in the figure. Note that the relative amplitude of the component related to respiration decreases during an episode of apnea.

8.2.7 Feature selection

In total, 28 features were extracted from each ECG minute, namely, the standard deviation of RR (S_1), the standard deviation of the three EDR signals (S_2), 5 serial correlation coefficients (r), the relative power of the second principal component (PC), and 18 features computed using orthogonal subspace projections (F). The reason to compute $18\times F$ is that three EDR algorithms are used to estimate the respiration. In order to determine the discriminative power of each respiratory signal and of each feature, two different statistical tests are used, namely, the F-score [48] and the Kruskal-Wallis test.

The F-score for a given feature z_i, with $i = \{1, \ldots, 28\}$, and the number of positive and negative instances denoted by n_+ and n_-, respectively, is computed as

$$\mathrm{F}_s(i) \equiv \frac{(\bar{z}_i^{(+)} - \bar{z}_i)^2 + (\bar{z}_i^{(-)} - \bar{z}_i)^2}{\mathrm{Z}_i^+ + \mathrm{Z}_k^-}, \tag{8.2}$$

with

$$\mathrm{Z}_i^+ = \frac{1}{n_+ - 1} \sum_{k=1}^{n_+} (z_{k,i}^{(+)} - \bar{z}_i^{(+)})^2$$

$$\mathrm{Z}_k^- = \frac{1}{n_- - 1} \sum_{k=1}^{n_-} (z_{k,i}^{(-)} - \bar{z}_i^{(-)})^2, \tag{8.3}$$

where $\bar{z}_i$, $\bar{z}_i^{(+)}$, and $\bar{z}_i^{(-)}$ correspond to the average of the values of the ith feature for the whole set, the positive and the negative sets, respectively. The larger the F-score, the larger the discriminative power of the feature.

For each derived feature, one p-value (using the Kruskal-Wallis test), and one F-score are computed. In addition, these values are also computed for the transformed features obtained after applying the logarithm. This logarithmic transformation is used to obtain features with distributions closer to normal.

8.2.8 Classification

Once a set of features was selected for classification using the approach described in the previous section, segments with $\omega > 0.9$ were selected as candidates for the training set, which was then extracted using the modified fixed-size method (FS2) proposed in section 2.3.2. The value of ω was selected using the values indicated in Figure 8.3, where it is clear that when a threshold

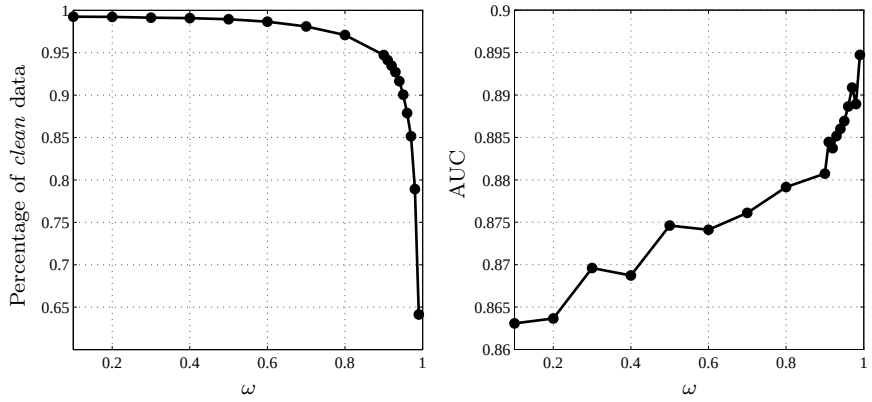

Figure 8.3: *(left) Percentage of the amount of data labeled as clean in the Physionet dataset with respect to the contamination level ω. (right) Area under the receiver operating characteristic curve (AUC) for different thresholds of ω. Note that when $\omega = 0.9$, about 95% of the data is considered as "clean", and the slope of the performance changes.*

of 0.9 is applied on the values of ω, about 95% of the data is kept. In addition, the figure also shows that the slope of the curve AUC vs. ω changes for $\omega > 0.9$. Therefore, by selecting the training set from segments with $\omega > 0.9$, a maximum performance is achieved by only removing about 5% of the data.

The FS2 method finds a subset of data points that maximizes the Renyi entropy. In other words, the subset contains parameters derived from "clean" ECG segments, and it represents the underlying distribution of the dataset. This subset is used to perform training and validation, and the remaining data points are later used to test the apnea classifiers.

Different classifiers are used in this study, and they are listed as follows:

1. Linear Discriminant Analysis (LDA)
2. Support Vector Machines (SVM) using linear, polynomial and RBF kernels
3. Least-Squares SVM (LS-SVM) using linear, polynomial and RBF kernels [187].

Their performances are evaluated using sensitivity, specificity, area under the receiver operating characteristic (ROC) curve, and accuracy.

8.3 Results

This section presents the comparison between the different EDR signals and the real respiratory efforts measured from the abdomen and thorax. A comparison between the different parameters extracted from either signal was performed for the two datasets separately, namely the Physionet, and the Leuven datasets. Furthermore, the features with the highest F-score and significant p-values were selected and will be presented below. In addition, different classifiers were evaluated, and a final algorithm to classify sleep apnea was identified. All experiments were carried out in MATLAB R2013a on an Intel(R) Core(TM) i7, 2.4GHz, 8GB RAM, running Windows 8.1.

8.3.1 Comparison between respiratory signals: real and EDR

In order to assess the usability of the EDR signals for the classification of sleep apnea, it is necessary to guarantee that the derived features do not differ significantly from those derived from the real respiratory signals. Hence, the features derived by means of orthogonal subspace projections on the heart rate and each respiratory signal (real or estimated) were compared. Figure 8.4 shows the different relative powers F_i computed using (8.1), for each respiratory signal, and for the two datasets used in this study. The feature ID corresponds to the notation used in Table 8.2. No significant differences were found between the respiratory signals. Moreover, it is clear from the figure that the three EDR signals led to the computation of very similar parameters. No discrimination between normal and apnea episodes was made at this point.

It can be observed from Figure 8.4 that the respiratory information derived using the R-peak amplitude, R_{amp}, is similar to the one obtained with the real respiratory signals. Figure 8.5 shows different examples of the components derived from the heart rate using R_{amp}, for a normal segment and different types of apneas. Note that in the segments containing apneas clear changes in the derived respiratory signal can be observed, and the amplitude of the components independent of respiration (dashed lines) become larger. This indicates that the simplest EDR algorithm already allows to observe variations in the respiratory effort, which are typical during apnea events.

8.3.2 Feature selection

After the observations described in the previous section it is clear that one can rely on features derived from estimated respiratory signals. This becomes

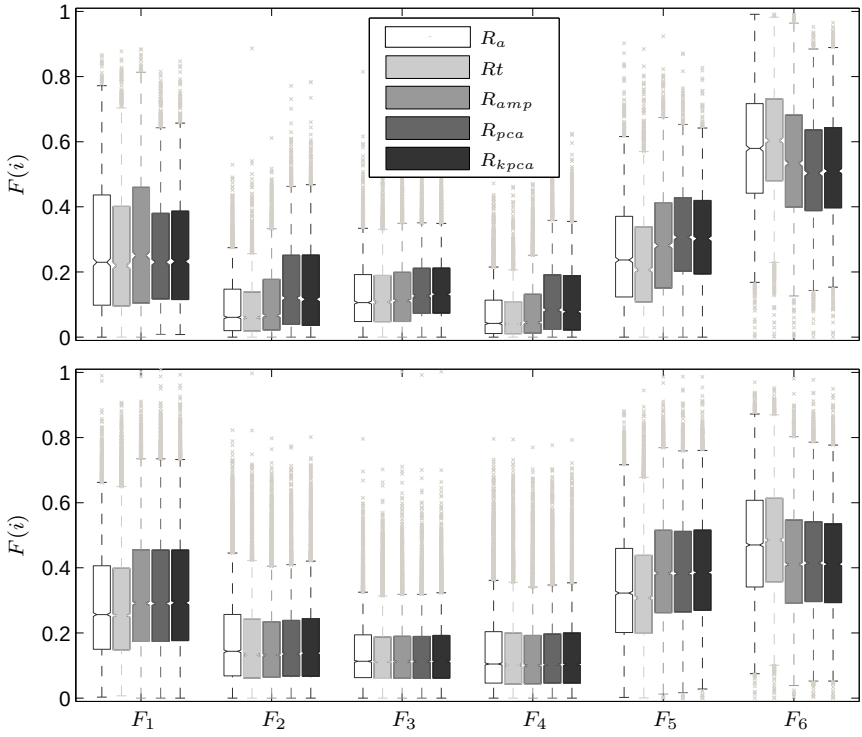

Figure 8.4: *Features extracted using orthogonal subspace projections in the Physionet (top) and Leuven (bottom) datasets. A comparison between the features obtained from the real respiratory signals and the three different EDR algorithms is indicated. No significant differences were found between the features. Note that the three EDR signals R_{amp}, R_{pca} and R_{kpca}, lead to very similar features.*

of high importance, especially in datasets such as Physionet, where not all recordings include real measurements of respiration. With this in mind, only the features derived from the EDR were taken into account in the selection of the final feature set. In addition, no significant differences were found between different EDR. For this reason, the features computed from R_{amp} were favored due to the simplicity of the computation of the EDR using only the amplitude of the R-peaks. It was found that the features with F-scores larger than 0.3, and p-values lower than 0.05 for the distinction between apneas and normal segments were: F_1, F_2, PC, and the logarithms of S_1, S_2 and r_5. Remember that F_1, F_2 and S_2 were computed using R_{amp}. Figure 8.6 shows the different features for segments with and without apneas. Note that for the Leuven dataset, a

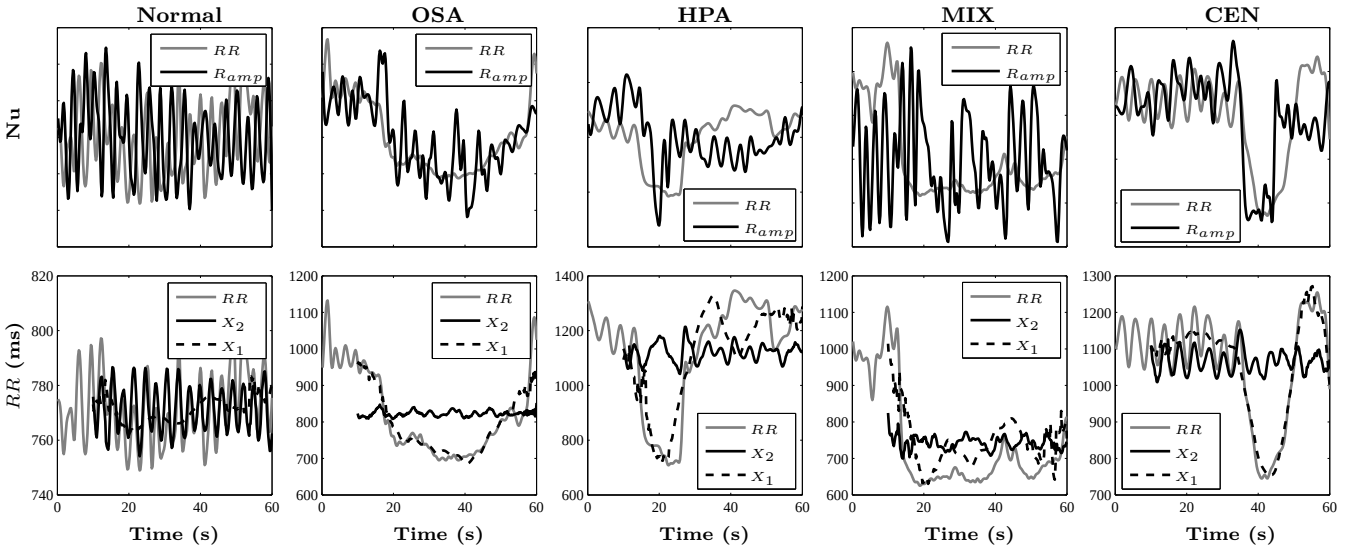

Figure 8.5: *Examples of the heart rate components derived using R_{amp}. Note that the dynamics of the EDR signal (R_{amp}) change in segments containing apnea events, which is an indication that the simplest EDR algorithm is sufficient to detect changes of the respiratory effort.*

discrimination between the different respiratory events is indicated. It is clear that the proposed features already provide a remarkable differentiation between respiratory events.

8.3.3 Apnea classification

In total, 34 324 ECG minutes were segmented on the Physionet dataset, and 5956 on the Leuven dataset. As mentioned in Section 8.2.3, the similarity between the ACF is used to determine a level of contamination of the ECG segment. All segments with $\omega > 0.9$ were included in a subset of "clean" ECGs. From these clean ECGs (32 477 in Physionet, 5205 in Leuven) training sets of 6000 segments from Physionet, and 3000 segments from the Leuven dataset, were selected by means of the fixed-size method. The remaining segments were used as test set and the comparison between classifiers was performed on this set.

Table 8.3 presents the performance for all tested classifiers, and for both datasets. It is clear from the table that the LS-SVM classifier using an RBF kernel outperforms the other classifiers in both datasets. Using the best classifier, the apnea-hypopnea index (AHI) is computed for each patient and each dataset. Results are shown in Fig. 8.7, where the real AHI is plotted against the estimated one. Note that the largest differences are obtained for larger AHI in both datasets. Some of these differences can be explained by the fact that apnea events can occur several times per minute, and since the data was segmented on a minute-by-minute basis, all those events are grouped into one single episode. If a threshold of 10min/h in apnea is applied to both datasets, it is possible to separate normal from apnea patients in the Physionet dataset, and patients with a higher than 10 AHI in the Leuven dataset, with an accuracy of 100%.

When looking at the classification results of the Leuven dataset, it is possible to manually identify the distribution of different types of respiratory events within the class labeled as *apnea*. This class contained 78.7% of 229 OSA, 37.7% of 496 OSH, 24.14% of 55 HPA, 93.48% of 48 MIX, 50% of 16 CEN, and 8.9% of 1454 normal segments, from the test set. As can be seen, if the OSH and HPA are excluded from the analysis, the classifier achieves accuracies of more than 93%. These results will be discussed in the next section.

The results described above were obtained after selecting the training set using the "clean" ECG segments, which correspond to the segments with $\omega > 0.9$. In total 2598 out of 40 280 had $\omega < 0.9$.. If the training set is selected from the whole dataset, the AUCs using LS-SVM with an RBF kernel, are reduced to 80.2% and 84.1% for Physionet and Leuven test datasets, respectively. This

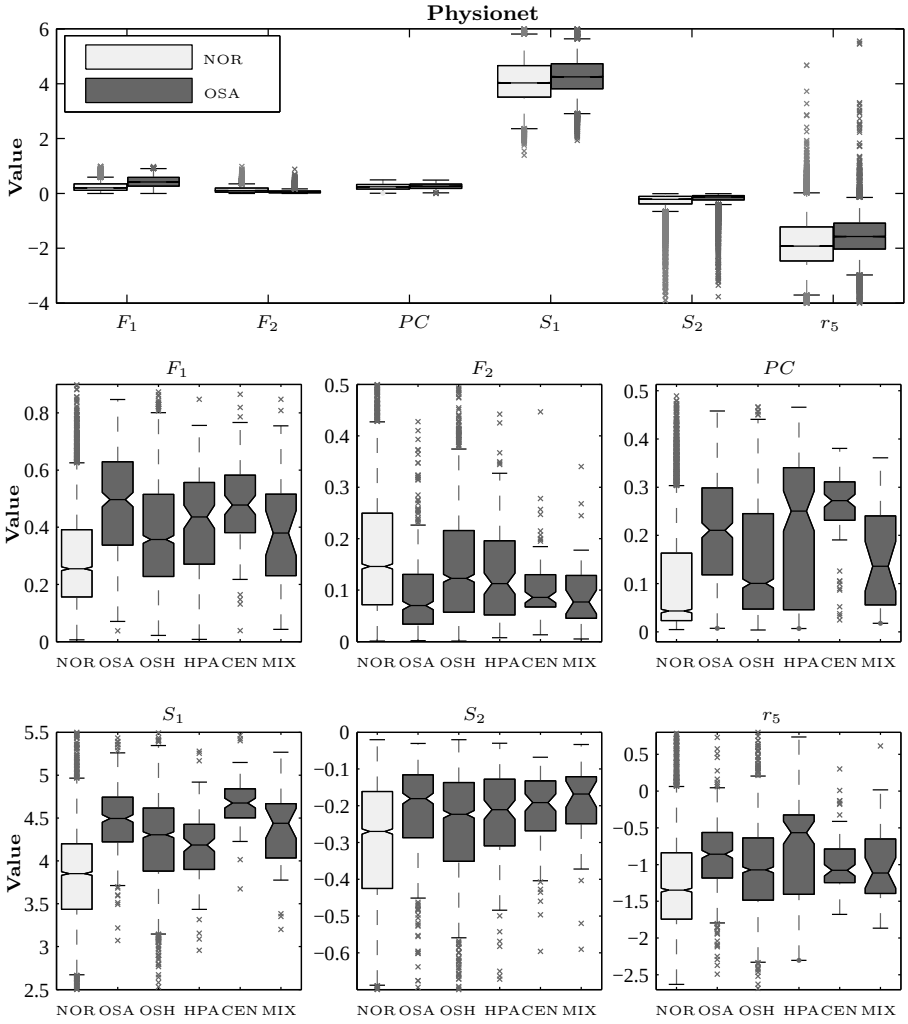

Figure 8.6: *Features selected as the most discriminative ones. The top panel indicates the features for the physionet dataset, and the other 6 panels show the features for the Leuven dataset and for all respiratory events. The F-scores of the features are larger than 0.3 and their p-values are lower than 0.05 for the distinction between normal segments and respiratory events. Note that the proposed features, namely F_1, F_2 and PC provide already a clear differentiation between normal and obstructive apneas. NOR stands for normal.*

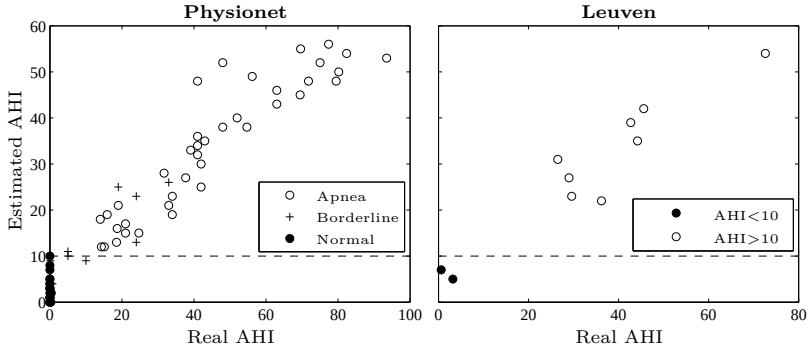

Figure 8.7: *Real and estimated apnea-hypopnea indices (AHI) for both datasets under investigation. Note that the maximum value for the estimated AHI is 60 due to the minute-by-minute analysis. A complete differentiation between the normal and apnea patients in the Physionet dataset can be made using a threshold of 10min/h (dashed line).*

shows that by including segments with a high confidence level in the training set, it is possible to reduce the bias in the classifier towards the presence of artifacts.

