KU LEUVEN

ARENBERG DOCTORAL SCHOOL
Faculty of Engineering Science

Epileptic Seizure Detection in a Home Environment

Thomas De Cooman

Supervisors:
Prof. dr. ir. S. Van Huffel
Prof. dr. L. Lagae

Dissertation presented in partial fulfillment of the requirements for the degree of Doctor of Engineering Science (PhD): Electrical Engineering

December 2018

Epileptic Seizure Detection in a Home Environment

Thomas DE COOMAN

Examination committee:
Prof. dr. ir. A. Bultheel, chair
Prof. dr. ir. S. Van Huffel, supervisor
Prof. dr. L. Lagae, supervisor
Prof. dr. ir. B. Vanrumste
Prof. dr. ir. C. Van Hoof
Prof. dr. J. Arends
(University of Technology Eindhoven)
Prof. dr. B. Ceulemans
(University of Antwerpen)
Prof. dr. ir. H. Sørensen
(Technical University of Denmark)

Dissertation presented in partial fulfillment of the requirements for the degree of Doctor of Engineering Science (PhD): Electrical Engineering

December 2018

© 2018 KU Leuven – Faculty of Engineering Science
Uitgegeven in eigen beheer, Thomas De Cooman, 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, fotokopie, 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, electronic or any other means without written permission from the publisher.

Acknowledgments

During my master thesis, I realized how much I loved the art of constructing algorithms. Even reading about simple sorting algorithms triggered some dopamine in my brains. Unfortunately, my master in Computer Science at the UGent did not result in a sufficient amount of algorithmic dopamine. Therefore I wanted to go for a master thesis topic in which the focus would go into the development of a new algorithm, preferably on a topic close to my interests and/or with an "impact on society" as millennials would say. The choices were limited back then, but I'm grateful to have been able to do master thesis on epileptic seizure detection using heart rate. This choice back then was the start of a journey into biomedical signal processing and machine learning, resulting in this PhD book.

First of all, my thanks to Prof. Sabine Van Huffel for allowing me to do this PhD thesis on the subject of my master thesis, even when the initial project for which I was assumed to be paid was not accepted. By giving me sufficient freedom in my research and giving mostly positive constructive feedback on my progression I was able to feel myself comfortable enough to do the research I wanted to do within the scope of my topic. Your efforts to try and make everyone feel at ease at work are really appreciated. I hope soon one of the BioMed algorithms will lead to real-life usage in order to improve the quality of lives of patients.

This thesis would not be possible without help from the different clinical partners, not only for their datasets, but also for their guidance from a clinical point of view. I'd like to thank all clinical collaborators from UZ Leuven (Prof. Lieven Lagae, Prof. Wim Van Paesschen, Prof. Katrien Jansen and Evy Cleeren) and UZ Antwerp (Prof. Berten Ceulemans and Anouk Van de Vel). You allowed me to understand the clinical part of my research, as well to understand the needed requirements for the patients. Anouk, as most of my clinical questions were referred to you, special thanks to you for refreshing the clinical part of my research once every couple of weeks/months. I hope we can still exchange

e-mails trying to keep each other up to date on what is happening at the other research side of epileptic seizure detection .

Also thank you Prof. Helge Sørensen for allowing me to do a research visit at DTU in Denmark. A lot of my ideas of this thesis came up during my stay in Denmark, so it had a greater impact than just those three months in the peaceful and green Denmark. Also thank you Prof. Troels Kjær for your clinical guidance throughout my visit and the nice dinner we had at your house. Also thanks to Markus and Mehdi for making my visit in Denmark socially nicer as well on a daily basis.

During the last year of my PhD, I was able to work with probably the largest available dataset for seizure detection using wearable signals. Thank you prof. Johan Arends for introducing me to the Dutch Tele-Epilepsy consortium and for believing in the fact that my personalized research approach can be of added value in existing seizure detection devices. This means a lot to me knowing all your experience on this topic. I hope that after your retirement, the collaboration with the consortium can continue in order to try and improve the quality of life of epilepsy patients.

Thank you Prof. Bart Vanrumste and Prof. Chris Van Hoof for your much appreciated feedback as jury members during my preliminary defense. Also thank you to the chair Prof. Adhemar Bultheel for hosting both defences in your relaxing way.

I would also like to thank the different partners from the SeizeIT consortium for the currently ongoing and future collaboration. You allowed to further validate my previously developed algorithms on new data, which is a crucial part of research. I hope it will lead to a device that will improve the quality of life of patients with more subtle seizures, which were mostly ignored by industry up until now.

As I mentioned above, this all started with my master thesis on epileptic seizure detection using the heart rate at the UGent. Thank you Pieter van Mierlo and Evelien Carrette for guiding me back then and making me hooked on the topic. It inspired me to continue researching this topic as you can read further in this book.

Of course, massive thanks to all the BioMed people. I know I might be a bit of a shy guy at first. But after a while, I felt comfortable enough here to express myself like when being around friends and family (perhaps a bit too comfortable at some times for some of you). I will not name all of you to avoid forgetting someone (I know the pain), but thanks to each individual currently in BioMed and the ones who already left since I started some years ago. Due to the home-work distance and busy lifestyle/kids, I was not able to get to know

all of you as well as I wanted, and some of you I know better than I wanted;) Sorry for bothering you lately with GDPR and IP-related stuff, I hope you still like me after all my assignments to you. Kirsten, Griet, Laura and Vanya, thank you for making me feel comfortable in the office during the first years of my PhD and guiding me on the right track. Alexander and Carolina (and later on Neetha), thank you for getting me off the right track with silly stuff whenever it was needed in the new office. Griet, thank you for sharing all the laughter/pains/practicalities throughout the entire PhD and the large amount of coffee breaks. Caro, thank you for the guidance (both work and offwork) and your everlasting kindness. I really hope you will be able to reach in the near future what you are aiming for. Sorry for making so many bad jokes as my sign of appreciation for this to you, but I cannot promise I will stop doing them in the future ;)

Bedankt mama en papa, om me ondanks mijn liever-lui-dan-moe jaren en jaren van zelftwijfel toch verder te laten studeren waar ik aan begonnen ben. Het besef van de nood aan verandering kwam net op tijd, en ik denk dat het schrijven van deze doctoraatsthesis jullie vermoeden bevestigt dat er wel degelijk wat in me zat.

Eindigen dient altijd gedaan te worden met de belangrijskte personen. Nele, bedankt om me tijdens al deze jaren te steunen in goede en kwade tijden. We hebben tijdens de afgelopen jaren veel beleefd op persoonlijk en zakelijk vlak, en dit alles was niet mogelijk zonder jouw steun en liefde. Nona en Cosette, bedankt om me er voldoende aan te doen herinneren dat, ondanks alle inspanningen, mijn thesis toch maar gewoon mijn "werk" is, en dat er nog altijd belangrijkere dingen in het leven zijn. Ik betwijfel dat jullie dit ooit lezen, maar indien jullie dit toch doen, laat het me maar weten, dan toveren jullie allicht een zoveelste glimlach op mijn gezicht :)

Thomas
Leuven, December 2018

Abstract

The quality of life of refractory epilepsy patients is strongly affected by the unpredictable occurrences of epileptic seizures. It can be improved by home monitoring solutions, specially automated seizure warning systems. These systems automatically detect ongoing seizures and warn the patient's relatives or caregivers about this event. The required algorithms for automated real-time detection of seizures are the focus of this thesis. The goal of this thesis is to try and improve the performance of these algorithms by using different approaches, in which the electrocardiography (ECG) is used as an example in most cases.

ECG-based seizure detection which is currently not thoroughly investigated in the literature. As a first solution, a new patient-independent algorithm for real-time seizure detection is proposed in this thesis. It uses a new preprocessing procedure that detects strong heart rate increases caused by sympathetic activations in the autonomic nervous system. The features that are extracted using this procedure are then classified by a support vector machine classifier. This algorithm showed to outperform algorithms from the literature by causing 50% less false alarms.

The automated detection of epileptic seizure is however a very patient-specific problem, and thus requires personalization in order to achieve an optimal performance. The problem of efficient and fast personalization of these algorithms is however currently not addressed in the literature. Therefore, different solutions for personalizing ECG-based seizure detection using a limited amount of patient data are proposed in this thesis. Each proposed personalization procedure has its specific target use case in practice, taking into account potential real-life issues. First, a transfer learning approach is used to personalize the initially proposed patient-independent seizure detector using a limited amount of annotated patient data. It can easily adapt during an offline training procedure, resulting in a robust personalized approach when a limited amount of seizures are acquired. This approach however requires accurately annotated data, which is only available in a hospital setting. Therefore, another

personalization approach is proposed, which in real-time heuristically adapts to the patient's heart rate characteristics by using annotations made by the patients themselves. As these annotations can sometimes be incorrect, an additional procedure was integrated into the algorithm to deal with potential incorrect patient feedback. The proposed method showed also to quickly adapt to the patient, and is robust against large fractions of incorrect patient feedback.

Also unsupervised personalization approaches are proposed in this thesis for nocturnal seizure detection. Unsupervised approaches do not require seizure annotations, and are thus of added value to avoid missing accurate seizure annotations or even missing seizure occurrences in patients. This is specially interesting for patients with a low seizure frequency, as the supervised approaches require the occurrence of a couple of seizures in a patient. Two different unsupervised solutions are proposed in this thesis. A first algorithm includes a real-time nocturnal adaptation procedure, which already starts to personalize the algorithm within a couple of hours of monitoring. The second algorithm uses a newly proposed unsupervised transfer learning approach, which offline adapts more accurately to the patient's specific characteristics.

These personalized algorithms led to decreases in false alarm rate varying between 30-57% compared to the initially proposed patient-independent approach. Automatized personalization is thus advised in ECG-based seizure detection, as it can be implemented without much extra effort from the patients/clinicians, requiring only a limited amount of (annotated) patient data. The proposed ECG-based algorithms are specially useful as part of multimodal algorithms for the detection of a broad range of seizure types.

In the last part of the thesis, different modalities including ECG, electromyography and accelerometry are compared for the detection of generalized tonic-clonic seizures. Different sensors are placed over the body to study the impact of sensor placement on the performance. Also a multimodal algorithm was proposed, which combines the inputs from different sensors for an improved seizure detection performance. This approach requires epileptic activity in the used sensors, but also takes into account potential delays in epileptic activity between the different sensor modalities. Combinations of sensors from different modalities placed on different parts of the body showed to lead to the best performance. The false alarm rate is decreased by at least 90% in 57% of the multimodal combinations compared to the unimodal performance. These results demonstrate that algorithms combining different modalities should be used in order to obtain a sufficiently high performance for usage in practice, preferably by incorporating a personalized approach.

Beknopte samenvatting

De levensstandaard van refractaire epilepsiepatiënten is sterk verminderd door het plotse optreden van aanvallen. Deze kan verbeterd worden door het gebruik van thuismonitoring, met in het bijzonder automatische waarschuwingssystemen. Deze detecteren automatisch aanvallen en alarmeren hierbij meteen de familieleden van de patiënt. De hiervoor benodigde algoritmes voor automatische real-time aanvalsdetectie zijn de kern van dit doctoraatsonderzoek.

Verschillende algoritmes voor automatische aanvalsdetectie worden voorgesteld in deze thesis. De grootste focus gaat hierbij naar het detecteren van aanvallen met behulp van electrocardiografie (ECG) doordat onderzoek hiernaar in de literatuur momenteel beperkt is. Een patiënt-onafhankelijk algoritme wordt hierbij als eerste oplossing voorgesteld. Deze maakt gebruik van een nieuwe voorverwerkingsmethode die sterke hartslagtoenames detecteert die veroorzaakt zijn door veranderingen in het autonoom zenuwstelsel. Verschillende kenmerken worden uit deze hartslagtoenames geëxtraheerd, waarna ze geclassificeerd worden met behulp van een support vector machine classificator. Dit algoritme zorgt voor 50% minder valse alarmen in vergelijking met de literatuur.