8.4 Discussion

This study had three main objectives, namely, the evaluation of the two novel features for the detection of sleep apnea, the comparison between three algorithms to derive respiratory information from the ECG, and the assessment of their added value in the classification of apnea. Bearing in mind the results presented in the previous section, it is clear that both features proposed here allow for a high discrimination between normal segments and respiratory events. Furthermore, the classification performance achieved with these two features is comparable with many algorithms proposed in the literature, where many more features are derived from the RR interval time series and the EDR. For instance, for purely automatic algorithms based on single-lead ECG, accuracies of about 85 to 90% were reported in [54, 156] on the Physionet dataset. We believe that the difference in accuracy between this work and the literature is due to the size of the training set. In [54, 156], 35 recordings were used to train the classifiers, while in this work "only" 6000 ECG segments were selected as training set. The reason for this is that the algorithms based on SVM formulations require the computation of a kernel matrix, of which the size is limited by the hardware specifications. In this work, the size was limited

Table 8.3: *Performances of the tested classifier on the Physionet and on the Leuven datasets. All respiratory features were computed using R_{amp}. The final feature set was: $\{F_1, F_2, PC, \log S_1, \log S_2, \log r_5\}$*

Classifier		**Physionet**				**Leuven**			
		Sens.	Spec.	Acc	AUC	Sens.	Spec.	Acc	AUC
LDA		71.74%	71.20%	71.43%	78.46%	81.22%	79.71%	79.86%	87.69%
SVM	Linear	74.63%	68.51%	71.16%	78.5%	**86.46%**	74.6%	75.83%	87.74%
	Polynomial	78.36%	68.22%	72.6%	80.4%	88.21%	76.57%	77.78%	89.90%
	RBF	78.2%	70.55%	73.86%	81.85%	82.97%	81.78%	81.9%	88.41%
LS-SVM	Linear	71.48%	76.26%	71.78%	78.46%	80.35%	80.01%	80.05%	87.7%
	Polynomial	73.4%	73.43%	73.43%	81.18%	81.66%	81.7%	81.68%	89.17%
	RBF	**84.71%**	**84.69%**	**84.74%**	**88.07%**	78.81%	**84.56%**	**83.95%**	**89.97%**

to 6000×6000. Nevertheless, even with a smaller training set, an accuracy of 84.74% was achieved for the Physionet dataset.

Concerning the feature derived from the morphology of the ECG, namely PC, it was observed that indeed, the relative power of the second component is higher due to variations in the QRS complexes caused by an increased sympathetic modulation.

Turning now to the feature derived using orthogonal subspace projections it is clear from the results that the information shared between heart rate and respiration, in the traditional LF and HF bands is significantly reduced during episodes of apnea. However, when looking at frequencies lower than 0.07Hz it was surprising to see that changes in respiration could also be observed in the heart rate. This could be the result of baseline variations in the respiration due to a reduced or complete absence of respiration, which cause a rise in the heart rate at very low frequencies. It is important to remember that the analysis was performed on a minute-by-minute basis, which can be understood as a limitation in the frequencies that can be described. This is to say that only frequencies between 0.03Hz and 0.07Hz can be analyzed by this feature.

Another point that requires attention is the fact that a comparison between three different algorithms to derive the EDR signal was performed. In addition, the features extracted from these signals were validated against those computed using the real respiration. It was remarkable to see how those features did not differ from each other. With this in mind it is clear that the usability of EDR signals on applications like sleep apnea detection is valid and significant. However, since no differences were observed between the three EDR algorithms, it is possible to rely on the simplest approach in order to reduce the complexity of the whole methodology. This last observation has clear advantages when developing portable devices, where less sensors are needed. Hence, if an accurate respiration can be estimated from the ECG, there is no need for respiratory belts.

It is not a surprise that the most difficult respiratory events to classify from the ECG signal are the hypopneas. As mentioned in section 8.1, these events are associated with a reduced airflow and a reduced respiratory effort, and can be accompanied by a drop in oxygen saturation, or by a reduction of more than 50% of the flow during at lest 10s, or by a reduced airflow with arousal or desaturation. Additionally, the difference between normal breathing and (central) hypopnea is not always clear [147]. This is because these respiratory events are not always accompanied by arousal, desaturation, snoring or thoraco-abdominal paradox. Therefore, they may appear very similar to normal segments. For this reason, in order to get a more reliable classification of this type of events different physiological signals such as airflow, oxygen saturation

and CO_2 can be considered to complement the proposed methodology. The classifiers used in this work only discriminated between normal and respiratory events, and all apneas and hypopneas were considered to belong to one class labeled as *apnea*. The differentiation made before was done manually in order to determine the distribution of respiratory events within the group classified as *apnea*.

8.5 Conclusions

This chapter presented a methodology for the automatic detection of sleep apnea from single-lead ECG signals. Two novel features were proposed for the characterization of the ECG signal. One describes the changes in the morphology of the QRS complexes caused by an increased sympathetic modulation during apnea events, and one quantifies the information shared between respiration and heart rate at different frequency ranges. These two features together with two well known parameters used in heart rate variability, namely the standard deviation and the serial correlation coefficient allowed to discriminate respiratory events from normal activity. The discrimination was performed using a LS-SVM classifier with an RBF kernel and it proved to perform better than other classifiers in the separation of obstructive apneas from normal segments. The accuracies achieved in this study are comparable to the best results reported thus far for fully automated algorithms, namely in the range between 85% and 90%.

One clear disadvantage of the proposed methodology is the fact that only a limited amount of segments can be used in the training phase. This is due to the need to compute a kernel matrix in the LS-SVM classifier. The size of this matrix is limited by the computational resources, and therefore many relevant normal or apnea examples can be missing from the training set. As a result, the performance of the classifier can be affected.

Another limitation of this study, and in general, of studies involving only the ECG signal, is that no clear ECG changes can be observed during most of hypopneas, not even for a human expert. In addition, the definition of these events has been modified a number of times over the last decades. Hence, in order to achieve a high accuracy in the detection of "all" respiratory events, the methodology proposed in this chapter needs to be complemented by the analysis of other modalities (e.g. SpO_2). Furthermore, to apply this methodology on different datasets, the definitions of hypopneas used during the annotations of the events for each specific dataset, need to be taken into account when evaluating the classifier's performance.

Apart from providing a label to each ECG minute, the methodology allows to estimate the contamination level of the recording by means of the autocorrelation function. This level (i.e. weight) was used in the training phase of the classifier, but it can also be used to determine the level of confidence on the results obtained for each ECG segment. In other words, manual inspection of the highly contaminated segments can be recommended based on these weights.

A different contribution of this work, was the comparison of three algorithms to derive respiratory information from the ECG. It was shown that the different parameters computed from the respiratory signals, real or estimated, were comparable for the detection of sleep apnea. This proves that EDR signals can achieve similar performances than real signals measured with respiratory belts. Therefore, the use of only ECG sensors can guarantee up to 85% accuracies in the discrimination of both apneas and hypopneas together. In addition, these performances can be achieved using only the amplitude of the R-peaks. Hence, the added value of more elaborated EDR algorithms was negligible in this particular application. This can be seen as an advantage in home monitoring systems that require algorithms with low computational cost.

Chapter 9

Conclusions and Future directions

9.1 Conclusions

This research presented different algorithms to extract diagnostic information from the electrocardiographic (ECG) signal. Novel features and algorithms were proposed for the analysis of the cardiac and respiratory activities derived from the ECG signal. Furthermore, their applicability and interpretability was investigated in two main diagnostic problems: detection of epilepsy and detection of sleep apnea.

Figure 9.1 illustrates the general ECG-based procedure proposed in this research. The grey boxes indicate the different contributions made during this work, and the output corresponds to either the detection of epileptic seizures, or the detection of sleep apnea events. It is worth noting that for each particular application, a set of features and algorithms were selected and validated.

This book was divided in two parts, algorithms and applications, and conclusions were drawn for each part separately, and they will be described next.

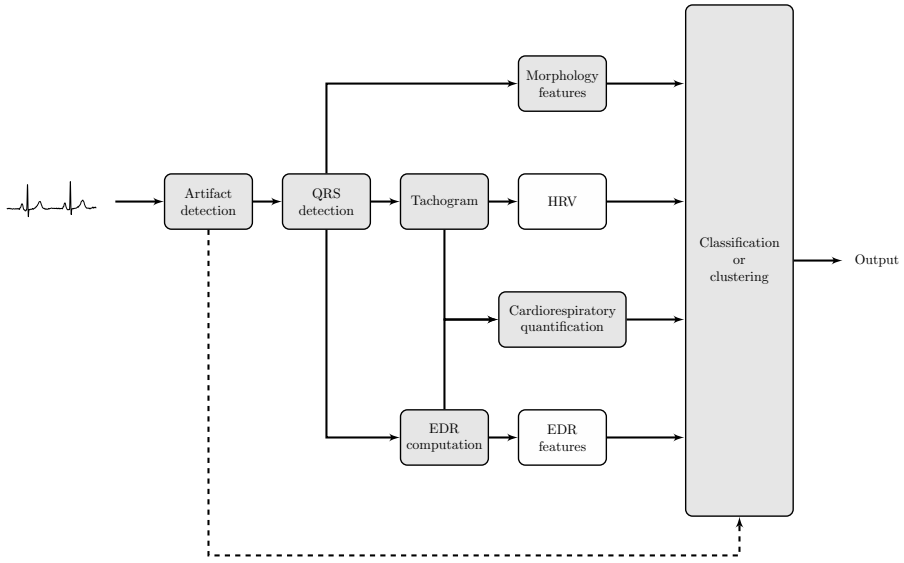

Figure 9.1: *Diagram of an ECG monitoring system that allows to extract cardiac and respiratory information. The gray boxes indicate the different tasks, in which many of the contributions of this research were applied.*

9.1.1 Algorithms

The contributions of this work can be divided in two main fields, machine learning, and ECG processing. Furthermore, all methodologies proposed in this research can be easily implemented and used in several ECG applications. For this reason, a set of graphical user interfaces and Matlab functions were developed and their descriptions can be found in the appendices. These tools are available for future research on ECG.

Machine learning

In the field of machine learning, one main contribution can be identified, and it is the novel model selection approach for kernel principal component analysis. The new approach based on distance distributions was proven to outperform many state-of-the-art approaches for the selection of the kernel parameter and the number of components in denoising applications. This was particularly clear in lower dimensional spaces, where other methods seem to fail. The approach was evaluated in different toy examples and real life applications

like the ECG-derived respiratory signal, and denoising of digits. The last application is typically used for the evaluation of denoising algorithms, hence, its use in this research was a must. Furthermore, it was shown that the impact of this new approach goes beyond biomedical applications, since it solves a problem that many in the field of machine learning attempted to solve in the past.

A disadvantage of this approach is that it requires the implementation of a Monte Carlo simulation, which can be seen as a bottleneck when online applications are considered. Future studies may consider to avoid this drawback.

A different contribution in this field corresponds to the modification of the fixed-size LS-SVM algorithm for the selection of training sets that best describe the underlying distributions of the input data. This modification finds a better subset of training points than the conventional fixed-size algorithm, in less iterations.

ECG processing

General algorithms for the pre-processing and analysis of the ECG signal were proposed in this research. First, an algorithm for the detection of ECG artifacts based on the autocorrelation function was proposed. This algorithm proved to detect the location of contaminated segments in long-term ECG signals with a high accuracy. In addition, it allowed to compute a *contamination level* for each segment, which can later be used to highlight the epochs that need visual inspection. By doing so, it can be possible to reduce the amount of time that doctors need to annotate the data. In online monitoring systems, however, this approach can be used to discard segments that exceed a certain contamination threshold. In other words, this level can be related to the reliability of the results obtained for each particular ECG segment. An important advantage of this algorithm is the simplicity of its computation, and for this reason, it can be easily implemented in ambulatory systems. This will need to be confirmed in future investigations that involve hardware and software implementations. One disadvantage of this proposed algorithm is that it requires a training period with relatively high signal-to-noise ratio, in order to set the reference for a "clean" ECG signal. This can be solved by performing a supervised training phase before the actual monitoring period.

Apart from the artifact detection algorithm, this research proposed a methodology for the detection of the QRS complexes. It is a modification of the well-known Pan-Tompkins algorithm, in which no derivative, square, or integrator filter is used. Instead, the envelopes of the ECG are used to remove many

variations unrelated to the QRS complexes. This algorithm was tested on a public dataset and its performance was comparable with those obtained with algorithms in the literature that use filters, wavelets, or other elaborated methodologies.

After detecting the location of the R-peaks, an evaluation of three different algorithms for the extraction of respiratory information from the ECG was performed. This information corresponds to the ECG-derived respiratory signal computed using the R-peak amplitude, principal component analysis (PCA), and kernel PCA. It was found that for nearly stationary ECG segments, the KPCA algorithm outperforms the other two in reconstructing the original respiratory signal. However, when transients and non stationarities were included, the algorithm based on the R-peak amplitude proved to be less sensitive and to produce more stable estimations of the respiratory signal. For these reasons, the R-peak amplitude was the preferred methodology in this research. Nevertheless, results obtained with the other two methods were also indicated. This is very good news, since the R-peak algorithm is the easiest to compute and it does not require demanding mathematical computations like the eigenvalue decomposition. Hence, its implementation in ambulatory systems can be straightforward.

Once the respiratory signal was estimated from the ECG, and the heart rate signal (i.e. tachogram) was computed, the cardiorespiratory interactions were quantified using three different methodologies. These were phase rectified signal averaging, information dynamics, and subspace projections. It was shown that these three methodologies offer complementary interpretations of the interactions between the heart rate and respiration, and they are significantly affected by changes in the autonomic control. Their applicability to the detection of epilepsy and sleep apnea was investigated, and the main conclusions will be described below.

9.1.2 Applications

The algorithms developed and evaluated in the first part of this research were used to analyze cardiorespiratory interactions in epilepsy and sleep apnea. In epilepsy they were used to uncover many unknown effects that epilepsy has on the autonomic nervous system, and to develop algorithms for the detection of epileptic seizures.

Epilepsy

The focus of this research was on childhood epilepsy, and on the effect that it has on the cardiac and respiratory control mechanisms of the autonomic nervous system. It was found that children suffering from West syndrome and absence epilepsy have an attenuated autonomic tone and a reduced cardiorespiratory coupling, which can compromise their defense mechanism against asphyxia and hypoxia. A different type of epilepsy was also investigated, namely temporal lobe epilepsy, where it was found that autonomic dysfunctions are only experienced during seizure activity. Nevertheless, it was clearly demonstrated that children suffering from epilepsy tend to have a more fixed heart rate, which can be an indication of a deeper underlying mechanism in epilepsy. This could be responsible for autonomic dysfunctions leading to SUDEP.

All these autonomic effects of epilepsy displayed a transient behavior towards the seizure onset. In other words, the decoupling between respiration and heart rate does not occur drastically at the seizure onset. This effect was later used to develop seizure detection algorithms that allowed to detected partial epileptic seizures with a PPV larger than 80%, in a window of 30s around the EEG onset of the seizures. Furthermore, a PPV of 83% was achieved for generalized seizures, which until today was not achieved by any other algorithm based solely on ECG analysis. However, the signals recorded during generalized seizures that involved movements were not included in this evaluation. This was done because the derivation of cardiac and respiratory information from such signals was not reliable, since movement artifacts affected the ECG signal beyond repair. Examples of these seizures are some tonic-clonic seizures that deeply involve motor activity, hence artifacts in the ECG make it impossible to correctly detect the position of the R-peaks. Therefore, other modalities (e.g. EMG, ACM) will need to be included to be able to detect such seizures. Other seizures that were not included in this study were the myoclonic seizures, which are characterized by very short jerks, and are in general terms not considered life-threatening events.

The cardiorespiratory analysis and the seizure detection systems used at this point, included most of the algorithms presented in the first part of this book and a novel feature proposed in this research, that can be used to detect changes in the morphology of the QRS complexes. These changes might be caused by variations of the heart rate, or other possible effects of epilepsy on the morphology of the ECG.

The interactions between respiration and heart rate were studied using PRSA, information dynamics, and subspace projections. All three methodologies displayed complementary information about the reaction of the ANS to the

occurrence of epileptic seizures. For instance, the speed of the reaction of the parasympathetic branch of the ANS could be observed, and it was found to be slower in epileptic children when compared with control subjects.