Automatisch aanvalsdetectie is echter een zeer patiënt-specifiek probleem, en dus is personalisatie van de algoritmes nodig voor een optimaal resultaat. De huidige literatuur bevat echter onvoldoende onderzoek naar algoritmes voor een snelle en efficiënte personalisatie van deze algoritmes door gebruik te maken van slechts een beperkte hoeveelheid patiënt-specifieke data. Daarom worden er verschillende mogelijkheden voor personalisatie uitgewerkt in deze thesis. Elk personalisatiealgoritme heeft zijn specifieke use case waarbij deze optimaal tot zijn recht komt. In een eerste algoritme wordt transfer learning toegepast om het eerder uitgewerkte patiënt-onafhankelijke algoritme te personaliseren. Deze wordt periodiek hertraind, en is in staat om een robuste classificator te bekomen door gebruik te maken van slechts enkele geannoteerde aanvalsdata. Dit algoritme vereist nauwkeurig geannoteerde data, wat enkel mogelijk is bij opnames in een ziekenhuis. Daarom wordt er ook een tweede

personalisatiealgoritme voorgesteld. Deze past zich in real-time aan aan de hartslagkarakteristieken van de patiënt door aanvalsannotaties van de patiënt zelf te gebruiken. Aangezien deze annotaties niet altijd correct zijn, werd een extra procedure toegevoegd aan het algoritme om potentiële slechte annotaties automatisch te verwijderen. Het voorgestelde algoritme personaliseert snel, en blijkt robust te zijn tegen deze inaccurate patiëntannotaties.

Ook algoritmes voor personalisatie die ongeannoteerde patiënt-specifieke data gebruiken worden uitgewerkt in deze thesis. Deze zijn vooral van toegevoegde waarde in geval van het ontbreken van nauwkeurig geannoteerde patiënt data en bij het gebrek aan patiënt-specifieke aanvalsdata. Deze algoritmes zijn dus vooral interessant voor het lange-termijn monitoren van patiënten met een lage aanvalsfrequentie aangezien de hierboven vermeldde algoritmes minstens enkele patiënt-specifieke aanvalsdata vereisen. Twee verschillende methodes worden hiervoor voorgesteld. Een eerste algoritme past zich in real-time aan aan de patiënt, waarbij personalisatie van het algoritme reeds start na enkele uren. Het tweede algoritme maakt gebruik van een nieuw voorgestelde transfer learning methode, die offline getraind wordt om zich zo accurater aan te passen aan de hartslagkarakteristieken van de patiënt.

Deze gepersonaliseerde algoritmes leidden tot 30-57% minder valse alarmen dan het voorgestelde patiënt-onafhankelijke algoritme. Automatische personalisatie is dus aangeraden voor aanvalsdetectie op basis van ECG aangezien deze algoritmes in de praktijk gebruikt kunnen worden zonder (veel) extra inspanning van de patiënt/zorgkundige. De algoritmes zijn reeds in staat zich aan te passen aan de patiënt na enkele dagen monitoring. De voorgestelde algoritmes zijn vooral van toegevoegde waarde als onderdeel van een multimodaal algoritme, waarbij ze het meeste bijdragen bij de detectie van klinisch subtielere aanvallen.

In het laatste deel wordt aanvalsdetectie met behulp van verschillende modaliteiten (ECG, electromyografie en accelerometrie) met elkaar vergeleken voor de detectie van gegeneraliseerde tonisch-clonische aanvallen. Verschillende sensoren worden hierbij over het lichaam geplaatst om de impact van de sensorlocaties te bestuderen. Ook een multimodaal algoritme dat de informatie van verschillende sensoren combineert wordt hierbij uitgewerkt. Dit multimodaal algoritme vereist epileptische activiteit in elke gebruikte sensor, maar houdt hierbij rekening met de eventuele tijdsverschillen tussen epileptische activiteit in de verschillende modaliteiten. Combinaties van verschillende modaliteiten geplaatst op verschillende locaties leiden hierbij tot de beste resultaten. Het aantal valse alarmen daalt hierbij in 57% van de combinaties met minstens 90% ten opzichte van de unimodale algoritmes. Deze resultaten tonen aan dat algoritmes die gebruik maken van verschillende modaliteiten nodig zijn om een voldoende hoge performantie te bekomen voor gebruik in de praktijk, bij voorkeur door gebruik te maken van automatisch gepersonaliseerde algoritmes.

Nomenclature

Variables and Symbols

a, b, C	Vectors
$\mathbf{A}, \mathbf{B}, \mathbf{C}$	Matrices
$\|\|w\|\|$	Euclidean norm of vector w
$\mathbf{A}^T$	Transpose of matrix A
$\mathbf{A}_{i,j}$	ij-th entry of matrix A
$f(.)$	Function
$\min_x f(x)$	Minimization over x. Minimal value $f(x)$ is returned
$\arg\min_x f(x)$	Minimization over x. Optimal x is returned
$\max_x f(x)$	Maximization over x. Maximal value $f(x)$ is returned
$\arg\max_x f(x)$	Maximization over x. Optimal x is returned
$\mathcal{L}$	Laplacian
$\frac{\partial y}{\partial x}$	Derivative of y with respect to x
$K(x_i, x_j)$	Kernel function evaluated on vectors x_i and x_j
$\varphi(x)$	Feature map applied to x
σ^2	Bandwidth of an RBF kernel

Units

μV	microvolt
bpm	beats per minute
GB	Gigabyte
GHz	Gigahertz
h	hours
Hz	Hertz
mg	micro units of Earth's gravity
ms	milliseconds
s	seconds

Abbreviations

ACF	Autocorrelation function
ACM	Accelerometer
AED	Anti-epileptic drugs
ANS	Autonomic nervous system
ART	Arterial blood pressure
AVN	Atrioventrical node
B	Bihemispheric
BP	Blood pressure
C	Clonic
CI	Confidence interval
CNS	Central nervous system
CP	Complex partial
CVP	Central venous pressure
CWT	Continuous wavelet transform
DBS	Deep brain stimulation
ECG	Electrocardiogram
ECoG	Electrocorticogram
EDA	Electrodermal activity
EDR	Electrocardiogram-derived respiration
EEG	Electroencephalogram
EMG	Electromyogram
EOG	Elektrooculogram
FAR	False alarm rate
FIAS	Focal impaired awareness seizures
FL	Frontal lobe
FN	False negative
FP	False positive
GSR	Galvanic skin response
GPS	Global positioning system
HF	High frequency
HK	Hyperkinetic
HR	Heart rate
HRI	Heart rate increase
HRV	Heart rate variability
ICU	Intensive care unit
ID	Identifier
ILAE	International League Against Epilepsy
KKT	Karush-Kuhn-Tucker
LA	Left ankle
LF	Low frequency

LOPO	Leave-one-patient-out
LOPO CT	Leave-one-patient-out crosstesting
LOPO CV	Leave-one-patient-out crossvalidation
LW	Left wrist
MIX	Mixed model
MM	Multimodal
MRI	Magnetic resonance imaging
MVC	Maximal voluntary contraction
OC SVM	One-class support vector machine
PAP	Pulmonary arterial pressure
PD	Patient data
PI	Patient-independent
PPG	Photoplethysmography
PPV	Positive predictive value
PS	Patient-specific
PSD	Power spectral density
QoL	Quality of life
RA	Right ankle
RAM	Random access memory
RBF	Radial basis function
RD	Recording duration
RMSSD	Root mean square of the successive differences
RSA	Respiratory sinus arrhythmia
RW	Right wrist
SAN	Sinoatrial node
SC	Subclinical
SD	Seizure duration
SDNN	Standard deviation of NN intervals
SDSD	Standard deviation of successive differences
Se	Sensitivity
sEMG	Surface electromyography
SG	Secondary generalized
SP	Simple partial
SU	Subtle
SUDEP	Sudden unexplained death in epilepsy
SV	Support vector
SVM	Support vector machine
T	Tonic
TC	Tonic-clonic
TL	Transfer learning
TLE	Temporal lobe epilepsy
TN	True negative

TP	True positive
UC	Unclassified
vEEG	Video-Electroencephalography
VLF	Very low frequency
VNS	Vagus nerve stimulation
WBAN	Wireless body area network

Contents

Abstract v

Beknopte samenvatting vii

Nomenclature ix

Contents xiii

List of Figures xxi

List of Tables xxxi

1 General introduction 1

1.1 Problem statement 1

1.2 Research objectives 2

1.3 Chapter overview and personal contributions 3

1.3.1 Introduction 3

1.3.2 Patient-independent seizure detection 3

1.3.3 Personalized seizure detection approaches 5

1.3.4 Multimodal seizure detection 7

1.3.5 Appendices 7

1.4 Collaborations 8

2 Epilepsy, the heart and epilepsy monitoring solutions 11

2.1 Introduction to epilepsy 11

2.1.1 Definitions and diagnose 11

2.1.2 Seizure classification 13

2.1.3 Treatment 15

2.1.4 Mortality in epilepsy 17

2.2 The heart 18

2.2.1 Physiology of the heart 18

2.2.2 Electrocardiography 19

2.2.3 The autonomic nervous system 20

2.2.4 Heart rate variability 21

2.3 Epilepsy and the autonomic nervous system 22

2.4 Long-term monitoring of epilepsy patients 25

2.4.1 Automated seizure warning system 25

2.4.2 Modalities for automated seizure detection 27

2.4.3 Seizure logging system 36

2.4.4 Seizure prediction system 36

2.4.5 SUDEP monitoring and prevention 37

3 Patient-independent ECG-based seizure detection 39

3.1 Introduction to ECG-based seizure detection 40

3.2 Data Acquisition 42

3.3 Methodology 42

3.3.1 Online R peak detection 43

3.3.2 Extraction of significant heart rate increases 44

3.3.3 Feature extraction 45

3.3.4 Feature selection 47
3.3.5 Classification 48
3.3.6 Evaluation criteria 53
3.3.7 Tested HR-based seizure detectors from the literature . 53
3.4 Results and discussion 54
3.4.1 *HRI-EXTRACT* 54
3.4.2 Feature selection & γ optimization 55
3.4.3 Fastened algorithm *fHRI-SVM* 55
3.4.4 Results for *fHRI-SVM* 56
3.4.5 Comparison between *HRI-SVM* and *fHRI-SVM* 58
3.4.6 Comparison between *fHRI-SVM* and the literature . . . 59
3.4.7 Inter-patient variability 63
3.4.8 Relationship between clinical phenomena and seizure detection results 65
3.4.9 Applications of HR-based seizure detection 66
3.4.10 Remaining issues 68
3.5 Conclusion 69

4 Supervised transfer learning for personalization of ECG-based seizure detection 71

4.1 Introduction to personalization 71
4.2 Data Acquisition 73
4.3 Methodology 73
4.3.1 Preprocessing 75
4.3.2 Feature extraction 75
4.3.3 Patient-independent classification 76
4.3.4 Personalized classification through transfer learning . . 77
4.3.5 Alternative automatic personalization solutions 80

4.3.6 Algorithm evaluation 81
4.3.7 Impact of number of seizures in training 81
4.4 Results 82
4.5 Discussion 85
4.5.1 Performance comparison of the PI and TL approach . . 85
4.5.2 Comparison of different personalization approaches . . . 86
4.5.3 Impact of number of seizures on personalization 88
4.5.4 General discussion 89
4.6 Conclusion 90

5 Real-time personalization using patient feedback 91

5.1 Introduction 92
5.2 Data Acquisition 93
5.3 Methodology 94
5.3.1 Patient-independent seizure detection 94
5.3.2 Patient-specific seizure detection 96
5.3.3 Real-time adaptive seizure detection 97
5.3.4 Feedback procedure 100
5.3.5 Evaluation criteria and additional tests 101
5.4 Results 102
5.4.1 Performance of the different seizure detection classifiers 102
5.4.2 Seizure sensitivity 103
5.4.3 Impact of incorrect user feedback on real-time adaptive seizure detection 106
5.5 Discussion 106
5.5.1 Comparison between the different approaches 106
5.5.2 Seizure sensitivity 108
5.5.3 Impact of incorrect user feedback on performance 108

5.5.4 General discussion and possible improvements 109
5.6 Conclusion 110

6 Unsupervised nocturnal seizure detection through real-time personalization 111

6.1 Introduction 112
6.2 Data acquisition 113
6.3 Methodology 113
6.3.1 Preprocessing 113
6.3.2 Adaptive classification 114
6.3.3 Initial patient-independent classification 117
6.3.4 Performance evaluation 117
6.4 Results 118
6.5 Discussion 121
6.5.1 Sensivity and detection delay 121
6.5.2 False alarm rate 122
6.5.3 General discussion and possible improvements 122
6.6 Conclusion 123

7 Personalizing nocturnal seizure detection using unsupervised transfer learning 125

7.1 Introduction 126
7.2 Data acquisition 126
7.3 Methodology 128
7.3.1 Preprocessing and feature extraction 128
7.3.2 Patient-independent reference classifier 128
7.3.3 Unsupervised personalization 129
7.3.4 Simulation for evaluating the impact of training size . . 134