Despite the promising results obtained with the proposed methodologies, further investigations on the non-linear interactions between heart rate and respiration need to be performed. This is crucial since these interactions are known to be highly non-linear. However, the linear approximations obtained in this study, already provided new insights into the autonomic effects, and hopefully they will be taken into account for the design of new treatment strategies in epilepsy.

A clear advantage of the methodologies used in this research is that they are easy to implement and can be adapted to ambulatory systems or online detection platforms. This is important, since a lot of efforts focused on the development of closed-loop systems that allow to abort or predict the occurrence of an epileptic seizure. Nevertheless, before these algorithms can be directly implemented in such systems, it is necessary to consider a training phase for each type of seizure, especially those involving motor activity. In fact, a unified framework, where other modalities such as accelerometers and EMG are included, must be developed in order to detect with high accuracy that type of epileptic seizures.

Not only online detection systems can be considered. These algorithms can also be useful for offline detection of seizures, which is normally used for diagnosis and followup of therapies. In this context, the seizure detection system, together with the algorithm for noise level estimation of the ECG can be used in parallel to highlight the possible "medical" events that need visual inspection. This will allow to indicate not only dubious events but also events that could probably be missed by a human expert. As a result, the time spent on manually annotating and verifying the signals can be reduced.

Sleep apnea

A methodology for the automatic detection of sleep apnea from single-lead ECG signals was proposed. Two novel features were used, one that describes the changes in the morphology of the QRS complexes, and one that quantifies the information shared between respiration and heart rate. These two features showed to complement the information obtained from HRV analysis, and they allowed to classify apnea and hypopnea events with an accuracy larger than 85%. These performances are comparable with other methodologies in the literature that require the computation of more elaborated features.

In this research, a comparison between three algorithms to derive respiratory information from the ECG was performed. It was observed that the estimated respiratory signals achieved comparable performances than the real signals measured using respiratory belts. This is very important, since this proves that the EDR signal can be used to classify sleep apnea events. Furthermore, when the performances of the three EDR algorithms were compared, no significant differences were found between them. Therefore, it was shown that with the simplest of all them, namely the R-peak amplitude, accuracies of about 84% were achieved for the detection of sleep apnea. Hence, the added value of more elaborated algorithms is negligible, which can be seen as an advantage when developing home monitoring systems that typically require algorithms with low computational cost.

The algorithm for estimation of the contamination level in the ECG was used in this application to improve the performance of the classifiers. This algorithm was used to select a clean training set, and to identify the segments where manual inspection is recommended. As mentioned in the previous application, this can significantly reduce the time spent by doctors in annotating and verifying the signals (i.e. polysomnographic signals).

9.2 Future directions

Different extensions and future directions of this research can be identified. These will be described below, where a differentiation is made between further developments and improvements of algorithms, and further considerations for applications in epilepsy and sleep apnea.

9.2.1 Algorithms

- In the field of machine learning, the model selection criterion based on distance distribution can be adapted to other unsupervised methods like clustering, and to supervised techniques like regression. In addition, tests on signals with color noise must be performed in order to assess the relationship between this model selection and the type of noise present in the data.
- Concerning the algorithm for estimating the contamination level in the ECG, a way to automatically detect the optimal threshold to be used for the detection of artifacts must be investigated.

- The methodologies used to quantify the interactions between respiration and heart rate assume either stationarity or are only able to describe the linear interactions between the systems, or both. Because of this, they may not give a complete and accurate picture of the processes involved. For this reason, the field of Network Physiology has been recently introduced [20, 95], and it aims to study different systems as an interconnected network, where all (i.e. not only linear) interactions between these physiological signals are taken into account. In this way, a proper assessment of the effect of epilepsy on autonomic controls can be achieved. These important and fundamental issues must be addressed in order to gain a deeper insight in the (non-linear) processes involved.

9.2.2 Applications

- A dataset containing epileptic seizures and apnea events must be used in order to fine tune the features and parameters needed to detect either events. In addition, it is necessary to realize that two different populations were used in this study, namely children suffering from epilepsy, and adults with sleep apnea. Therefore, in order to develop a full ECG-based monitoring system that can work for both applications at once it is necessary to investigate both disorders on the same population. This is important because it is known that children suffering from epilepsy often experience seizures accompanied by sleep apnea episodes [149], and SUDEP occurs more often during sleep [66].

- In this research, unsupervised techniques were used to detect epileptic seizures. However, semi-supervised methods can be used and can probably lead to a better discrimination between seizures and normal activity. This is the case, because semi-supervised techniques allow to give some labels to the learning problem. In other words, the parts of the signals during seizures can be labeled, while other parts cannot. This also applies for sleep apnea classification. The reason for this is that both applications rely on the accurate annotation of the medical expert. However, it is known that human errors are included in these annotations, and as a result some events can be missed. Consequently, a complete supervised scheme learns from erroneous data, which does not lead to accurate discriminatory functions.

- Efforts must be made on the integration of other modalities in the detection of epileptic seizures. This was clear from the results of generalized seizures, where signals like the EMG and accelerometers can contribute to the improvement of the performance of detection systems.

- The field of apnea detection can benefit from the analysis of signals like the photoplethysmogram (PPG) and accelerometers. These can be included in the development of more comfortable systems to monitor sleep. Nevertheless, the validation of such modalities will need to be done using polysomnography (PSG), paying special attention to the synchronization of those new modalities with the gold standard, PSG.

Appendix A

Telehealth in cardiology

Telehealth has become a very important tool that allows the monitoring of heart failure patients in a home environment. However, little is known about the effect that such monitoring systems have on patients' compliance, evolution and self-care behavior. Especially, the effect that the selected user interface has on these factors is unknown. This study aims to investigate this, and to determine some practicalities that must be considered when designing and implementing a telehealth programme for heart failure. To achieve this, daily measurements of blood pressure, pulse, SpO_2 and weight were collected from 534 patients suffering from heart failure. In addition, they were asked to fill in the European heart failure self-care behavior scale questionnaire and the EQ-5D quality of life questionnaire, before and after the monitoring period. Two telehealth systems were used, namely, the Motiva platform provided by Philips and the standalone unit provided by Docobo, the Doc@Home system. Significant differences were found between both systems concerning the compliance and adherence of patients. Moreover, a general, positive effect of telehealth was identified thanks to increased self-awareness of patients when managing their condition. These findings are supported by behavioral changes and a better understanding of heart failure from the patients' perspective.

A.1 Introduction

Heart failure (HF) is recognized as one of the most common diseases in Western Europe, with more than 10% of the people older than 65 years suffering from it [134]. HF represents the most common cause of readmission in industrialized

countries, and more than 30% of patients hospitalized due to HF will die within 1 year [104]. It is well known that HF is often accompanied by several comorbidities such as coronary artery disease, hypertension, atrial fibrillations, and valvular heart disease [133]. Therefore, it is not surprising that this disease strongly impacts social and economic systems in a large scale. For instance, it has been reported [134] that the costs associated with heart failure represented more than 2% of the total cost of health care. Therefore, the use of alternative approaches such as telehealth tools that help improving the HF outcome, and reduce the costs, is imminent. Even though some telehealth programs have been implemented for the monitoring of heart failure patients [35] [222] [37] [58], their effect on the entire health care system is not yet understood. For example, it has been reported that telehealth seems to be more effective for more critical patients and patients with higher risk of hospitalization [145]. For this reason, it is essential to identify effective protocols that allow to better manage heart failure [39]. In this respect, the Recap (Regional Care Portals) project aims to improve and promote the implementation of telehealth systems in heart failure (Brainport Development, 2014). This project brings 13 partners from the North West Europe region together, to work towards the analysis of new legal, financial, organizational and technological solutions in telehealth. These solutions aim to improve the communication between patients and care givers, and to increase the awareness of patients with respect to the importance of self-managing their condition. With this in mind, it is expected that new telehealth solutions could have a positive effect on the diagnosis and treatment of HF.

This appendix deals with technological factors that need to be considered during the design and implementation of telehealth programs in heart failure. For this reason, three partners worked together on the implementation of the telehealth program. One of the partners involved in this task is the North East London NHS Foundation Trust (NELFT) in the UK, where HF patients were recruited and followed up. Health Enterprise East also in the UK, is the second partner involved in this work, which took care of the business planning and commercial research on telehealth systems monitoring HF patients. The third partner is the KU Leuven in Belgium, where the data analysis took place. In addition, the definition of technological practicalities need to be considered during the implementation of telehealth programs was performed.

In this study, two different telehealth systems were implemented, namely the Motiva platform offered by Philips (The Netherlands), and the standalone unit provided by Docobo (UK). Docobo model numbers are as follows: A&D UA-767PBT-C40 BT BP Monitor (3PM-A02 SN), A&D UC-321PBT-C40 BT Scales (3PM-A03 SN), HealthHUB 2 Roaming GPRS BT Kit (DCU-002-RGPRS-BT), and Nonin 9560 Bluetooth Oximeter (3PM-N06 SN).

The main differences between these two technologies are the user interface, mobility, and the educational information. Therefore, it is possible to assess differences between these systems, and determine whether education and mobility have an impact on the evolution of HF patients. Another important point investigated in this work is the effect of telehealth on the self-care behavior of patients. With this in mind, the general effect of education and self-management will be clarified and hopefully later used in the implementation of new telehealth programmes.

A.2 Methods

Ethics

The study was approved by Health research authority NRES east of England with ref no 12/EE/0137. Each patient signed a written consent form before starting the telehealth programme.

Patient population

The dataset used in this study consists of daily measurements of blood pressure, pulse, blood oxygen saturation (SpO_2), and weight, collected from HF patients remitted to the North East London NHS Foundation Trust (NELFT) in the UK, from November 2012 to November 2014. All eligible patients were older than 18 years and they had a confirmed diagnosis of left ventricular systolic dysfunction (LVSD) supported by a positive ECG. In addition, at least one of the following conditions was satisfied: sudden increase in weight of more than 1.5kg in 24 hours; systolic blood pressure lower than 90 mm Hg; sudden increase in shortness of breath; episodes of palpitation/tachycardia without collapse; change of medication within 48 hours of discharge from acute; or high Hospital Anxiety Depression Self-Assessment (HADS) score [223]. On the other hand, if any of the following conditions applied, the patient was excluded from the study: diagnosis without ECG; clinically stable NYHA I to II classification (New York Heart Association); Nil oedema; normotension (120/80 mmHg); end of life with NYHA larger than II; or insufficient cognitive understanding to complete questionnaires or to use the telehealth equipment.

Study design

Once a patient was enrolled into the program, the Hospital Anxiety and Depression Scale (HADS) was administered. Next, a demonstration of the use of the telehealth program was provided by the HF nurses, together with the explanation of the EU Recap project and the benefits of the whole program. In total 534 HF patients were eligible, with mean age 69.1yr±12.6yr. No record was kept of how many patients had declined to participate in the first place. Figure A.1 presents the flow of participants in the study. All patients were randomized into two groups, one receiving the standalone unit offered by Docobo (135 patients), and one receiving the Motiva telehealth platform offered by Philips (399 patients). The groups were unevenly split because more Philips Motiva equipment than Docobo systems were available for this study, hence more of the latter to give to patients. During the first days of monitoring, the patients were trained at home, to measure the 5 physiological parameters mentioned above, twice a day, on a daily basis. After 42 days of constant monitoring, the telehealth equipment was returned to NELFT and it was allocated to a different patient. A monitoring of 42 days was selected as the appropriate study period because research has shown heart failure readmissions are known to occur within 6 weeks at minimum to 6 months at maximum [1].

Apart from the physiological measurements, patients were asked to fill out the European heart failure self-care behaviour scale (EHFScBS) [97] and the EQ-5D [90] quality of life questionnaires, at the beginning and end of the monitoring period. This was done in order to assess changes in self-care behaviour and quality of life followed by the telehealth program. Additionally, the HADS was also administered at the end of each trial. Of the 534 patients, 411 (mean age 68.5yr±12.2yr) completed the 42 day program, 123 were females and 288 were males, and 105 patients received Docobo units and 306 patients received the Motiva system. All data transferred to KU Leuven for analysis was completely anonymized.

Interventions like changes in medication, titrations, and withdrawals were annotated in the data. However, no information about type of medication or specific doses was available. It is important to mention that all of these interventions were the same for both groups, and they were made mainly by the nurses with input from cardiologists and GPs where appropriate. Additionally, both groups were monitored by the same healthcare providers.

Table A.1: *Main characteristics of the telehealth systems*

Docobo	**Motiva**
Stand-alone unit	Interactive telehealth platform
No video streaming	Personal healthcare channel
No educational information	Educational information to control and understand the condition
Blood pressure (systolic/diastolic)	
Pulse rate	
Oxygen saturation (SpO_2)	
Body weight	

The telehealth systems

As indicated previously, two different monitoring systems were deployed, namely, the Docobo unit and the Motiva platform. Table A.1 indicates the main characteristics of both systems. The most important difference between these two systems relies on the the functions that they offer to the patients. For instance, the Motiva platform offers more educational material, a more familiar interface, since it is a personal healthcare television channel. In this channel, patients can read reminders sent from the health providers, and access videos where the importance of a healthy diet, physical activity, and medication is highlighted. Even though this information is crucial to encourage patients to stay healthy, no record is kept of when and how often the patients access and review this material.

On the other hand, the Docobo system used for this study is a stand-alone unit that does not provide educational information, but offers more mobility to the patients. The main goal of this study is to determine differences on the physiological parameters, and answers to questionnaires, for patients with different systems. In the same way, indirect effects of the user interface can be assessed.

Data analysis

In order to quantify the evolution of all HF patients that completed the 42 days monitoring period, the slope (m) for each parameter was computed using linear regression. Each patient was then characterized by 5 slopes that describe

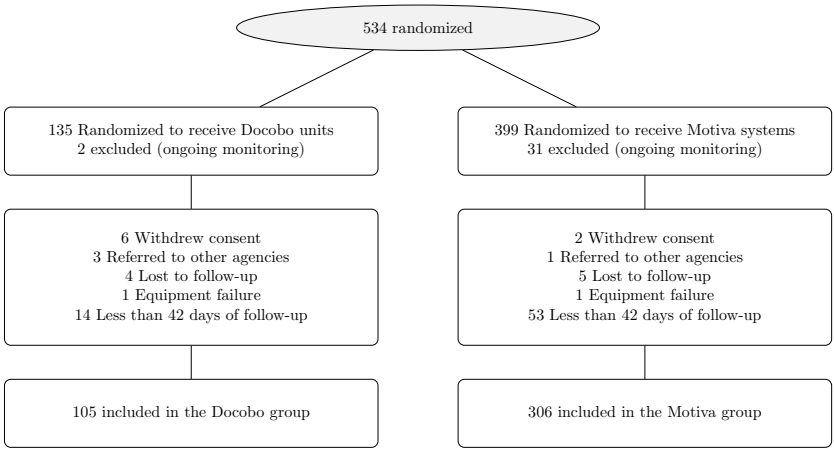

Figure A.1: *Heart failure patients on the telehealth program.*

the tendency for each parameter from day 1 to day 42. The slopes for each group were compared, and differences were assessed by means of the Kruskal-Wallis test with a 95% confidence interval. The reason to use this test is that it does not assume normality in the variables, as for instance one-way Anova. Furthermore, the quality of life and the self-care behaviour for patients with decreased/increased trends were analysed. From this analysis it was possible to determine the possible impact of the different telehealth systems on the compliance of patients to the program. The later was also associated with the amount of missing data for each telehealth system, and the amount of additional data that was transfer to the telehealth program. The additional data refers to the amount of days that patients measured themselves more times than the ones required by the program (twice a day).

The amount of interventions were also compared between the groups, as well as the differences between initial and final HADS for each group. In addition, differences in age and gender were also considered. In order to evaluate differences between proportions, the Chi-square test was implemented.

All calculations were carried out in MATLAB R2013a on an Intel(R) Core(TM) i7, 2.4GHz, 8GB RAM, running Windows 8.1.

A.3 Results

Initial common complains

The first observations correspond to the initial common issues reported by the patients during the installation of the telehealth systems and during the first days of monitoring. For instance, patients using the Docobo unit have problems adhering to the monitoring times. Additionally, due to the characteristics of the unit, they need to rely on a carer to be able to collect and transmit the data. Therefore, initial issues concerning the transmission of the data on a daily basis are commonly reported. On the other hand, patients using the Motiva system complain about the portability of the unit. Furthermore, the installation of the platform is sometimes not straightforward due to the different TV technologies available in the patients' homes. This causes problems on the transmission of the first readouts of the physiological parameters.

This information was picked up from the free text patients wrote into the surveys and from speaking to the patients. Hence, no statistics can be made on these observations.

Patient's compliance

An important factor to consider when implementing a telehealth system is the amount of data transmitted to the hospital. This is required for the design of algorithms that can predict decompensation in heart failure. In this study for instance, patients needed to measure themselves twice a day. However, it was observed that patients using the Motiva system measured themselves significantly ($p \ll 0.05$) more times a day than patients using the Docobo unit. This is shown in Figure A.2, where the amount of missing data per group is also indicated.

Effect of telehealth on the physiological parameters

The initial comparison was performed between the "baseline" parameters for both groups under investigation. Figure A.3 indicates the physiological parameters at day 1, and the age of all patients. At this point, no significant differences ($p > 0.05$) were found.