7.3.5 Performance evaluation 135
7.4 Results 136
7.5 Discussion 138
7.5.1 Performance comparison of the different approaches . . 138
7.5.2 Impact of seizure type on performance 141
7.5.3 Impact amount of data on performance 142
7.5.4 General discussion 143
7.6 Conclusion 144

8 Multimodal Generalized Tonic-Clonic Seizure Detection 145

8.1 Introduction 146
8.1.1 Tonic-clonic seizure detection using EMG 146
8.1.2 Tonic-clonic seizure detection using accelerometry . . . 147
8.1.3 Multimodal seizure detection 148
8.2 Data acquisition 149
8.3 Methodology 149
8.3.1 Seizure detection using EMG 149
8.3.2 Seizure detection using accelerometry 151
8.3.3 Seizure detection using ECG 152
8.3.4 Personalized seizure detection using ECG 152
8.3.5 Multimodal methodology 152
8.3.6 Algorithm evaluation 154
8.4 Results 154
8.4.1 Unimodal results 154
8.4.2 Multimodal results 157
8.4.3 Unimodal personalization in a multimodal combination 158
8.5 Discussion 160

8.5.1 Comparison of different unimodal approaches 160
8.5.2 Multimodal tonic-clonic seizure detection 162
8.5.3 Added value of unimodal personalization in a multimodal algorithm 164
8.5.4 General discussion 165
8.6 Conclusion 167

9 Conclusions and future directions 169

9.1 Conclusions 169
9.1.1 ECG-based patient-independent seizure detection 169
9.1.2 Personalization of ECG-based seizure detection 170
9.1.3 Multimodal seizure detection 173
9.2 Future directions 174

A Features extracted for automated seizure detection 179

A.1 ECG-based seizure detection features extracted during *HRI-EXTRACT* 179
A.1.1 Features extracted during the HRI 179
A.1.2 Features extracted from pre-HRI period 182
A.2 EMG-based seizure detection features 183
A.3 ACM-based seizure detection features 184

B Heart beat detection in multimodal data 187

B.1 Introduction 188
B.2 Datasets and challenge information 189
B.3 Methodology 189
B.3.1 Automatic relevant signal detection using power spectral density similarity 190
B.3.2 Individual peak detection 191

B.3.3 Combination of individual peak locations 193
B.3.4 Heart beat detection using signal type identifiers 195
B.4 Results 196
B.4.1 Automatic signal detection 196
B.4.2 Multimodal heart beat detection 198
B.5 Discussion 199
B.5.1 Automatic signal detection 199
B.5.2 Multimodal heart beat detection 201
B.6 Conclusion 203

Bibliography **205**

Curriculum vitae **227**

List of publications **229**

List of Figures

1.1 Schematic overview of the chapters in this thesis. ECG: electrocardiogram, EMG: electromyogram, ACM: accelerometer, PI: patient-independent, PS: patient-specific. 4

1.2 Overview of the different use cases for personalization investigated in this thesis. Upper left (Chapter 4): Using accurately annotated data from the hospital to offline personalize seizure detection using supervised transfer learning (applicable day and night). Upper right (Chapter 5): Using real-time patient feedback to previously generated alarms to adapt to the patient-specific characteristics using a heuristic support vector machine updating procedure (applicable day and night). Lower right (Chapter 6): Real-time personalization without requiring seizure annotations and user input using a simplified classification procedure (only applicable during the night). Lower left (Chapter 7): Offline personalization (can be done during the day) using unsupervised transfer learning without requiring seizure annotations and user feedback (seizure detection only applicable during the night). . 6

2.1 Effect of age on epilepsy prevalence. Taken from [85]. 12

2.2 New ILAE seizure classification guidelines 2017. Taken from [61]. 13

2.3 Overview of the different lobes in the brain. Taken from [134]. 14

2.4 Overview of the physiology of the heart. RA: right atrium, LA: left atrium, RV: right ventricular, LV: left ventricular, SAN: sinoatrial node, AVN: atrioventricular node. Taken from [108]. 18

2.5 Placement of sensors for a standard 12-lead ECG. Taken from [203]. 19

2.6 Example of an ECG segment to illustrate the different ECG morphologies. Taken from [194]. 20

2.7 Illustration of the most found ictal heart rate pattern. The dashed lines indicate the interval where typically the seizure onset is located with respect to this heart rate pattern. 23

2.8 Examples of four ictal heart rate patterns from four different patients. (a) Pattern including monotonically increasing heart rate, heart rate stabilization and fast decreasing heart rate. An example of strong post-ictal heart rate changes is illustrated in this figure. (b) A similar pattern as in (a), but with high frequency alterations during the heart rate increase and a much less steeper decrease in heart rate. The heart rate increase also starts pre-ictally, reaching the maximal heart rate around seizure onset. (c) The heart rate increases linearly towards a peak heart rate, then slightly decreases and increases again towards the same peak heart rate. (d) Similar pattern as in (c), but the first reached heart rate peak is much lower than the maximal heart rate peak achieved in the second increase. The heart rate decreases nearly to bradycardia after the heart rate increase. . 24

2.9 Example of EEG signals during a seizure. 28

2.10 Example of ictal ECG and PPG activity acquired synchronously. PPG is obtained in this example from the wrist. Epileptic activity can be recognized by the increased frequency of the heart cycles. 30

2.11 Example of changes in the accelerometry along the different axes during a generalized tonic-clonic seizure. The accelerometer signal is acquired from the wrist in this example. The clonic phase of the seizure starts around 5 s after seizure onset. 31

2.12 Example of changes in a biceps EMG signal during a generalized tonic seizure. 32

2.13 Comparison of the two typical options for multimodal seizure detection. (a) Illustration of late integration, in which feature extraction and classification occur per individual modality. Afterwards, the unimodal results are combined into 1 decision output. (b) Illustration of early integration, in which the extracted features from different modalities are combined in one feature vector. This feature vector is the input for a multimodal classifier, which decides on whether there is epileptic activity present or not. Note that in both procedures the classifier steps can be replaced by any kind of decision steps (e.g. rule-based approaches). 35

3.1 Illustration of ictal heart rate changes during temporal lobe seizures. 40

3.2 Example of an ictal HRI from patient 17. The features required for the extraction of significant HRIs are defined in this figure. $\overline{HR}_{pre}$ indicates the average HR over the 60s before the start of the HRI. HR_{start} indicates the HR at the beginning of the HRI, HR_{peak} the peak HR at the end of the HRI. Δt_{HRI} indicates the length of the HRI (in seconds), ΔHR indicates the amount of increased HR. The EEG seizure onset is located at 33922s in this example. 46

3.3 Overview of the proposed patient-independent seizure detection procedure followed in this chapter. 53

3.4 Example of ictal noise in the ECG. 54

3.5 Histogram of the detection delay between the seizure onset and the alarm time for the proposed seizure detection algorithms. . 58

3.6 Sensitivity performance for seizures with and without tachycardia for the different tested algorithms. 60

3.7 Boxplots showing the inter-patient variability in sensitivity (a), false alarm rate (b), PPV (c) and detection delay(d) for the different tested algorithms. 61

3.8 Patient-specific sensitivity, FP/hour, PPV and detection delay results for the different tested patient-independent seizure detection algorithms. 62

3.9 Sensitivity versus FP/hour results for the different tested algorithms. For fHRI-SVM and the method from [38] the different performance results were found by altering the SVM bias term b. For the Osorio and vanElmpt algorithm, the parameter values from the papers are altered. For Osorio, each line contains a constant delay term D with varying heart rate threshold values [142]. For vanElmpt, each line contains a constant second time window Δt_2, with changing heart rate threshold values [188]. We refer to the corresponding papers for details of these algorithms and their parameters. 63

3.10 Comparison of a seizure and a non-seizure HRI originating from the same patient. Both HRIs were detected during sleep, within only a few hours of each other. 64

3.11 Sensitivity results for each class of seizure severity for the different tested algorithms. The average sensitivity over all algorithms for the 3 severity classes is also shown in this figure. 66

4.1 Illustration of the effect of transfer learning on the performance compared to traditional machine learning approaches. 73

4.2 Overview of the proposed supervised transfer learning approach for personalized heart rate-based seizure detection. PS TL SVM: patient-specific transfer learning support vector machine; PI SVM: patient-independent support vector machine. 77

4.3 Sensitivity, FAR and F_3-score per patient with patient average (P-Av) and overall average (T-Av) performances for the patient-independent (PI) and transfer learning (TL) algorithm 83

4.4 Boxplots of sensitivity, FAR and F_3-score for the patient-independent (PI), fully patient-specific (PS), mixed (MIX) and transfer learning (TL) algorithm. 83

4.5 Impact of the number of seizures on the F_3-score performance for the full patient-specific (PS) approach, mixed approach and transfer learning (TL) approach. In case 0 seizures are available for training, the performance of the patient-independent algorithm is shown. 84

4.6 Examples of atypical ictal heart rate behavior within a patient. (a) Example of 3 typical and 1 atypical ictal heart rate patterns for patient 13. (b) Example of a atypical ictal small heart rate increase and 3 normal ictal heart rate increases for patient 6. The vertical line in both figures indicate the seizure onset of the different seizures. The shown tachogram signals are filtered using the median filter discussed in Section 4.3.1. 86

4.7 Visualization of SVM boundary of the patient-independent (PI), patient-specific (PS) and transfer learning (TL) model for different D values (0.01 and 1). 87

5.1 Illustration of added value of real-time personalization compared to periodic offline personalization. The online adaptive algorithm allows for a lot of performance increase (illustrated with the marked area) during the different periodic updates, certainly in the beginning of the recordings of each patient. 92

5.2 Overview of the proposed real-time adaptive seizure detector. . 97

5.3 Sensitivity and FAR results per patient for the different classifiers. 103

5.4 Illustration of convergence for the adaptive classifier. (a) Average FP/h per 12 hours of recording in function of time for patients 1 and 10 for the three different approaches. (b) Example of the SVM parameter convergence for patient 10. The different lines in the w_t figure indicate the terms for the 3 features. 104

5.5 (a)Effect of time after the start of recordings on the false alarm rate (FAR) of the patient-independent (PI), patient-specific (PS) and adaptive (Ad.) algorithm averaged over all patients and per 12 hours for the first 5 days of data of each patient. (b) Reduction in false alarm rate for the adaptive algorithm compared to patient-independent algorithm in function of time after the start of the recording (averaged over all patients). The rhythmic behavior of the FAR results in (a) are caused by synchronized day-night differences in all recordings, causing large differences in FAR rates between two subsequent averages. 104

5.6 Impact factors on seizure sensitivity for *HRI-EXTRACT* and the proposed adaptive seizure detection classifier. (a) Impact of seizure type on sensitivity. The numbers between brackets indicate the number of seizures per type. (b) Impact of seizure duration on sensitivity. 105

5.7 Boxplots of the seizure duration for each seizure type analyzed in the dataset. 105

5.8 Impact of incorrect feedback on the performance. (a) Effect of different percentages of incorrect user feedback on sensitivity, FP/hour rate and PPV for the proposed adaptive seizure detection approach. The variation in performance between different runs is shown by the variation bars. (b) Comparison of the PPV values for the proposed updating procedure and the traditional SVM updating procedure in case 0, 0.1, 0.5, 1, 5 and 10% of user feedback is incorrect. 107

6.1 Illustration of the proposed adaptive seizure detection algorithm. The decision boundary is formed by the ellipse: data points outside the ellipse cause a warning if they also have a peak heart rate higher than the peak heart rate of the center of the ellipse. In this example, a new seizure data point needs to be classified (marked as a star). 116

6.2 Overview of the proposed real-time adaptive seizure detection algorithm. The part in the dashed box shows the initial offline procedure that only needs to be done once. HRI: heart rate increase, PI: patient-independent. 117

6.3 Boxplots for sensitivity (a) and detection delay (b) for the different clinical seizure types. TC: Tonic-clonic, HK: hyperkinetic, T: tonic, C: clonic, SU: subtle seizures. 120

6.4 False alarm rate per type for patients for the original patient-independent (PI) algorithm and the proposed adaptive (Ad.) algorithm. GS: generalized seizures, FL: frontal lobe, TL: temporal lobe. 120

7.1 Overview of the proposed unsupervised transfer learning approach for personalized heart rate-based seizure detection. PS TL SVM: patient-specific transfer learning support vector machine; PI SVM: patient-independent support vector machine. 133

7.2 Illustration of the used procedure for automated estimation of parameter ν in (7.3). $peak_{index}$ indicates the patient-independent (PI) SVM output for which the highest likelihood is found in the available patient-specific data. The cutoff value (the vertical red line) is found by subtracting $peak_{index}$ with a fixed value δ. The blue area indicates the data points which are seen as outliers for this patient. Parameter ν is then estimated as the blue area over the entire area under the likelihood curve. 134