The next step was to compare the evolution of the physiological parameters for both groups. This was done using the slopes using the 42 days of monitoring

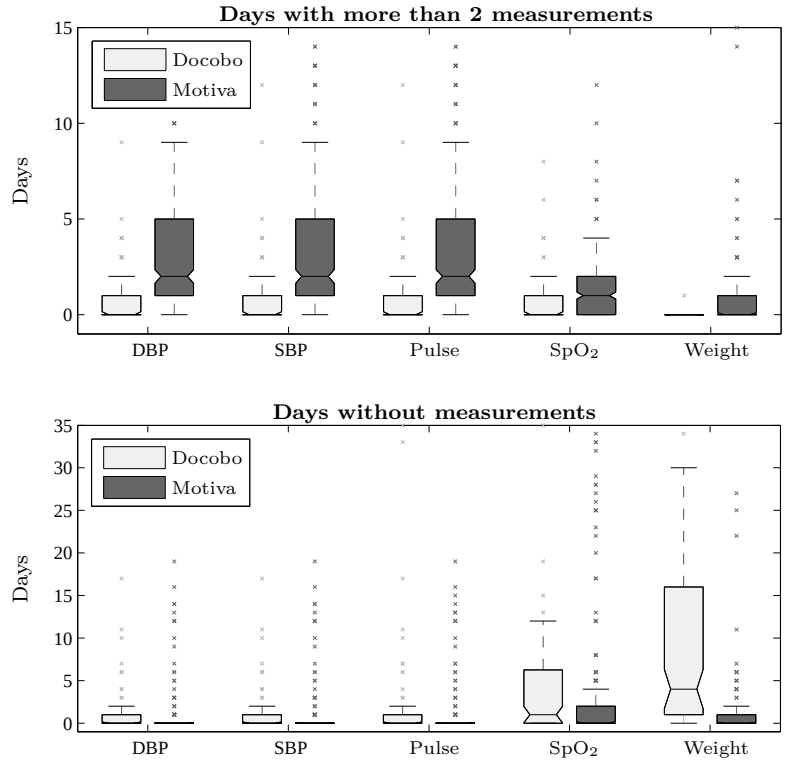

Figure A.2: *Amount of data transmitted to the hospitals for each telehealth system. Note that patients with Motiva unit seem to comply better with the programme. Differences in DBP, SBP and Pulse for more than 2 measurements a day were significantly different ($p < 0.05$). Concerning the days without measurement, Weight indicated a significant difference ($p < 0.05$). The central mark of the boxplots is the median, and the edges correspond to the 25th and 75th percentiles.*

period. Figure A.4 shows the differences for each physiological parameter. A decreasing trend was found in DBP, pulse and weight, however, only the first one resulted significant with $p = 0.001$.

Effect of telehealth on quality of life

The questionnaire used to assess changes in quality of life (QoL) consists of 5 questions that evaluate the following factors: anxiety/depression;

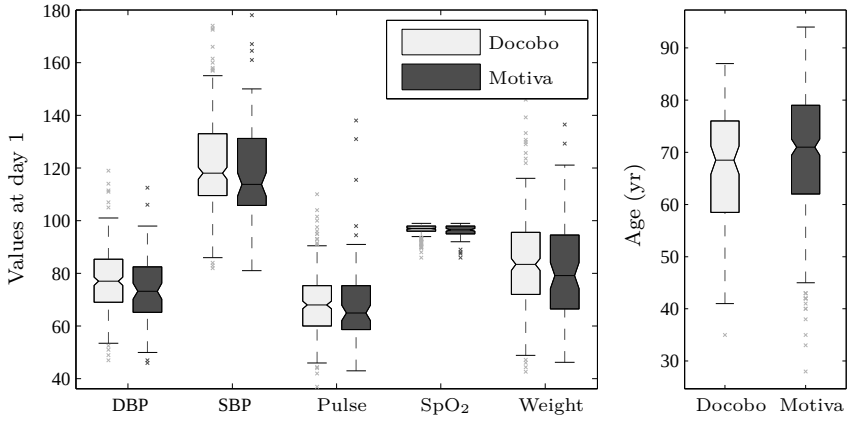

Figure A.3: *Physiological parameters at day 1 (left) and age (right) of the patients in both groups under investigation. The central mark of the boxplots is the median, and the edges correspond to the 25th and 75th percentiles.*

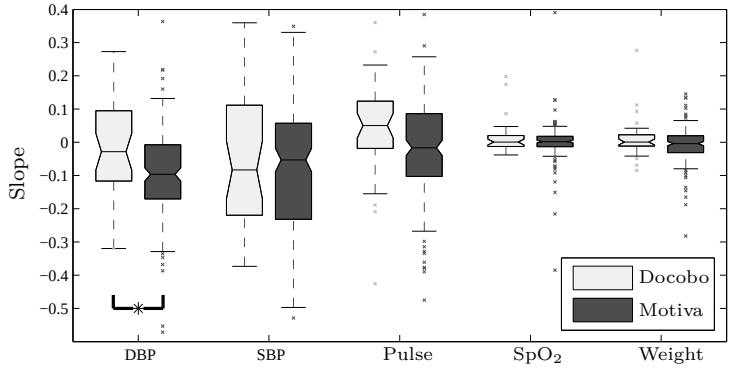

Figure A.4: *Slopes for each physiological parameter and for each telehealth system. Diastolic blood pressure tends to decrease in patients using Motiva ($p = 0.001$). Pulse and weight also tend to decrease in those patients. However, no significant differences were found. The central mark of the boxplots is the median, and the edges correspond to the 25th and 75th percentiles.*

pain/discomfort; usual activities (e.g. work, study, housework, and family or leisure activities); self-care; and mobility [90]. Each of these factors gets scored according to the severity of the problems experienced by the patients. As mentioned before, a score was obtained at the beginning and end of the monitoring period. Therefore, it was possible to determine if changes in the QoL were caused by the telehealth system, or if the user interface had any

influence on this. No significant differences ($p > 0.05$) were found between the factors at the beginning and end for either all patients together or for the two groups separately.

A different way of evaluating the anxiety/depression factor is by means of the HADS. Here, these scores at the beginning and end were also evaluated, and no significances ($p > 0.05$) were found.

When looking at patients that indicated an improved QoL at the end of the monitoring period, no relationship was found between this improvement and their evolution during 42 days of monitoring. This is to say that these patients did not show a consistent evolution in the physiological parameters.

Effect of telehealth on self-care behaviour

To assess changes on self-care behaviour, the European heart failure self-care behaviour scale [97] was used, and it consists of the following 12 items:

I1. I weigh myself every day

I2. If I get short of breath, I take it easy

I3. If my shortness of breath increases, I contact my doctor or nurse

I4. If my feet/legs become more swollen than usual, I contact my doctor or nurse

I5. If I gain 2 kg in 1 week, I contact my doctor or nurse

I6. I limit the amount of fluids I drink (not more than 1.52 l/day)

I7. I take a rest during the day

I8. If I experience increased fatigue, I contact my doctor or nurse

I9. I eat a low salt diet

I10. I take my medication as prescribed

I11. I get a flu shot every year

I12. I exercise regularly

Each item was scored from "(1) I completely agree" to "(5) I don't agree at all". As with the QoL questionnaire, two scores were obtained, one at the beginning

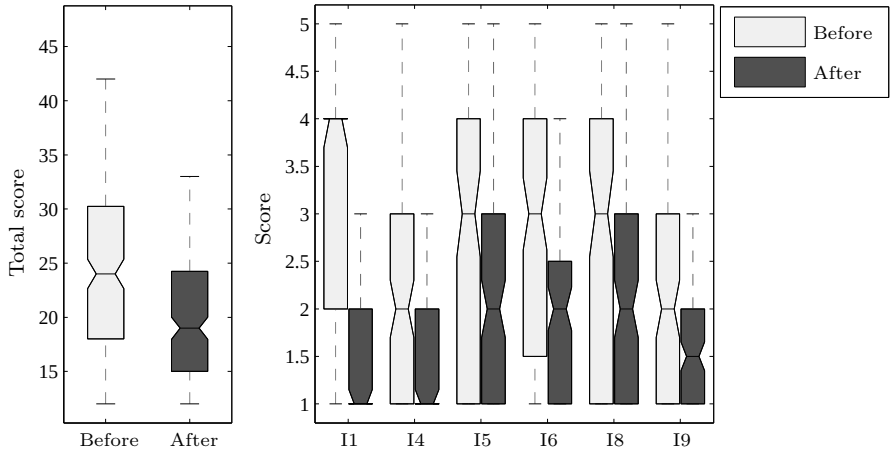

Figure A.5: *Scores of the self-care behaviour questionnaire for all the patients under investigation. All differences were significant ($p < 0.05$). The central mark of the boxplots is the median, and the edges correspond to the 25th and 75th percentiles.*

and one at the end of the telehealth program. Here significant differences ($p < 0.05$) in some of the evaluated items were found between the groups. Moreover, the telehealth program shows a general improvement in the self-care behaviour. This can be seen in Figure A.5, where the comparison between the scores at the beginning and end are indicated for the total score ($p < 0.05$) and for the items with significant differences. These results were equal for both telehealth systems, where no significant differences were found.

An important observation here is the fact that there are significant differences ($p < 0.05$) in the scores for I1 and I6, which correspond to measuring body weight every day and limiting the amount of daily fluid intake. These two factors play a key role in the monitoring of HF. Hence, when the patients with decreasing trend in DBP (136 patients), in SBP (131 patients) and in weight (91 patients) are taken, and their scores at the beginning and end are compared, it is possible to observe an improvement in the self-care behavior. This improvement is more clearly observed in the daily weighing and daily fluid intake (see Figure A.6).

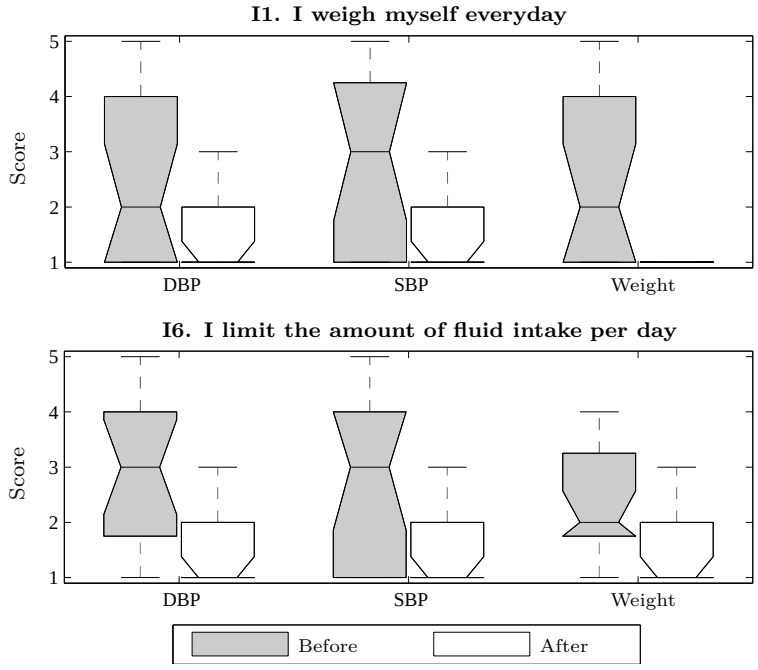

Figure A.6: *Scores of items 1 and 6. Significant differences ($p < 0.05$) were found between the scores for patients with decreasing trends in DBP, SBP and Weight. The central mark of the boxplots is the median, and the edges correspond to the 25th and 75th percentiles.*

Interventions

Different interventions were annotated in the dataset, and some significant differences were found between both telehealth systems. It was observed that patients using the Docobo unit received 2% more diuretic titration, 3% less titration of beta blockers, and 2% less titration of angiotensin-converting enzyme inhibitors than Motiva patients. Unfortunately, no information about the exact dose of medication nor differentiation between up- and down-titration was available.

Another important observation that can be made, is that 6 out of 135 patients using Docobo withdrew the programme, while only 2 out of 399 withdrew when using Motiva platform. After evaluating these proportions using the Chi-square test, a significant difference was found with ($p = 0.004$).

A.4 Discussion

The results of this study show that the characteristics of the telehealth platform have an effect on the compliance of the patients to the telehealth program. This is supported by the amount of missing data, and by the withdrawal rates of both systems. This is to say, a system offering educational information and a more user friendly interface such as the Motiva system, may increase patients' compliance and adherence to the program, and this on its turn may have long term effects on the evolution of patients' condition. However, in this study, after only 42 days of monitoring, it was already possible to see differences in the evolution of blood pressure. Nevertheless, it may be speculated that these differences could have been caused by differences in titration between both telehealth systems. Unfortunately, no information on medication nor on exact doses was available during this study. Therefore, in order to draw stronger conclusions on the effect of either telehealth platform on patients' outcome, it is necessary to implement longer trials that provide information about medication changes during the monitoring period. In addition, longer-term trials will also allow to assess changes in the quality of life of patients using telehealth.

Having considered the effect of the user interfaces on the evolution of the physiological parameters, it is now important to comment on the behavioral changes observed in this study. It was shown that a general improvement in the self-awareness, self-management and assistance-seeking was achieved after the telehealth program. This is a remarkable finding, because it proves that by implementing a telehealth program, it is possible to increase the awareness of patients of the importance of self-managing their conditions. For instance, daily weighing and limitation of daily fluid intake became an important factor for patients during the monitoring period. Hence, significant improvements in a long term can be achieved, such as better adherence to telehealth programs that can result in a reduction of heart failure decompensations.

Notwithstanding patients were more aware of the importance of monitoring their weight, it was observed that measurements of this parameter were missing for most of the patients. In addition, even patients using the telehealth system more times than needed during the day, showed a reduction in the amount of daily measurements of weight. This was also the case for SpO_2. These findings are relevant, since weight is a very important parameter in heart failure, and it has been shown to be probably the most useful physiological parameter to predict decompensation in heart failure [222] [47]. Therefore, the fact that patients tend to forget, or simply avoid to take daily measurements of weight, must be considered when developing telehealth systems, and when designing algorithms for risk prediction. It is worth noting that attention must be paid to the way weight measurements are integrated into the telehealth system. For

example, a way where all sensors are integrated into one compact platform could be developed, as well as new protocols to attract and motivate patients to monitor themselves on a daily basis.

Two important points must be considered here, first is the fact that the Motiva unit uses a "well-known" interface, namely the television, and second is the mean age of the population under investigation, 69.1yr±12.6yr. This is the case because elderly people of today are more used to the TV technology, while the heart failure population of 10 or 20 years will already be used to more portable devices like standalone units. Therefore, it is obvious to think that patients using the Motiva system are reminded more often to take daily measurements, while patients using standalone units tend to comply less with the monitoring times, only because of technological differences. However, these results will probably change in the future due to the fact that society is relying more and more on portable devices. These findings are in agreement with the increasing needs for personalized health informatics [160], where each telehealth system is specifically tuned to work for each particular patient.

A.5 Conclusions

This study assessed changes in the evolution of physiological parameters such as blood pressure, pulse, blood oxygen saturation and weight, and in the self-care behaviour caused by the implementation of a telehealth program on heart failure patients. Two telehealth systems were used to monitor patients, namely the Motiva platform offered by Philips, and the Docobo unit. The main difference between both is the user interface, and the amount of educational information that is presented to the patients. Significant differences were found between both systems, as well as a possible positive effect of education on the compliance and adherence of the patients to the telehealth programme. This effect will need to be confirmed in future studies, where information of how often patients access educational material is available.

A general improvement in the self-care behaviour was achieved by the use of telehealth, which appeared to also have a positive effect on the evolution of the physiological parameters. Therefore, an improved heart failure outcome can be achieved, which can reduce the load in the healthcare system. However, this is something that needs to be confirmed by longer-term trials, where hospitalizations are included as endpoints.

One of the main limitations of this study was the lack of information about types of medication and dose changes. Hence, no links between changes in the physiological parameters and medication could be made. In addition, the short-

term duration of this study, namely 42 days per patient, limited the possibility to observe clinical events. Therefore, no conclusions about the effect of each telehealth system on the outcome could be reached.

The results obtained in this study allow to define some practicalities and technicalities that need to be considered when designing and implementing a telehealth system.

1. The telehealth device needs to be considered, because it has a direct influence in the compliance of the patients to the program. For some patients the mobility of the unit is important, while for others the fact that they do not need to use/learn a "new" technology already improves the compliance to the monitoring times.

2. The educational and additional information presented to the patients should be included in the telehealth program.

3. A way to attract and motivate patients to take daily measurements of weight needs to be reconsidered. This is necessary for the design of algorithms that predict decompensation due to heart failure.

4. It is of paramount importance to include information about medication and doses in order to assess the direct effect of telehealth on the evolution and final outcome of the patients.

Appendix B

Graphical user interfaces

Three different graphical user interfaces (GUIs) were developed during this research. The first one includes the artifact detection algorithm based on the autocorrelation function (ACF) of the ECG segments. It allows to manually set the threshold of the weights computed using the degrees of the ACF. It also gives the possibility to select the length of the ECG segments to be analyzed, and three different filters (high-pass, low-pass, notch filter) are available to be applied to the original signals. This GUI is shown in Figure B.1.

The second GUI allows to visually inspect the positions of the R-peaks (see Figure B.2). Filters are also available in this GUI. Furthermore, each time a heart beat is added or removed, the tachogram is updated automatically.

The last GUI shown in Figure B.3 is dedicated to the computation of the EDR signals using the R-peak amplitude, PCA, and kPCA. It also allows to compute the different standard HRV parameters, and additionally, it computes the frequency parameters LF and HF, using wavelet transform.

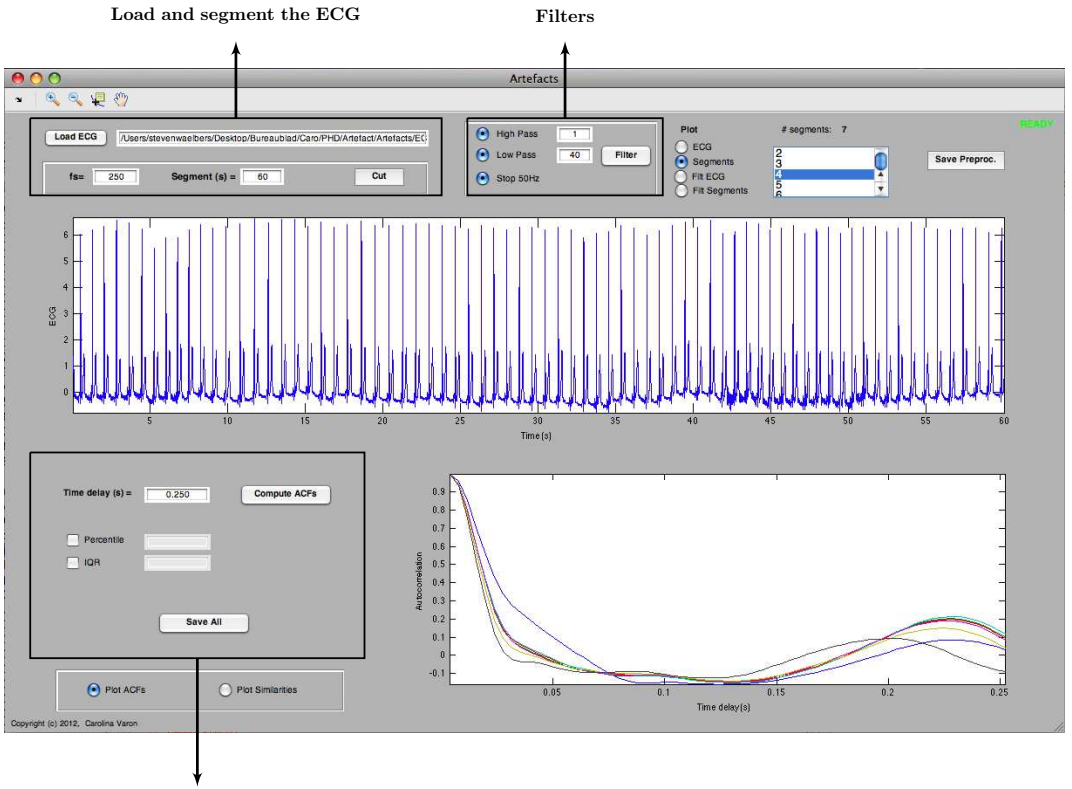

Figure B.1: Graphical user interface for the computation of the contamination level of segments of the ECG.

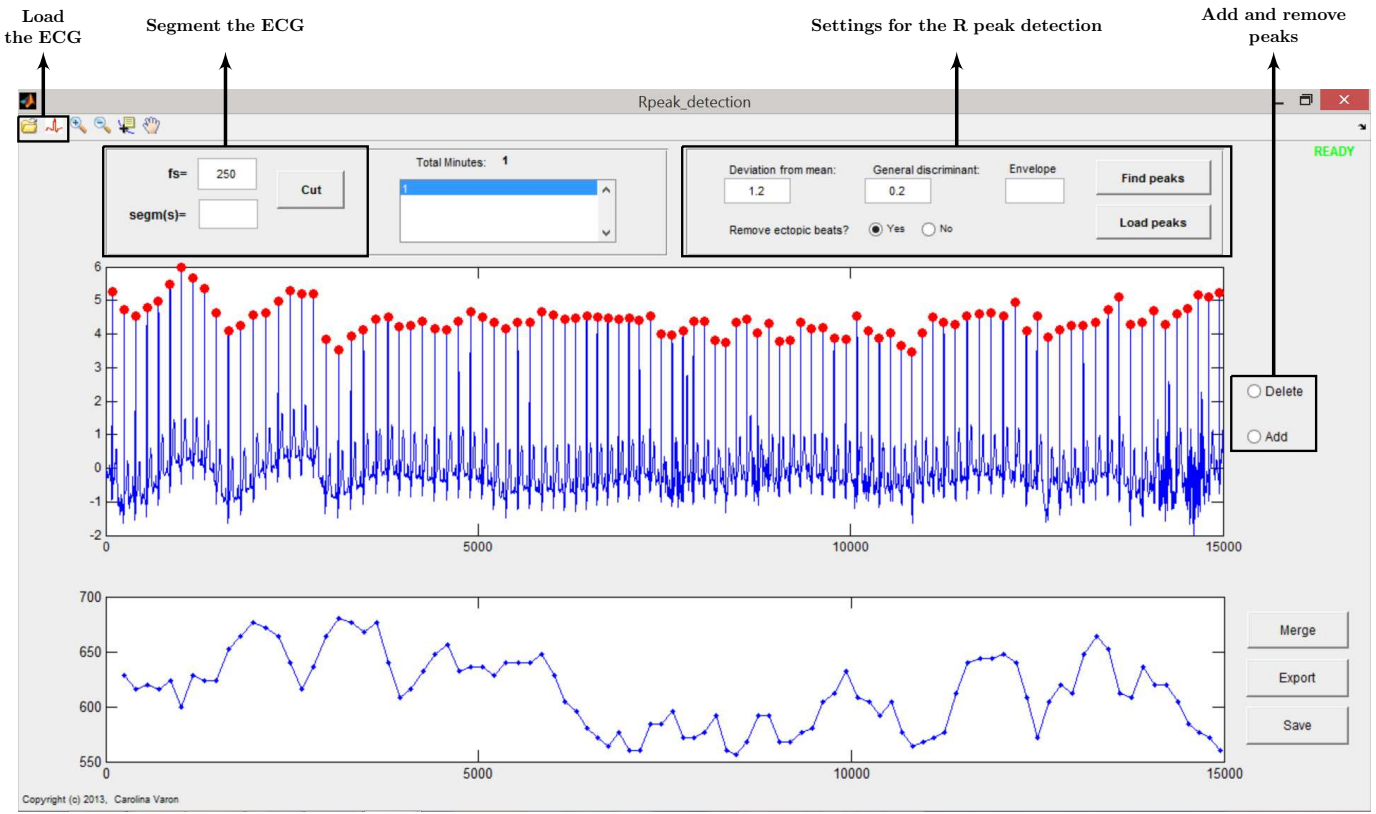

Figure B.2: Graphical user interface for the manual inspection of R-peaks.

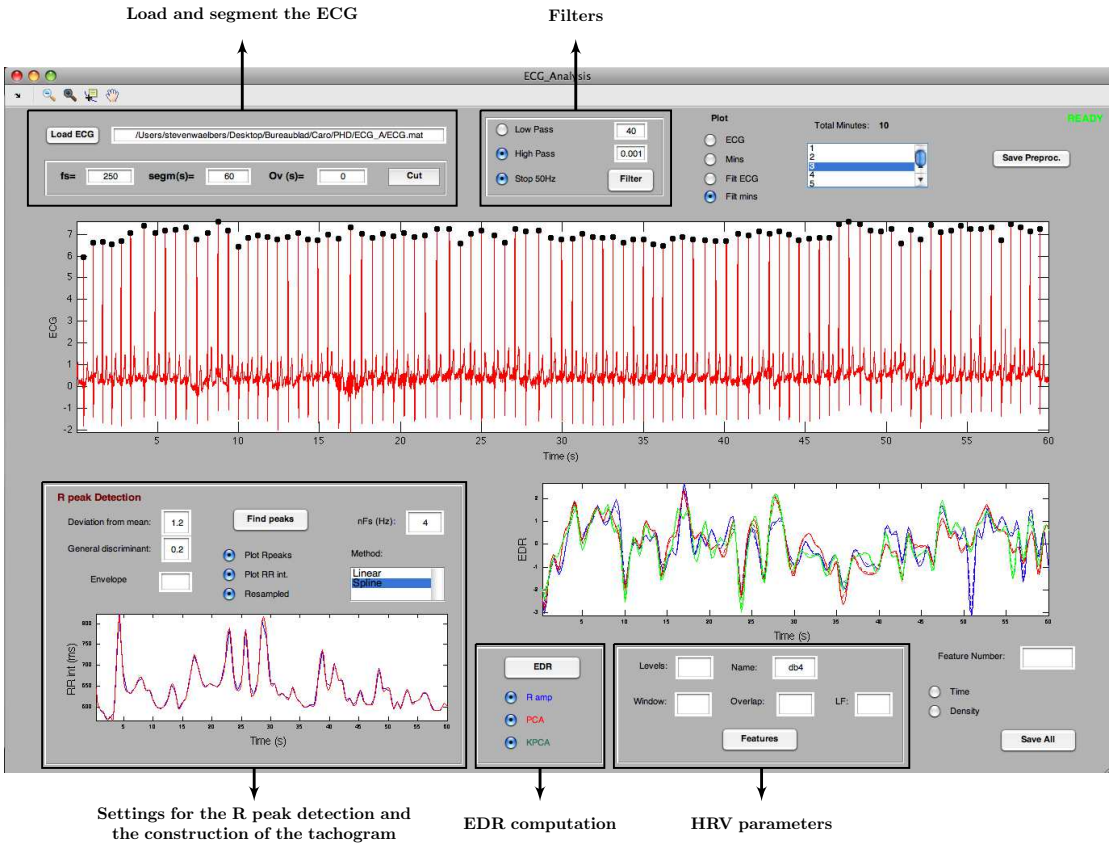

Figure B.3: Graphical user interface for the computation of the EDR and HRV parameters.

Bibliography

[1] National institute for health and care excellence. Services for people with chronic heart failure, 2014.

[2] World health organization. WHO|Epilepsy, 2015.

[3] Acharya, U. R., Joseph, K. P., Kannathal, N., Lim, C. M., and Suri, J. S. Heart rate variability: a review. *Medical and Biological Engineering and Computing 44*, 12 (2006), 1031–1051.

[4] Afonso, V. X., Tompkins, W. J., Nguyen, T. Q., and Luo, S. ECG beat detection using filter banks. *IEEE Transactions on Biomedical Engineering 46*, 2 (1999), 192–202.

[5] Al-Aweel, I., Krishnamurthy, K., Hausdorff, J., Mietus, J., Ives, J., Blum, A., Schomer, D., and Goldberger, A. Postictal heart rate oscillations in partial epilepsy. *Neurology 53*, 7 (1999), 1590–1590.

[6] Alzate, C., and Suykens, J. A. K. A weighted kernel PCA formulation with out-of-sample extensions for spectral clustering methods. In *International Joint Conference on Neural Networks (IJCNN'06)* (2006), IEEE, pp. 138–144.

[7] Alzate, C., and Suykens, J. A. K. Kernel component analysis using an epsilon-insensitive robust loss function. *IEEE Transactions on Neural Networks 19*, 9 (2008), 1583–1598.

[8] Alzate, C., and Suykens, J. A. K. A regularized kernel CCA contrast function for ICA. *Neural Networks 21*, 2 (2008), 170–181.

[9] Alzate, C., and Suykens, J. A. K. Sparse kernel models for spectral clustering using the incomplete cholesky decomposition. In *IEEE International Joint Conference on Neural Networks* (2008), IEEE, pp. 3556–3563.

[10] Alzate, C., and Suykens, J. A. K. Multiway spectral clustering with out-of-sample extensions through weighted kernel PCA. *IEEE Transactions on Pattern Analysis and Machine Intelligence 32*, 2 (2010), 335–347.

[11] Alzate, C., and Suykens, J. A. K. Sparse kernel spectral clustering models for large-scale data analysis. *Neurocomputing 74*, 9 (2011), 1382–1390.

[12] Arzeno, N. M., Deng, Z.-D., and Poon, C.-S. Analysis of first-derivative based QRS detection algorithms. *IEEE Transactions on Biomedical Engineering 55*, 2 (2008), 478–484.

[13] Augustyniak, P. Continuous noise estimation using time-frequency ECG representation. In *Computing in Cardiology, 2011* (2011), IEEE, pp. 133–136.

[14] Bach, F. R., and Jordan, M. I. Kernel independent component analysis. *The Journal of Machine Learning Research 3* (2003), 1–48.

[15] Bancaud, J., Henriksen, O., Rubio-Donnadieu, F., Seino, M., Dreifuss, F. E., and Penry, J. K. Proposal for revised clinical and electroencephalographic classification of epileptic seizures. *Epilepsia 22*, 22 (1981), 489–501.

[16] Barros, A. K., Mansour, A., and Ohnishi, N. Removing artifacts from electrocardiographic signals using independent components analysis. *Neurocomputing 22*, 1 (1998), 173–186.

[17] Bartsch, R., Kantelhardt, J. W., Penzel, T., and Havlin, S. Experimental evidence for phase synchronization transitions in the human cardiorespiratory system. *Physical review letters 98*, 5 (2007), 054102.

[18] Bartsch, R. P., Liu, K. K., Ma, Q. D., and Ivanov, P. C. Three independent forms of cardio-respiratory coupling: Transitions across sleep stages.

[19] Bartsch, R. P., Schumann, A. Y., Kantelhardt, J. W., Penzel, T., and Ivanov, P. C. Phase transitions in physiologic coupling. *Proceedings of the National Academy of Sciences 109*, 26 (2012), 10181–10186.

[20] Bashan, A., Bartsch, R. P., Kantelhardt, J. W., Havlin, S., and Ivanov, P. C. Network physiology reveals relations between network topology and physiological function. *Nature communications 3* (2012), 702.

[21] Bauer, A., Kantelhardt, J. W., Barthel, P., Schneider, R., Mäkikallio, T., Ulm, K., Hnatkova, K., Schömig, A., Huikuri, H., Bunde, A., et al. Deceleration capacity of heart rate as a predictor of mortality after myocardial infarction: cohort study. *The lancet 367*, 9523 (2006), 1674–1681.

[22] Bauer, A., Kantelhardt, J. W., Bunde, A., Barthel, P., Schneider, R., Malik, M., and Schmidt, G. Phase-rectified signal averaging detects quasi-periodicities in non-stationary data. *Physica A: Statistical Mechanics and its Applications 364* (2006), 423–434.

[23] Bauer, A., Morley-Davies, A., Barthel, P., Müller, A., Ulm, K., Malik, M., and Schmidt, G. Bivariate phase-rectified signal averaging for assessment of spontaneous baroreflex sensitivity: pilot study of the technology. *Journal of electrocardiology 43*, 6 (2010), 649–653.

[24] Baumgartner, C., Lurger, S., and Leutmezer, F. Autonomic symptoms during epileptic seizures. *Epileptic disorders: international epilepsy journal with videotape 3*, 3 (2001), 103–116.

[25] Bellman, R. *Adaptive control processes: a guided tour*. Princeton University Press Princeton, 1961.

[26] Benitez, D., Gaydecki, P., Zaidi, A., and Fitzpatrick, A. The use of the hilbert transform in ECG signal analysis. *Computers in biology and medicine 31*, 5 (2001), 399–406.

[27] Berg, A. T., Berkovic, S. F., Brodie, M. J., Buchhalter, J., Cross, J. H., Van Emde Boas, W., Engel, J., French, J., Glauser, T. A., Mathern, G. W., et al. Revised terminology and concepts for organization of seizures and epilepsies: report of the ILAE commission on classification and terminology, 2005-2009. *Epilepsia 51*, 4 (2010), 676–685.

[28] Berntson, G. G., Cacioppo, J. T., and Quigley, K. S. Respiratory sinus arrhythmia: autonomic origins, physiological mechanisms, and psychophysiological implications. *Psychophysiology 30*, 2 (1993), 183–196.

[29] Berntson, G. G., Thomas Bigger, J., Eckberg, D. L., Grossman, P., Kaufmann, P. G., Malik, M., Nagaraja, H. N., Porges, S. W., Saul, J. P., Stone, P. H., and Van der Molen, M. W. Heart rate variability: Origins, methods, and interpretive caveats. *Psychophysiology 34*, 6 (1997), 623–648.

[30] Beyer, K., Goldstein, J., Ramakrishnan, R., and Shaft, U. When is "nearest neighbor" meaningful? In *Database Theory-ICDT'99*. Springer, 1999, pp. 217–235.

[31] Billman, G. E., and Hoskins, R. S. Time-series analysis of heart rate variability during submaximal exercise. evidence for reduced cardiac vagal tone in animals susceptible to ventricular fibrillation. *Circulation 80*, 1 (1989), 146–157.

[32] Blumenfeld, H. Consciousness and epilepsy: why are patients with absence seizures absent? *Progress in brain research 150* (2005), 271–603.

[33] Blumhardt, L., Smith, P., and Owen, L. Electrocardiographic accompaniments of temporal lobe epileptic seizures. *The Lancet 327*, 8489 (1986), 1051–1056.

[34] Boudaoud et al, S. Corrected integral shape averaging applied to obstructive sleep apnea detection from the electrocardiogram. *Journal on Advances in Signal Processing 2007*, 1 (2007), 032570.

[35] Boyne, J. J., Vrijhoef, H. J., Crijns, H. J., De Weerd, G., Kragten, J., and Gorgels, A. P. Tailored telemonitoring in patients with heart failure. *Effects of telemonitoring in patients with heart failure 14*, 7 (2013), 69.

[36] Brás, S., Ferreira, N., and da Silva Cunha, J. P. ECG artefact detection algorithm-an algorithm to improve long-term ECG analysis. In *BIOSIGNALS* (2012), pp. 329–333.

[37] Brunner-La Rocca, H. P., Buser, P. T., Schindler, R., Bernheim, A., Rickenbacher, P., and Pfisterer, M. Management of elderly patients with congestive heart failure-design of the trial of intensified versus standard medical therapy in elderly patients with congestive heart failure (TIME-CHF). *American heart journal 151*, 5 (2006), 949–955.

[38] Bsoul, M., Minn, H., and Tamil, L. Apnea medassist: real-time sleep apnea monitor using single-lead ECG. *IEEE Transactions on Information Technology in Biomedicine 15*, 3 (2011), 416–427.