7.3 Impact of seizure type on performance of the proposed unsupervised transfer learning approach. The FAR and sensitivity results are obtained for different values for δ (range:[-2.9,2.9], in steps of 0.2, only showing results with a FAR smaller than 2 FP/night). Sensitivity results are here obtained as seizure-based overall results. 137

7.4 Boxplots of the ictal peak heart rate (a) and ictal maximal heart rate gradient (b) for different seizure types as found by the used preprocessing procedure. TC: tonic-clonic, HK: hyperkinetic, T: tonic. 138

7.5 Impact of number of nights in training on performance of the proposed personalized transfer learning approach. 139

7.6 Example of the boundaries of the different classifiers. PI SVM: patient-independent SVM classifier, OC SVM+TL: one-class SVM classifier with transfer learning, OC SVM: one-class SVM without transfer learning using only patient-specific data. Data points in yellow indicate non-epileptic data points, pink data points indicate seizure data points. 140

7.7 Boxplots of the ictal peak heart rate (a) and ictal maximal heart rate gradient (b) for generalized tonic-clonic (TC) and hyperkinetic (HK) seizures for both the mainly adults patients of the dataset mentioned in this chapter and the pediatric patients described in Chapter 6. 141

8.1 Illustration of ictal activity in accelerometry (ACM), electromyogram (EMG) and electrocardiography (ECG) during a tonic-clonic seizure. The vertical red line indicates the clinical seizure onset. The tonic phase of this seizure ends at 24 s, after which the clonic phase starts. During the tonic phase, the high frequency activity of the EMG is increased. During the clonic phase, strong periodic patterns can be seen in the three ACM channels. During both tonic and clonic phase, the heart rate gradually increases (RR intervals decrease). 147

8.2 Sensor placement for the different modalities in the dataset used in this chapter. Accelerometers (ACM): placed at both wrists and ankles. Surface electromyogram (sEMG): located at both biceps, consists of two sensors, one placed at the belly of the muscle and one placed on the tendon. This results in one EMG signal per biceps. Electrocardiogram (ECG): two sensors placed at the chest, one near the right shoulder and one near the left hip, which lead to the lead II ECG signal. 150

8.3 Overview of the proposed multimodal (MM) seizure detection algorithm. 154

8.4 Sensitivity over false alarm rate ROC-like curve for the evaluated unimodal seizure detection algorithms. 156

8.5 Boxplots indicating the inter-patient FAR variability between the different unimodal seizure detection algorithms. ECG-pers.: personalized ECG-based seizure detection method described in Section 8.3.4. 157

8.6 Boxplots for detection delay for the different unimodal seizure detection algorithms. 158

8.7 Average false alarm rate (FAR) for the combinations with patient-independent (PI) and patient-specific (PS) ECG-based seizure detection. 159

8.8 Impact of different sensors on the multimodal F_3 performance. (a) F_3 score improvement by including a sensor to an unimodal algorithm grouped per modality. (b) Average multimodal F_3 score for each sensor involved in the multimodal setting. Differentiation is made for combinations with a different modality and combinations of the same modality. 163

8.9 Average F_3 scores for each sensor when combined with another sensor, averaged out for each type of modality. 163

9.1 Illustration of a multimodal master-slave setup. In this example, the ECG sensor is the master sensor, and the ACM and EMG sensors are the so-called slaves. The ACM and EMG sensors are only active when they are awoken by the master sensor after epileptic activity (marked in green) is detected. After a period T after the last epileptic activity in the master, the slaves are turned off again. A multimodal alarm is then generated if all three sensors contained epileptic activity (see second master activation). 176

9.2 Illustration of automated selection of sensors for a multimodal seizure detection system. 177

B.1 Example of a multimodal data segment. In this example the BP signal and the EOG signal with R peak interference can improve the performance of the heart beat detection in case the ECG signal would have been noisy. 188

B.2 Overview of the algorithms proposed in this study. In the MM_{NI} algorithm (a), only the signal type of the first signal is certainly known (ECG). In the MM_{WI} algorithm (b), all available signal IDs are used. Crosses indicate after which steps signals can be removed from further processing. 190

B.3 Illustration of the automatic signal detection procedure. The left figure shows four signals during a 6 seconds epoch of the selected noise-free ECG period. The PSD values of all signals are shown in the right figure. The dotted line in the right figure indicates the average heart rate frequency HR_{av} in the selected noise-free period in the ECG signal. The frequency segment that is used for computing the correlation coefficient is marked by the solid vertical lines. In this example only the BP and stroke volume (SV) signal are selected as signals of interest. 192

B.4 Illustration of the majority voting procedure for combining heart beats from different signals. Signals $vote_1$ and $vote_2$ indicate temporary versions of $vote$ for each signal separately. The peaks from the BP signal are indicated by diamond symbols, the arrows indicate the peaks for this signal after alignment. Signal $vote$ is defined as the sum of $vote_1$ and $vote_2$. The real R peak annotations are marked by an asterisk. 193

B.5 Illustration of the combination of peak locations of different signals. The red circles indicate the initial set of R peak locations *peaks*. In this case, peak locations from ECG (narrow peaks) and BP (wide peaks) are combined. The chained lines show the estimated R peak locations $estloc_2^n$, which lay near the ECG peaks. The ECG peak locations are thus added to *peaks* in this example. 194

B.6 Percentages of signals with given identifier that are detected by the proposed method for automatic signal detection. The first bars indicate the sensitivity percentages of each signal type after inspecting the PSD similarity (see Section B.3.1), the seconds bars indicate the percentages of used signals after removing the signals with too high variability in delay values (see Section B.3.2). 197

List of Tables

3.1 Overview of the used dataset in this chapter. 43

3.2 Comparison of F_β results for different typical heart rate-based seizure detection results during training. All F_β-scores, PPV and sensitivity (Se) are shown in percentages. 48

3.3 Results for fHRI-SVM in case a fixed γ is used for all patients during SVM training. 55

3.4 Results per patient for *fHRI-SVM*. 57

3.5 Comparison of the results for the different evaluated algorithms. Bold numbers indicate the algorithms with similar sensitivity or FP/h rate as the proposed algorithm. Values for sensitivity (Se) and positive predictive value (PPV) are given in percentages, detection delay is expressed in seconds. 58