[39] Bui, A. L., Horwich, T. B., and Fonarow, G. C. Epidemiology and risk profile of heart failure. *Nature Reviews Cardiology 8*, 1 (2010), 30–41.

[40] Caicedo, A., Tachtsidis, I., Papademetriou, M., and Van Huffel, S. Decoupling the influence of systemic variables in the peripheral

and cerebral haemodynamics during ECMO procedure by means of oblique and orthogonal subspace projections. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2012), IEEE, pp. 6153–6156.

[41] Caicedo, A., Varon, C., Hunyadi, B., Papademetriou, M., Tachtsidis, I., and Van Huffel, S. Signal decomposition based on wavelet regression and oblique subspace projections: Theoretical framework, algorithms and applications in biomedical signal processing. *IEEE Transactions on Biomedical Engineering a revised version has been submitted* (2015).

[42] Caicedo, A., Varon, C., and Van Huffel, S. Cardiopulmonary coupling in sleep apnea assessed by means of orthogonal subspace projection. In *the 5th International Conference of Advances in Medical Signal and Information Processing (MEDSIP)* (2012).

[43] Campana, L. M., Owens, R. L., Clifford, G. D., Pittman, S. D., and Malhotra, A. Phase-rectified signal averaging as a sensitive index of autonomic changes with aging. *Journal of Applied Physiology 108*, 6 (2010), 1668–1673.

[44] Candès, E. J., Li, X., Ma, Y., and Wright, J. Robust principal component analysis? *Journal of the ACM (JACM) 58*, 3 (2011), 11.

[45] Caples, S. M., Garcia-Touchard, A., and Somers, V. K. Sleep-disordered breathing and cardiovascular risk. *Sleep 30*, 3 (2007), 291.

[46] Charkoudian, N., and Rabbitts, J. A. Sympathetic neural mechanisms in human cardiovascular health and disease. In *Mayo Clinic Proceedings* (2009), vol. 84, Elsevier, pp. 822–830.

[47] Chaudhry, S. I., Wang, Y., Concato, J., Gill, T. M., and Krumholz, H. M. Patterns of weight change preceding hospitalization for heart failure. *Circulation 116*, 14 (2007), 1549–1554.

[48] Chen, Y.-W., and Lin, C.-J. Combining SVMs with various feature selection strategies. In *Feature extraction*. Springer, 2006, pp. 315–324.

[49] Chroni, E., Sirrou, V., Trachani, E., Sakellaropoulos, G. C., and Polychronopoulos, P. Interictal alterations of cardiovagal function in chronic epilepsy. *Epilepsy research 83*, 2 (2009), 117–123.

[50] Chung, F. R. *Spectral graph theory*, vol. 92. American Mathematical Soc., 1997.

[51] Clifford, G. D., Azuaje, F., and McSharry, P. *Advanced methods and tools for ECG data analysis*. Artech House, Inc., 2006.

[52] Clifford, G. D., and Tarassenko, L. Quantifying errors in spectral estimates of HRV due to beat replacement and resampling. *IEEE Transactions on Biomedical Engineering 52*, 4 (2005), 630–638.

[53] Cortes, C., and Vapnik, V. Support vector networks. *Machine learning 20*, 3 (1995), 273–297.

[54] De Chazal, P., Heneghan, C., Sheridan, E., Reilly, R., Nolan, P., and O'Malley, M. Automated processing of the single-lead electrocardiogram for the detection of obstructive sleep apnoea. *IEEE Transactions on Biomedical Engineering 50*, 6 (2003), 686–696.

[55] de Chazal, P., Penzel, T., and Heneghan, C. Automated detection of obstructive sleep apnoea at different time scales using the electrocardiogram. *Physiological measurement 25*, 4 (2004), 967.

[56] DeGiorgio, C. M., and DeGiorgio, A. C. SUDEP and heart rate variability. *Epilepsy research 90*, 3 (2010), 309.

[57] DeGiorgio, C. M., Miller, P., Meymandi, S., Chin, A., Epps, J., Gordon, S., Gornbein, J., and Harper, R. M. RMSSD, a measure of vagus-mediated heart rate variability, is associated with risk factors for SUDEP: the SUDEP-7 inventory. *Epilepsy & Behavior 19*, 1 (2010), 78–81.

[58] Dendale, P., De Keulenaer, G., Troisfontaines, P., Weytjens, C., Mullens, W., Elegeert, I., Ector, B., Houbrechts, M., Willekens, K., and Hansen, D. Effect of a telemonitoring-facilitated collaboration between general practitioner and heart failure clinic on mortality and rehospitalization rates in severe heart failure: the TEMA-HF 1 (telemonitoring in the management of heart failure) study. *European journal of heart failure 14*, 3 (2012), 333–340.

[59] Devinsky, O. Effects of seizures on autonomic and cardiovascular function. *Epilepsy currents 4*, 2 (2004), 43–46.

[60] Eckert, D. J., Jordan, A. S., Merchia, P., and Malhotra, A. Central sleep apnea: pathophysiology and treatment. *Chest Journal 131*, 2 (2007), 595–607.

[61] Elgendi, M., Eskofier, B., Dokos, S., and Abbott, D. Revisiting QRS detection methodologies for portable, wearable, battery-operated, and wireless ECG systems. *PloS one 9*, 1 (2014), e84018.

[62] Englot, D. J., and Blumenfeld, H. Consciousness and epilepsy: why are complex-partial seizures complex? *Progress in brain research 177* (2009), 147–170.

[63] Faes, L., Nollo, G., Jurysta, F., and Marinazzo, D. Information dynamics of brain-heart physiological networks during sleep. *New J. Phys in press* (2014).

[64] Faes, L., and Porta, A. Conditional entropy-based evaluation of information dynamics in physiological systems. In *Directed information measures in neuroscience*. Springer, 2014, pp. 61–86.

[65] Faes, L., Widjaja, D., Van Huffel, S., and Nollo, G. Investigating cardiac and respiratory determinants of heart rate variability in an information-theoretic framework. In *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE* (2014), IEEE, pp. 6020–6023.

[66] Ferri, R., Curzi-Dascalova, L., Arzimanoglou, A., Bourgeois, M., Beaud, C., Elia, M., Musumeci, S. A., Tripodi, M., et al. Heart rate variability during sleep in children with partial epilepsy. *Journal of sleep research 11*, 2 (2002), 153–160.

[67] Fine, S., and Scheinberg, K. Efficient SVM training using low-rank kernel representations. *The Journal of Machine Learning Research 2* (2002), 243–264.

[68] Fisher, R. A. The use of multiple measurements in taxonomic problems. *Annals of eugenics 7*, 2 (1936), 179–188.

[69] Fisher, R. S., Acevedo, C., Arzimanoglou, A., Bogacz, A., Cross, J. H., Elger, C. E., Engel, J., Forsgren, L., French, J. A., Glynn, M., et al. ILAE official report: A practical clinical definition of epilepsy. *Epilepsia 55*, 4 (2014), 475–482.

[70] Fisher, R. S., Boas, W. v. E., Blume, W., Elger, C., Genton, P., Lee, P., and Engel, J. Epileptic seizures and epilepsy: definitions proposed by the international league against epilepsy (ILAE) and the international bureau for epilepsy (ibe). *Epilepsia 46*, 4 (2005), 470–472.

[71] Fraden, J., and Neuman, M. QRS wave detection. *Medical and Biological Engineering and computing 18*, 2 (1980), 125–132.

[72] Franco, P., Szliwowski, H., Dramaix, M., and Kahn, A. Polysomnographic study of the autonomic nervous system in potential victims of sudden infant death syndrome. *Clinical autonomic research 8*, 5 (1998), 243–249.

[73] Friesen, G. M., Jannett, T. C., Jadallah, M. A., Yates, S. L., Quint, S. R., and Nagle, H. T. A comparison of the noise sensitivity of nine QRS detection algorithms. *IEEE Transactions on Biomedical Engineering 37*, 1 (1990), 85–98.

[74] Froning, J. N., Olson, M. D., and Froelicher, V. F. Problems and limitations of ECG baseline estimation and removal using a cubic spline technique during exercise ECG testing: recommendations for proper implementation. *Journal of electrocardiology 21* (1988), S149–S157.

[75] Gang, Y., and Malik, M. Heart rate variability analysis in general medicine. *Indian pacing and electrophysiology journal 3*, 1 (2003), 34.

[76] Garcia, M., DÂĂŹGiano, C., Estelles, S., Leiguarda, R., and Rabinowicz, A. Ictal tachycardia: its discriminating potential between temporal and extratemporal seizure foci. *Seizure 10*, 6 (2001), 415–419.

[77] García-González, M., Argelagós, A., Fernández-Chimeno, M., and Ramos-Castro, J. Differences in QRS locations due to ECG lead: relationship with breathing. In *XIII Mediterranean Conference on Medical and Biological Engineering and Computing 2013* (2014), Springer, pp. 962–964.

[78] García-González, M., Argelagós, A., Fernández-Chimeno, M., and Ramos-Castro, J. Differences in QRS locations due to ecg lead: relationship with breathing. In *XIII Mediterranean Conference on Medical and Biological Engineering and Computing 2013* (2014), Springer, pp. 962–964.

[79] Garcia III, A. J., Koschnitzky, J. E., Dashevskiy, T., and Ramirez, J.-M. Cardiorespiratory coupling in health and disease. *Autonomic Neuroscience 175*, 1 (2013), 26–37.

[80] Gigout, S., Louvel, J., Rinaldi, D., Martin, B., and Pumain, R. Thalamocortical relationships and network synchronization in a new genetic model "in mirror" for absence epilepsy. *Brain research 1525* (2013), 39–52.

[81] Girolami, M. Orthogonal series density estimation and the kernel eigenvalue problem. *Neural Computation 14*, 3 (2002), 669–688.

[82] Goldberger, A. L., Amaral, L. A., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., and Stanley, H. E. Physiobank, physiotoolkit, and physionet components of a new research resource for complex physiologic signals. *Circulation 101*, 23 (2000), e215–e220.

[83] Golub, G. H., and van Van Loan, C. F. Matrix computations (johns hopkins studies in mathematical sciences).

[84] Guyton, A., and Hall, J. Textbook of medical physiology, 11th.

[85] Hahn, G., Sipinkova, I., Baisch, F., and Hellige, G. Changes in the thoracic impedance distribution under different ventilatory conditions. *Physiological measurement 16*, 3A (1995), A161.

[86] Handforth, A., DeGiorgio, C., Schachter, S., Uthman, B., Naritoku, D., Tecoma, E., Henry, T., Collins, S., Vaughn, B., Gilmartin, R., et al. Vagus nerve stimulation therapy for partial-onset seizures a randomized active-control trial. *Neurology 51*, 1 (1998), 48–55.

[87] Harnod, T., Yang, C. C., Hsin, Y.-L., Shieh, K.-R., Wang, P.-J., and Kuo, T. B. Heart rate variability in children with refractory generalized epilepsy. *Seizure 17*, 4 (2008), 297–301.

[88] Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J., and Tibshirani, R. *The elements of statistical learning*, vol. 2. Springer, 2009.

[89] Hayano, J., and Yasuma, F. Hypothesis: respiratory sinus arrhythmia is an intrinsic resting function of cardiopulmonary system. *Cardiovascular research 58*, 1 (2003), 1–9.

[90] Herdman, M., Gudex, C., Lloyd, A., Janssen, M., Kind, P., Parkin, D., Bonsel, G., and Badia, X. Development and preliminary testing of the new five-level version of EQ-5D. *Quality of life research 20*, 10 (2011), 1727–1736.

[91] Hinneburg, A., Aggarwal, C. C., and Keim, D. A. What is the nearest neighbor in high dimensional spaces? In *Proceedings of 26th International Conference on Very Large Data Bases, September 10-14, 2000, Cairo, Egypt* (2000), pp. 506–515.

[92] Horn, J. L. A rationale and test for the number of factors in factor analysis. *Psychometrika 30*, 2 (1965), 179–185.

[93] Huhn, E., Lobmaier, S., Fischer, T., Schneider, R., Bauer, A., Schneider, K., and Schmidt, G. New computerized fetal heart rate analysis for surveillance of intrauterine growth restriction. *Prenatal diagnosis 31*, 5 (2011), 509–514.

[94] Hull, J. J. A database for handwritten text recognition research. *IEEE Transactions on Pattern Analysis and Machine Intelligence 16*, 5 (1994), 550–554.

[95] Ivanov, P. C., and Bartsch, R. P. Network physiology: mapping interactions between networks of physiologic networks. In *Networks of networks: the last frontier of complexity*. Springer, 2014, pp. 203–222.

[96] Iyengar, N., Peng, C., Morin, R., Goldberger, A. L., and Lipsitz, L. A. Age-related alterations in the fractal scaling of cardiac interbeat interval dynamics. *American Journal of Physiology-Regulatory Integrative and Comparative Physiology 40*, 4 (1996), R1078.

[97] Jaarsma, T., Strömberg, A., Mårtensson, J., and Dracup, K. Development and testing of the european heart failure self-care behaviour scale. *European Journal of Heart Failure 5*, 3 (2003), 363–370.

[98] Jain, A. K., Murty, M. N., and Flynn, P. J. Data clustering: a review. *ACM computing surveys (CSUR) 31*, 3 (1999), 264–323.

[99] Jansen, K., Vandeput, S., Van Huffel, S., and Lagae, L. Cardiac autonomic dysfunction in West syndrome. *Epilepsy research 102*, 3 (2012), 167–172.

[100] Jansen, K., Varon, C., Huffel, S., and Lagae, L. Early respiratory dysfunction as a biomarker for epileptic encephalopathy. *Acta Neurologica Scandinavica 128*, 6 (2013), 381–385.

[101] Jansen, K., Varon, C., Van Huffel, S., and Lagae, L. Ictal and interictal respiratory changes in temporal lobe and absence epilepsy in childhood. *Epilepsy research 106*, 3 (2013), 410–416.

[102] Jansen, K., Varon, C., Van Huffel, S., and Lagae, L. Peri-ictal ECG changes in childhood epilepsy: implications for detection systems. *Epilepsy & Behavior 29*, 1 (2013), 72–76.

[103] Jolliffe, I. *Principal component analysis*. Wiley Online Library, 2005.

[104] Jong, P., Vowinckel, E., Liu, P. P., Gong, Y., and Tu, J. V. Prognosis and determinants of survival in patients newly hospitalized for heart failure: a population-based study. *Archives of internal medicine 162*, 15 (2002), 1689–1694.

[105] Jørgensen, K. W., and Hansen, L. K. Model selection for gaussian kernel PCA denoising. *IEEE Transactions on Neural Networks and Learning Systems 23*, 1 (2012), 163–168.

[106] Kantelhardt, J. W., Bauer, A., Schumann, A. Y., Barthel, P., Schneider, R., Malik, M., and Schmidt, G. Phase-rectified signal averaging for the detection of quasi-periodicities and the prediction of cardiovascular risk. *Chaos: An Interdisciplinary Journal of Nonlinear Science 17*, 1 (2007), 015112.

[107] Kaufman, L., and Rousseeuw, P. J. *Finding groups in data: an introduction to cluster analysis*, vol. 344. John Wiley & Sons, 2009.

[108] Kesper, K., Canisius, S., Penzel, T., Ploch, T., and Cassel, W. ECG signal analysis for the assessment of sleep-disordered breathing and sleep pattern. *Medical & biological engineering & computing 50*, 2 (2012), 135–144.

[109] Khandoker, A. H., Palaniswami, M., and Karmakar, C. K. Support vector machines for automated recognition of obstructive sleep apnea syndrome from ECG recordings. *IEEE Transactions on Information Technology in Biomedicine 13*, 1 (2009), 37–48.

[110] Kluge, K. A., Harper, R. M., Schechtman, V. L., Wilson, A. J., Hoffman, H. J., and Southall, D. P. Spectral analysis assessment of respiratory sinus arrhythmia in normal infants and infants who subsequently died of sudden infant death syndrome. *Pediatric research 24*, 6 (1988), 677–682.

[111] Köhler, B.-U., Hennig, C., and Orglmeister, R. The principles of software QRS detection. *IEEE Engineering in Medicine and Biology Magazine 21*, 1 (2002), 42–57.

[112] Kong, X., Li, K., Yang, Q., Wenyin, L., and Yang, M.-H. A new image quality metric for image auto-denoising. In *IEEE International Conference on Computer Vision (ICCV)* (2013), IEEE, pp. 2888–2895.

[113] Köppen, M. The curse of dimensionality. In *5th Online World Conference on Soft Computing in Industrial Applications (WSC5)* (2000), pp. 4–8.

[114] Kowey, P. R., and Kocovic, D. Z. Ambulatory electrocardiographic recording. *Circulation 108*, 5 (2003), e31–e33.

[115] Krishnan, R., Natarajan, B., and Warren, S. Two-stage approach for detection and reduction of motion artifacts in photoplethysmographic data. *IEEE Transactions on Biomedical Engineering 57*, 8 (2010), 1867–1876.

[116] Kulessa, B., Srokowski, T., and Drozdz, S. Long-time autocorrelation function of ECG signal for healthy versus diseased human heart. *Acta Physica Polonica B 34* (2003), 3.

[117] Kwok, J.-Y., and Tsang, I. W. The pre-image problem in kernel methods. *IEEE Transactions on Neural Networks 15*, 6 (2004), 1517–1525.

[118] La Rovere, M. T., Pinna, G. D., Maestri, R., Mortara, A., Capomolla, S., Febo, O., Ferrari, R., Franchini, M., Gnemmi, M., Opasich, C., et al. Short-term heart rate variability strongly predicts sudden cardiac death in chronic heart failure patients. *circulation 107*, 4 (2003), 565–570.

[119] Langley, P., Bowers, E. J., and Murray, A. Principal component analysis as a tool for analyzing beat-to-beat changes in ecg features: application to ECG-derived respiration. *IEEE Transactions on Biomedical Engineering 57*, 4 (2010), 821–829.

[120] Larsen, P., Tzeng, Y., Sin, P., and Galletly, D. Respiratory sinus arrhythmia in conscious humans during spontaneous respiration. *Respiratory physiology & neurobiology 174*, 1 (2010), 111–118.

[121] Lee, J., and Devinsky, O. The role of autonomic dysfunction in sudden unexplained death in epilepsy patients. *Reviews in neurological diseases 2*, 2 (2004), 61–69.

[122] Lespinats, S., Verleysen, M., Giron, A., and Fertil, B. DD-HDS: A method for visualization and exploration of high-dimensional data. *IEEE Transactions on Neural Networks 18*, 5 (2007), 1265–1279.

[123] Leutmezer, F., Schernthaner, C., Lurger, S., Pötzelberger, K., and Baumgartner, C. Electrocardiographic changes at the onset of epileptic seizures. *Epilepsia 44*, 3 (2003), 348–354.

[124] Lizier, J. T. *The local information dynamics of distributed computation in complex systems*. Springer, 2012.

[125] Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. Local measures of information storage in complex distributed computation. *Information Sciences 208* (2012), 39–54.

[126] Lombardi, F. Clinical implications of present physiological understanding of HRV components. *Cardiac Electrophysiology Review 6*, 3 (2002), 245–249.

[127] Lucchini, M., Widjaja, D., Varon, C., Jansen, K., and Huffel, S. V. Phase-rectified signal averaging to evaluate ANS development in premature infants. In *Proceedings of the International Conference on Bio-inspired Systems and Signal Processing (BIOSIGNALS)* (2013), pp. 203–208.

[128] Luders, H. O., Najm, I., Nair, D., Widdess-Walsh, P., and Bingman, W. The epileptogenic zone: general principles. *Epileptic disorders 8* (2006), S1.

[129] Malik, M., Bigger, J. T., Camm, A. J., Kleiger, R. E., Malliani, A., Moss, A. J., and Schwartz, P. J. Heart rate variability: standards of measurement, physiological interpretation and clinical use. *European heart journal 17*, 3 (1996), 354–381.

[130] Malik, M., Farrell, T., and Camm, A. J. Circadian rhythm of heart rate variability after acute myocardial infarction and its influence on the prognostic value of heart rate variability. *The American journal of cardiology 66*, 15 (1990), 1049–1054.

[131] Malpas, S. C. Sympathetic nervous system overactivity and its role in the development of cardiovascular disease. *Physiological reviews 90*, 2 (2010), 513–557.

[132] Massin, M. M., Maeyns, K., Withofs, N., Ravet, F., and Gérard, P. Circadian rhythm of heart rate and heart rate variability. *Archives of disease in childhood 83*, 2 (2000), 179–182.

[133] McMurray, J. J., and Pfeffer, M. A. Heart failure. *The Lancet 365*, 9474 (2005), 1877–1889.

[134] McMurray, J. J., and Stewart, S. Epidemiology, aetiology, and prognosis of heart failure. *Heart 83*, 5 (2000), 596–602.

[135] Messenheimer, J. A., Quint, S. R., Tennison, M. B., and Keaney, P. Monitoring heart period variability changes during seizures. I. methods. *Journal of Epilepsy 3*, 1 (1990), 47–54.

[136] Meyer, C., and Keiser, H. Electrocardiogram baseline noise estimation and removal using cubic splines and state-space computation techniques. *Computers and Biomedical Research 10*, 5 (1977), 459–470.

[137] Mika, S., Schölkopf, B., Smola, A. J., Müller, K.-R., Scholz, M., and Rätsch, G. Kernel PCA and de-noising in feature spaces. In *NIPS* (1998), vol. 11, Citeseer, pp. 536–542.

[138] Mitchell, T. M. *Machine learning*. McGraw-Hill Science/Engineering/Math, 1997.

[139] Mitra, V., Wang, C.-J., and Banerjee, S. Text classification: A least square support vector machine approach. *Applied Soft Computing 7*, 3 (2007), 908–914.

[140] Moody, G. B., and Mark, R. G. QRS morphology representation and noise estimation using the karhunen-loeve transform. In *Computers in Cardiology 1989* (1989), IEEE, pp. 269–272.

[141] Moody, G. B., and Mark, R. G. The impact of the MIT-BIH arrhythmia database. *Engineering in Medicine and Biology Magazine 20*, 3 (2001), 45–50.

[142] Moody, G. B., Mark, R. G., Bump, M. A., Weinstein, J. S., Berman, A. D., Mietus, J. E., and Goldberger, A. L. Clinical validation of the ecg-derived respiration (EDR) technique. *Computers in cardiology 13* (1986), 507–510.

[143] Moody, G. B., Mark, R. G., Zoccola, A., and Mantero, S. Derivation of respiratory signals from multi-lead ECGs. *Computers in cardiology 12* (1985), 113–116.

[144] Müller, A., Morley-Davies, A., Barthel, P., Hnatkova, K., Bauer, A., Ulm, K., Malik, M., and Schmidt, G. Bivariate phase-rectified signal averaging for assessment of spontaneous baroreflex sensitivity: normalization of the results. *Journal of electrocardiology 45*, 1 (2012), 77–81.

[145] Nakamura, N., Koga, T., and Iseki, H. A meta-analysis of remote patient monitoring for chronic heart failure patients. *Journal of telemedicine and telecare 20*, 1 (2014), 11–17.

[146] Nei, M. Cardiac effects of seizures. *Epilepsy Currents 9*, 4 (2009), 91–95.

[147] of an American Academy of Sleep Medicine Task Force., T. R. Sleep-related breathing disorders in adults: recommendations for syndrome definition and measurement techniques in clinical research. *Sleep 22*, 5 (1999), 667–689.

[148] Okada, M. A digital filter for the QRS complex detection. *IEEE Transactions on Biomedical Engineering*, 12 (1979), 700–703.

[149] O'Regan, M. E., and Brown, J. K. Abnormalities in cardiac and respiratory function observed during seizures in childhood. *Developmental Medicine & Child Neurology 47*, 1 (2005), 4–9.

[150] Pallas-Areny, R., Colominas-Balague, J., and Rosell, F. The effect of respiration-induced heart movements on the ECG. *IEEE Transactions on Biomedical Engineering 36*, 6 (1989), 585–590.

[151] Pan, J., and Tompkins, W. J. A real-time QRS detection algorithm. *IEEE Transactions on Biomedical Engineering*, 3 (1985), 230–236.

[152] Pandya, U., and Desai, U. Bluetooth electrocardiogram: Its evaluation in a hospital environment. *Pulse, IEEE 4*, 3 (2013), 35–40.

[153] Parati, G., Mancia, G., Di Rienzo, M., and Castiglioni, P. Point: Counterpoint: Cardiovascular variability is/is not an index of autonomic control of circulation. *Journal of applied physiology 101*, 2 (2006), 676–682.

[154] Pattinson, K. T., Mitsis, G. D., Harvey, A. K., Jbabdi, S., Dirckx, S., Mayhew, S. D., Rogers, R., Tracey, I., and Wise, R. G. Determination of the human brainstem respiratory control network and its cortical connections in vivo using functional and structural imaging. *Neuroimage 44*, 2 (2009), 295–305.

[155] Pellock, J. M., Hrachovy, R., Shinnar, S., Baram, T. Z., Bettis, D., Dlugos, D. J., Gaillard, W. D., Gibson, P. A., Holmes, G. L., Nordli, D. R., et al. Infantile spasms: a US consensus report. *Epilepsia 51*, 10 (2010), 2175–2189.

[156] Penzel, T., McNames, J., De Chazal, P., Raymond, B., Murray, A., and Moody, G. Systematic comparison of different algorithms for apnoea detection based on electrocardiogram recordings. *Medical and Biological Engineering and Computing 40*, 4 (2002), 402–407.

[157] Penzel, T., Moody, G. B., Mark, R. G., Goldberger, A. L., and Peter, J.-H. H. The apnea-ECG database. In *Computers in Cardiology 2000* (2000), IEEE, pp. 255–258.

[158] Penzel, T., Wessel, N., Riedl, M., Kantelhardt, J. W., Rostig, S., Glos, M., Suhrbier, A., Malberg, H., and Fietze, I. Cardiovascular and respiratory dynamics during normal and pathological sleep. *Chaos: An Interdisciplinary Journal of Nonlinear Science 17*, 1 (2007), 015116.

[159] Peppard, P. E., Young, T., Barnet, J. H., Palta, M., Hagen, E. W., and Hla, K. M. Increased prevalence of sleep-disordered breathing in adults. *American journal of epidemiology 177*, 9 (2013), 1006–1014.

[160] Pinciroli, F., Corso, M., Fuggetta, A., Masseroli, M., Bonacina, S., and Marceglia, S. Telemedicine and e-health. *Pulse, IEEE 2*, 3 (2011), 62–70.

[161] Piotrowskia, Z., and Rózanowski, K. Robust algorithm for heart rate (HR) detection and heart rate variability (HRV) estimation. *Acta Physica Polonica A 118*, 1 (2010), 131–135.

[162] Rangayyan, R. M. *Biomedical Signal Analysis: A Case-Study Approach*. Wiley-IEEE Press, 2002.

[163] Rissanen, J. Mdl denoising. *IEEE Transactions on Information Theory 46*, 7 (2000), 2537–2543.

[164] Rosenbaum, D. S., Jackson, L. E., Smith, J. M., Garan, H., Ruskin, J. N., and Cohen, R. J. Electrical alternans and vulnerability to ventricular arrhythmias. *New England Journal of Medicine 330*, 4 (1994), 235–241.

[165] Sameni, R., Shamsollahi, M. B., Jutten, C., and Clifford, G. D. A nonlinear bayesian filtering framework for ECG denoising. *IEEE Transactions on Biomedical Engineering 54*, 12 (2007), 2172–2185.

[166] Sameni, R., Shamsollahi, M. B., Jutten, C., and Clifford, G. D. A nonlinear bayesian filtering framework for ECG denoising. *IEEE Transactions on Biomedical Engineering 54*, 12 (2007), 2172–2185.

[167] Sathyaprabha, T., Satishchandra, P., Netravathi, K., Sinha, S., Thennarasu, K., and Raju, T. Cardiac autonomic dysfunctions in chronic refractory epilepsy. *Epilepsy research 72*, 1 (2006), 49–56.

[168] Scargle, J. D. Studies in astronomical time series analysis. II-statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal 263* (1982), 835–853.

[169] Schölkopf, B., Smola, A., and Müller, K.-R. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation 10*, 5 (1998), 1299–1319.

[170] Scholkopf, B., and Smola, A. J. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.

[171] Schreiber, T. Measuring information transfer. *Physical review letters 85*, 2 (2000), 461.

[172] Schreiber, T. Measuring information transfer. *Physical review letters 85*, 2 (2000), 461–464.

[173] Schumann, A. Y., Kantelhardt, J. W., Bauer, A., and Schmidt, G. Bivariate phase-rectified signal averaging. *Physica A: Statistical Mechanics and its Applications 387*, 21 (2008), 5091–5100.

[174] Seber, G. A. *Multivariate observations*, vol. 252. John Wiley & Sons, 2009.

[175] Sevcencu, C., and Struijk, J. J. Autonomic alterations and cardiac changes in epilepsy. *Epilepsia 51*, 5 (2010), 725–737.

[176] Seyal, M., Bateman, L. M., Albertson, T. E., Lin, T.-C., and Li, C.-S. Respiratory changes with seizures in localization-related epilepsy: Analysis of periictal hypercapnia and airflow patterns. *Epilepsia 51*, 8 (2010), 1359–1364.

[177] Sharma, L., Dandapat, S., and Mahanta, A. Kurtosis-based noise estimation and multiscale energy to denoise ECG signal. *Signal, Image and Video Processing 7*, 2 (2013), 235–245.

[178] Sharmila, K., Krishna, E. H., Reddy, K. N., and Reddy, K. A. Application of multiscale principal component analysis (MSPCA) for enhancement of ecg signals. In *IEEE Instrumentation and Measurement Technology Conference (I2MTC)* (2011), IEEE, pp. 1–5.

[179] Shields, W. D. West's syndrome. *Journal of child neurology 17*, 1 suppl (2002), S76–S79.

[180] Shyu, L.-Y., Wu, Y.-H., and Hu, W. Using wavelet transform and fuzzy neural network for VPC detection from the holter ECG. *IEEE Transactions on Biomedical Engineering 51*, 7 (2004), 1269–1273.

[181] Silverman, B. W. *Density estimation for statistics and data analysis*, vol. 26. CRC press, 1986.

[182] Singh, K., Katz, E. S., Zarowski, M., Loddenkemper, T., Llewellyn, N., Manganaro, S., Gregas, M., Pavlova, M., and Kothare, S. V. Cardiopulmonary complications during pediatric seizures: a prelude to understanding SUDEP. *Epilepsia 54*, 6 (2013), 1083–1091.

[183] Sinha, K. Brain stem infarction: Clinical clues to localize them. *Journal, Indian Academy of Clinical Medicine 1*, 3 (2000), 213–21.

[184] Stolojescu, C. Estimation of noise in ECG signals using wavelets. In *11th International Conference on Optimization of Electrical and Electronic Equipment,(OPTIM)* (2008), IEEE, pp. 113–118.

[185] Surges, R., Scott, C. A., and Walker, M. C. Enhanced QT shortening and persistent tachycardia after generalized seizures. *Neurology 74*, 5 (2010), 421–426.

[186] Suykens, J. A. K., Alzate, C., Pelckmans, K., et al. Primal and dual model representations in kernel-based learning. *Statistics Surveys 4* (2010), 148–183.

[187] Suykens, J. A. K., Van Gestel, T., De Brabanter, J., De Moor, B., and Vandewalle, J. *Least squares support vector machines*. World Scientific, 2002.

[188] Suykens, J. A. K., Van Gestel, T., Vandewalle, J., and De Moor, B. A support vector machine formulation to PCA analysis and its kernel version. *IEEE Transactions on Neural Networks 14*, 2 (2003), 447–450.

[189] Suykens, J. A. K., and Vandewalle, J. Least squares support vector machine classifiers. *Neural processing letters 9*, 3 (1999), 293–300.

[190] Syeda-Mahmood, T., Beymer, D., and Wang, F. Shape-based matching of ECG recordings. In *29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (2007), IEEE, pp. 2012–2018.

[191] Szilágyi, L. Wavelet-transform-based QRS complex detection in on-line holter systems. In *First Joint BMES/EMBS Conference* (1999), vol. 1, IEEE, pp. 271–vol.

[192] Tang, J.-t., Yang, X.-l., Xu, J.-c., Tang, Y., Zou, Q., and Zhang, X.-k. The algorithm of R peak detection in ECG based on empirical mode decomposition. In *Fourth International Conference on Natural Computation, ICNC'08* (2008), vol. 5, IEEE, pp. 624–627.

[193] Thomas, R. J., Mietus, J. E., Peng, C.-K., Gilmartin, G., Daly, R. W., Goldberger, A. L., and Gottlieb, D. J. Differentiating obstructive from central and complex sleep apnea using an automated electrocardiogram-based method. *Sleep 30*, 12 (2007), 1756.

[194] Thomas, R. J., Mietus, J. E., Peng, C.-K., and Goldberger, A. L. An electrocardiogram-based technique to assess cardiopulmonary coupling during sleep. *Sleep 28*, 9 (2005), 1151–1161.

[195] Throne, R. D. Detecting ventricular fibrillation using efficient techniques for computing a normalized autocorrelation. *Computers in biology and medicine 23*, 4 (1993), 317–325.

[196] Tipping, M. E. Sparse kernel principal component analysis. In *Advances in Neural Information Processing Systems 13*, T. Leen, T. Dietterich, and V. Tresp, Eds. MIT Press, 2001, pp. 633–639.

[197] Tripathy, R. K., Acharya, A., and Choudhary, S. K. Gender classification from ECG signal analysis using least square support vector machine. *American Journal of Signal Processing 2*, 5 (2012), 145–149.

[198] Tzeng, Y., Sin, P. Y., and Galletly, D. C. Human sinus arrhythmia: inconsistencies of a teleological hypothesis. *American Journal of Physiology-Heart and Circulatory Physiology 296*, 1 (2009), H65–H70.

[199] Van de Vel, A., Cuppens, K., Bonroy, B., Milosevic, M., Jansen, K., Van Huffel, S., Vanrumste, B., Lagae, L., and Ceulemans, B. Non-EEG seizure-detection systems and potential SUDEP prevention: state of the art. *Seizure 22*, 5 (2013), 345–355.

[200] Van Elmpt, W. J., Nijsen, T. M., Griep, P. A., and Arends, J. B. A model of heart rate changes to detect seizures in severe epilepsy. *Seizure 15*, 6 (2006), 366–375.

[201] Vanninen, E., Tuunainen, A., Kansanen, M., Uusitupa, M., and Länsimies, E. Cardiac sympathovagal balance during sleep apnea episodes. *Clinical Physiology 16*, 3 (1996), 209–216.

[202] Vapnik, V. Pattern recognition using generalized portrait method. *Automation and remote control 24* (1963), 774–780.

[203] Vapnik, V. *Statistical learning theory*, vol. 2. Wiley New York, 1998.

[204] Varon, C., Alzate, C., and Suykens, J. A. K. Noise level estimation for model selection in kernel PCA denoising. *IEEE Transactions on Neural Networks and Learning Systems in press* (2015).

[205] Varon, C., Caicedo, A., Jansen, K., Lagae, L., and Van Huffel, S. Detection of epileptic seizures from single lead ECG by means of phase rectified signal averaging. In *36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (Aug 2014), pp. 3789–3790.

[206] Varon, C., Caicedo, A., Testelmans, D., Buyse, B., and Van Huffel, S. A novel algorithm for the automatic detection of sleep apnea from single-lead ECG. *IEEE Transactions on Biomedical Engineering in press* (2015).

[207] Varon, C., Jansen, K., Lagae, L., Faes, L., and Van Huffel, S. Transient behavior of cardiorespiratory interactions towards the onset of epileptic seizures.

[208] Varon, C., Jansen, K., Lagae, L., and Van Huffel, S. Cardiorespiratory analysis on children suffering from absence and complex partial seizures. In *Computing in Cardiology Conference (CinC), 2012* (2012), IEEE, pp. 113–116.

[209] Varon, C., Jansen, K., Lagae, L., and Van Huffel, S. Detection of epileptic seizures by means of morphological changes in the ECG. In *Computing in Cardiology Conference (CinC), 2013* (2013), IEEE, pp. 863–866.

[210] Varon, C., Montalto, A., Jansen, K., Lagae, L., Marinazzo, D., Faes, L., and Van Huffel, S. Interictal cardiorespiratory variability in temporal lobe and absence epilepsy in childhood. In *Proceedings of the 8th Conference of the European Study Group on Cardiovascular Osciliations* (2014).

[211] Varon, C., Montalto, A., Jansen, K., Lagae, L., Marinazzo, D., Faes, L., and Van Huffel, S. Interictal cardiorespiratory variability in temporal lobe and absence epilepsy in childhood. *Physiological measurement 36*, 4 (2015), 845–856.

[212] Varon, C., Morenikeji, A., Minter, J., Stapleton, M., Thomson, S., Jaecques, S., Brunner-La Rocca, H.-P., and Van Huffel, S. Telehealth on heart failure: results of the recap project. *Journal of Telemedicine and Telecare in press* (2015).

[213] Varon, C., Testelmans, D., Buyse, B., Suykens, J. A. K., and Van Huffel, S. Robust artefact detection in long-term ECG recordings based on autocorrelation function similarity and percentile analysis. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2012), IEEE, pp. 3151–3154.

[214] Varon, C., Testelmans, D., Buyse, B., Suykens, J. A. K., and Van Huffel, S. Sleep apnea classification using least-squares support vector machines on single lead ECG. In *Proceedings of the 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2013), IEEE, pp. 5029–5032.

[215] Watanabe, K., Negoro, T., Aso, K., and Matsumoto, A. Reappraisal of interictal electroencephalograms in infantile spasms. *Epilepsia 34*, 4 (1993), 679–685.

[216] Widjaja, D., Caicedo, A., Vlemincx, E., Van Diest, I., and Van Huffel, S. Separation of respiratory influences from the tachogram: a methodological evaluation. *PloS one 9*, 7 (2014), e101713.

[217] Widjaja, D., Varon, C., Dorado, A. C., Suykens, J. A. K., and Van Huffel, S. Application of kernel principal component analysis for single-lead-ECG-derived respiration. *IEEE Transactions on Biomedical Engineering 59*, 4 (2012), 1169–1176.

[218] Williams, C., and Seeger, M. Using the Nyström method to speed up kernel machines. In *Proceedings of the 14th Annual Conference on Neural Information Processing Systems* (2001), no. EPFL-CONF-161322, pp. 682–688.

[219] Wolf, M., Varigos, G., Hunt, D., and Sloman, J. Sinus arrhythmia in acute myocardial infarction. *The Medical Journal of Australia 2*, 2 (1978), 52–53.

[220] Yasuma, F., and Hayano, J.-i. Respiratory sinus arrhythmia: Why does the heartbeat synchronize with respiratory rhythm? *Chest Journal 125*, 2 (2004), 683–690.

[221] Young, T., Finn, L., Peppard, P. E., Szklo-Coxe, M., Austin, D., Nieto, F. J., Stubbs, R., and Hla, K. M. Sleep disordered breathing and mortality: eighteen-year follow-up of the Wisconsin sleep cohort. *Sleep 31*, 8 (2008), 1071.

[222] Zhang, J., Goode, K. M., Cuddihy, P. E., and Cleland, J. G. Predicting hospitalization due to worsening heart failure using daily weight measurement: analysis of the trans-european network-home-care management system (TEN-HMS) study. *European journal of heart failure 11*, 4 (2009), 420–427.

[223] Zigmond, A. S., and Snaith, R. P. The hospital anxiety and depression scale. *Acta psychiatrica scandinavica 67*, 6 (1983), 361–370.

[224] Zijlmans, M., Flanagan, D., and Gotman, J. Heart rate changes and ECG abnormalities during epileptic seizures: prevalence and definition of an objective clinical sign. *Epilepsia 43*, 8 (2002), 847–854.

Curriculum vitae

Carolina Varon was born in Ibagué, Colombia. She received the B.Sc. degree in electronic engineering from the Universidad de Ibagué, Ibagué, Colombia, in 2005. From 2002 to 2003, she was a Site Engineer in Agroindustrial del Tolima, Ibagué, Colombia, and in 2005 she joined Security Solutions, Bogotá, Colombia, where she was the Technical Support for Latin America until 2007.

In 2007 she moved to Leuven, Belgium, where two years later she received the M.Sc. degree in astronomy and astrophysics and in 2010 the M.Sc. degree magna cum laude in artificial intelligence, both from the Katholieke Universiteit Leuven. Since 2011 she has been working toward the Ph.D. degree in the Department of Electrical Engineering (ESAT-STADIUS) of the KU Leuven under the supervision of Prof. Sabine Van Huffel and Prof. Johan Suykens.

List of publications

Papers in international journals

1. **Varon C.**, Caicedo A., Testelmans D., Buyse B., Van Huffel S. A novel algorithm for the automatic detection of sleep apnea from single-lead ECG. *IEEE Transactions on Biomedical Engineering* 2015, in press.

2. **Varon, C.**, Alzate C., Suykens J.A.K. Noise Level Estimation for Model Selection in Kernel PCA Denoising. *IEEE Transactions on Neural Networks and Learning Systems* 2015, in press.

3. **Varon, C.**, Alao M., Minter, J., Stapleton, M., Thomson, S., Jaecques, S., Brunner-La Rocca, H., Van Huffel, S. Telehealth on heart failure: results of the Recap project. *Journal of telemedicine and telecare* 2015, in press.

4. **Varon, C.**, Montalto A., Jansen K., Lagae L., Marinazzo D., Faes L., and Van Huffel S. Interictal Cardiorespiratory Variability in Temporal Lobe and Absence Epilepsy in Childhood. *Physiological Measurement*, vol. 36, no. 4, pp. 845-856, 2015. *(Invited contribution).*

5. Caicedo A., **Varon C.**, Thewissen L., Naulaers G., Lemmers P., van Bel F., van Huffel S. Influence of the Maternal use of Labetalol on the Neurogenic Mechanism for Cerebral Autoregulation Assessed by Means of NIRS. *Advances in Experimental Medicine and Biology*, vol. 812, pp. 173-179, 2014.

6. Jansen K., **Varon C.**, Van Huffel S., Lagae L. Peri-ictal ECG changes in childhood epilepsy: implications for detection systems. *Epilepsy and Behavior*, vol. 29, pp. 72-76, 2013.

7. Jansen K., **Varon C.**, Van Huffel S., Lagae L. Ictal and interictal respiratory changes in temporal lobe and absence epilepsy in childhood. *Epilepsy research*, vol. 106, no. 3, pp. 410-416, 2013.

8. Jansen K., **Varon C.**, Van Huffel S., Lagae L. Early respiratory dysfunction as a biomarker for epileptic encephalopathy. *Acta Neurologica Scandinavica*, vol. 128, no. 6, pp. 381-385, 2013.

9. Widjaja D., **Varon C.**, Caicedo A., Suykens J.A.K., Van Huffel S. Application of Kernel Principal Component Analysis for Single Lead ECG-Derived Respiration. *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 4, pp. 1169-1176, 2012.

10. **Varon C.**, Alzate C., Suykens J.A.K., Debosscher J. Kernel Spectral Clustering of Time Series in the CoRoT Exoplanet Database. *Astronomy & Astrophysics*, vol. 531, no. A156, pp. 1-15, 2011.

Papers submitted in international journals

11. **Varon, C.**, Jansen K., Lagae L., and Van Huffel S. Detection of epileptic seizures from single lead ECG. *Journal of Electrocardiology*, submitted in April 2015. *(Invited contribution).*

12. Willemen T., **Varon C.**, Caicedo A., Haex B., Vander Sloten J., Van Huffel, S. Probabilistic cardiac and respiratory based classification of sleep and apneic events in subjects with sleep apnea. *Physiological measurement*, submitted in January 2015.

13. De Cooman T., Goovaerts G., **Varon C.**, Widjaja D., Willemen T., Van Huffel S. Heart Beat Detection in Multimodal Data Using Automatic Signal Type Recognition. *Physiological measurement*, submitted in February 2015. *(Invited contribution)*

Papers in proceedings of international conferences

14. **Varon C.**, Jansen K., Lagae L., Faes L., Van Huffel S. Transient behavior of cardiorespiratory interactions towards the onset of epileptic seizures. *In proceedings of the 41st Annual Computing in Cardiology (CinC)*, Cambridge MA, USA, Sep. 2014.

15. **Varon C.**, Minter J., Stapleton M., Thomson S., Jaecques S., Brunner-La Rocca H-P., Van Huffel S. Effect of Telehealth on Self-Care Behavior of Heart Failure Patients. *In proceedings of the 41st Annual Computing in Cardiology (CinC)*, Cambridge MA, USA, Sep. 2014.

16. **Varon C.**, Caicedo A., Jansen K., Lagae L., Van Huffel S. Detection of epileptic seizures from single lead ECG by means of phase rectified signal averaging. *In proceedings of the 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Chicago IL, USA, Aug. 2014.

17. Caicedo A., **Varon C.**, Alderliesten T., Lemmers P., Van Bel F., Naulaers G., Van Huffel S. Differences in the Cerebral Hemodynamic Regulation Mechanisms of Premature Infants with Intra-Ventricular Hemorrhage Assessed by Means of Phase Rectified Signal Averaging. *In proceedings of the 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Chicago IL, USA, Aug. 2014.

18. Caicedo A., **Varon C.**, Van Huffel S. Analysis of non-linear respiratory influences on sleep apnea classification. *In proceedings of the 41st Annual Computing in Cardiology (CinC)*, Cambridge MA, USA, Sep. 2014.

19. Widjaja D., **Varon C.**, Testelmans D., Buyse B., Faes L., Van Huffel S. Separating respiratory influences from the tachogram: methods and their sensitivity to the type of respiratory signal. *In proceedings of the 41st Annual Computing in Cardiology (CinC)*, Cambridge MA, USA, Sep. 2014.

20. Goovaerts G., **Varon C.**, Vandenberk B., Willems R., Van Huffel S. Tensor-based Detection of T Wave Alternans in Multilead ECG Signals. *In proceedings of the 41st Annual Computing in Cardiology (CinC)*, Cambridge MA, USA, Sep. 2014.

21. De Cooman T., Goovaerts G., **Varon C.**, Widjaja D., Van Huffel S. Robust detection of heart beats in multimodal data using automatic signal type recognition and heart beat location estimation. *In proceedings of the 41st Annual Computing in Cardiology (CinC)*, Cambridge MA, USA, Sep. 2014. *(Winner of phase II and third place in phase II of the Physionet/CinC challenge 2014.)*

22. Willemen T., **Varon C.**, Haex B., Vander Sloten J., Van Huffel S. Assessment of Different Methodologies to Include Temporal Information in Classifying Episodes of Sleep Apnea Based on Single-Lead Electrocardiogram. *In proceedings of the 41st Annual Computing in Cardiology (CinC)*, Cambridge MA, USA, Sep. 2014.

23. **Varon C.**, Montalto A., Jansen K., Lagae L., Marinazzo D., Faes L., Van Huffel S. Interictal cardiorespiratory variability in temporal lobe and absence epilepsy in childhood. *In proceedings of the 8th Conference of the European Study Group on Cardiovascular Oscillations (ESGCO)*, Trento, Italy, Apr. 2014.

24. **Varon C.**, Minter J., Stapleton M., Thomson S., Jaecques S., Sanders-van Wijk S., Brunner-La Rocca H-P., Van Huffel S. Telehealth for patients with heart failure: A comparison between telehealth systems. *In proceedings of the International eHealth, Telemedicine and health ICT Forum (Med-e-Tel)*, Luxembourg, Apr. 2014.

25. **Varon C.**, Jansen K., Lagae L., Van Huffel S. Detection of epileptic seizures by means of morphological changes in the ECG. *In proceedings of the 40th Annual Computing in Cardiology (CinC)*, Zaragoza, Spain, Sep. 2013.

26. **Varon C.**, Testelmans D., Buyse B., Suykens J.A.K., Van Huffel S. Sleep apnea classiĭnĄcation using least-squares support vector machines on single lead ECG. *In proceedings of the 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Osaka, Japan, Jul. 2013.

27. Lucchini M., Widjaja D., **Varon C.**, Jansen K., Van Huffel S. Phase-Rectified Signal Averaging to Evaluate ANS Development in Premature Infants. *In proceedings of the International Conference on Bio-inspired Systems and Signal Processing (BIOSIGNALS)*, Barcelona, Spain, Feb. 2013.

28. **Varon C.**, Testelmans D., Buyse B., Suykens J.A.K., Van Huffel S. Robust artefact detection in long-term ECG recordings based on autocorrelation function similarity and percentile analysis. *In proceedings of the 34th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, San Diego CA, USA, Aug. 2012.

29. **Varon C.**, Jansen K., Lagae L., Van Huffel S. Cardiorespiratory Analysis on Children Suffering from Absence and Complex Partial Seizures. *In proceedings of the 39th Annual Computing in Cardiology (CinC)*, Krakow, Poland, Sep. 2012.

30. Joosen P., Aerts W., **Varon C.**, Widjaja D., Vandeput S., Aubert A., Van Huffel S. Effect of posture on the cardiorespiratory system using canonical correlation analysis. *In proceedings of the 39th Annual Computing in Cardiology (CinC)*, Krakow, Poland, Sep. 2012.

31. Aerts W., Joosen P., Widjaja D., **Varon C.**, Vandeput S., Van Huffel S., Aubert A. Heart Rate and blood pressure variability under moon, mars and zero gravity conditions during parabolic flights. *In proceedings of the Joint Life Science Meeting "Life in Space for Life on Earth"*, Aberdeen, United Kingdom, Jun. 2012.

32. Widjaja D., **Varon C.**, Caicedo A., Van Huffel S. An improved ECG-Derived Respiration Method using Kernel Principal Component Analysis. *In proceedings of the 38th Annual Computing in Cardiology (CinC)*, Hangzhou, China, Sep. 2011.

Abstracts in proceedings of (inter)national conferences

33. **Varon C.**, Alao M., Minter J., Stapleton M., Thomson S., Jaecques S., Brunner-La Rocca H-P., Van Huffel, S. Telehealth on heart failure: Preliminary results of the Recap project. *In proceedings of the 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Chicago IL, USA, Aug. 2014.

34. **Varon C.**, Alao M., Minter J., Stapleton M., Thomson S., Jaecques S., Brunner-La Rocca H-P., Van Huffel, S. Telehealth on heart failure: Preliminary results of the Recap project. *In proceedings of the 13th Belgian day on biomedical engineering*, Brussels, Belgium, Nov. 2014.

35. Jansen, K., **Varon C.**, Van Huffel S., Lagae L. Morphological ECG changes in seizures: implications for seizure detection. *In proceedings of the European Paediatric Neurology Society (EPNS)*, Brussels, Belgium, Sep. 2013.

36. Jansen, K., **Varon C.**, Van Huffel S., Lagae L. Cardiac and respiratory autonomic dysfunction in childhood epilepsy. *In proceedings of the European Paediatric Neurology Society (EPNS)*, Brussels, Belgium, Sep. 2013.

37. **Varon C.**, Jansen K., Lagae L., Van Huffel S. Effect of interictal epileptic activity on the cardio-respiratory control. *In proceedings of the 11th Belgian day on biomedical engineering*, Brussels, Belgium, Dec. 2012.

38. Caicedo A., **Varon C.**, Van Huffel S. Orthogonal subspace projections in sleep apnea detection. *In proceedings of the 11th Belgian day on biomedical engineering*, Brussels, Belgium, Dec. 2012.

39. **Varon C.**, Jansen K., Lagae L., Van Huffel S. Long-term effects of absence and complex-partial seizures on cardiorespiratory dynamics. *In abstract book of the 40th ISOTT annual meeting*, Bruges, Belgium, Aug. 2012.

40. **Varon C.**, Jansen K., Lagae L., Van Huffel S. Respiratory pattern in infants with West symdrome. *In proceedings of the 5th International Conference of Advances in Medical Signal and Information Processing (MEDSIP)* , Liverpool, United Kingdom, Jul. 2012.

41. Caicedo A., **Varon C.** and Van Huffel S. Cardiopulmonary Coupling in Sleep Apnea Assessed by Means of Orthogonal Subspace Projection. *In proceedings of the 5th International Conference of Advances in Medical Signal and Information Processing (MEDSIP)* , Liverpool, United Kingdom, Jul. 2012.

42. **Varon C.**, Vandeput S., Suykens J.A.K., Van Huffel S. Sleep apnea detection using an LS-SVM classifier on single lead ECG. *In proceedings of the 10th Belgian day on biomedical engineering*, Brussels, Belgium, Dec. 2011.



FACULTY OF ENGINEERING SCIENCE
DEPARTMENT OF ELECTRICAL ENGINEERING (ESAT)
STADIUS
Kasteelpark Arenberg 10 box 2446
B-3001 Leuven
carolina.varon@esat.kuleuven.be