4.1 An overview of the dataset used in this chapter. RD: Recording Duration, SD: Seizure Duration, B: Bihemispheric, R: right, L: left, Nc: not clear, Te: Temporal, Fr: Frontal, Oc: Occipital. . 74

4.2 Results for the patient-independent (PI), fully patient-specific (PS), mixed (Mix) and transfer learning (TL) approach. Both patient average (P-Av) and overall average (T-Av) are shown in the table. 82

5.1 Overview of the used evaluation dataset which contains a subset of the EPILEPSIAE database. 95

5.2 Overview of the results of the different seizure detection classifiers. 102

6.1 Overview of the datasets used in this chapter. 114

6.2 Overview of the results of the proposed adaptive approach for different presumed seizure onsets. 119

7.1 Detailed overview of the dataset used for testing in this chapter. TC: tonic-clonic seizure, HK: focal hyperkinetic seizure, T>30: generalized tonic seizure with a duration of more than 30 s. . . 127

7.2 Overview of the overall results of the different applied methodologies in this chapter. PI: patient-independent, OC SVM: one-class SVM, OC SVM+TL: proposed one-class SVM with transfer learning. 136

7.3 Performance of the proposed transfer learning approach per seizure type. TC: tonic-clonic seizure, HK: focal hyperkinetic seizure, T>30: generalized tonic seizure with a duration of at least 30 s. 137

8.1 Overview of the used δ values for multimodal seizure detection (see (8.3)). 153

8.2 Overview of the unimodal and multimodal results. Bold values indicate the best value for a specific evaluation metric for the unimodal and multimodal algorithms. Se: sensitivity, FAR: false alarm rate, PI: patient-independent, PS: patient-specific. 155

B.1 Results (in %) for the different proposed methods for multimodal heart beat detection. Also the results for the other final official top 5 challenge entries and for the unimodal *gqrs* algorithm are shown. $MM_{NI} - limited$ and $MM_{WI} - limited$ indicate the limited versions of both algorithms which can only use two signals. The last column indicates the average performance over both discussed datasets. 198

B.2 Results for the different proposed methods for multimodal heart beat detection on the hidden challenge test set. The old version of the MM_{WI} algorithm is the algorithm as described in [39]. . 199

Chapter 1

General introduction

1.1 Problem statement

Epilepsy is a neurological disease that affects around 1% of the people worldwide. Patients for who their seizures can not be controlled by currently available treatments, suffer from a strongly decreased quality of life due to the unpredictable character of these seizures. Monitoring of these patients in their home environment can improve this, as it allows them to feel safer in their everyday life. Monitoring of these patients can also increase the effectiveness of the treatment, allowing to give more objective measurements of a patient's health status in terms of an automatically generated seizure diary.

Monitoring at home of refractory epilepsy patients became a topic of interest during the last decades. Automated detection of epileptic seizures allows the usage of automated seizure warning systems and seizure diaries, which are beneficial for both patients and clinicians. Except for a small range of seizure types, the results from these algorithms in the literature however currently lead to an insufficient accuracy to allow real usage in practice. The main goal of this thesis is to improve currently available seizure detection algorithms by proposing more advanced procedures for efficient algorithm personalization and multimodal combination in order to further close the bridge between research and actual usage of seizure warning systems.

Most research for seizure detection at home went into the detection of seizures including strong movements. Studies however showed that also strong changes in the heart rate can be found during seizures. Heart rate-based seizure detection algorithms are currently however not well investigated in the literature. As an

initial goal of this thesis, the patient-independent seizure detection using the heart rate is improved by implementing a new preprocessing procedure and applying machine learning.

Most algorithms from literature use so-called patient-independent seizure detection algorithms. These are easy to use in practice as they do not require personal data from the test subject, but in most cases this leads to suboptimal results due to the high inter-patient variability in epileptic seizures. Efficient personalized approaches are however not investigated thoroughly in the application of seizure detection at home. In this thesis, different personalization approaches will be proposed which are usable in different use cases in practice. The personalization will be performed by only using a limited amount of patient data in order to remain usable in practice. As a use case, the different personalized approaches are performed for heart rate-based seizure detection, but they are applicable to seizure detection using any type of modality.

In most cases, seizure detection in the literature is performed using only one modality. The combination of multiple modalities could however lead to a sufficiently high performance which could lead to actual usage of such systems in practice. In this thesis, also the comparison and combination of different modalities is investigated to indicate the added value of a multimodal algorithm for the detection of generalized tonic-clonic seizures in pediatric patients.

All proposed methods could be used in order to facilitate and improve different use cases of monitoring of refractory epilepsy patients in their home environment. The different research objectives of this thesis are summarized in Section 1.2. The proposed methods and the motivation for them are shortly described on a chapter-by-chapter basis in Section 1.3.

1.2 Research objectives

The goal of this thesis is to develop new algorithms for automated real-time seizure detection in a home environment. This is done through following research objectives:

1. Applying a machine learning approach to obtain a patient-independent heart rate-based seizure detection algorithm in epileptic patients. This is done in order to improve the current mainly rule-based methods from the literature (Chapter 3).

2. Investigate different solutions for automated personalization of heart rate-based seizure detection. These options include both unsupervised and

supervised personalization approaches, as well as offline and real-time personalization solutions (Chapters 4-7).

3. Formulation of transfer learning frameworks for supervised and unsupervised problems (Chapters 4 and 7).
4. Development of real-time adaptive classifiers for both supervised and unsupervised learning (Chapters 5 and 6).
5. Compare different unimodal algorithms for generalized tonic-clonic seizure detection in terms of performance and wearability in pediatric patients (Chapter 8).
6. Create a multimodal algorithm that accurately detects generalized tonic-clonic seizures with a limited amount of sensors (Chapter 8).

1.3 Chapter overview and personal contributions

The outline of the complete thesis is depicted in Figure 1.1. The content of the different chapters and the personal contributions are summarized below.

1.3.1 Introduction

Chapter 2 gives an introduction to several topics of this thesis. Different definitions related to epilepsy and seizure types are given. The working of the heart and the autonomic nervous system is described, as well as how they change during epileptic seizures. Also different (both existing and potential future) solutions for long-term monitoring of refractory epilepsy patients are described. An overview of the different modalities that can be used for automated seizure detection to facilitate these monitoring solutions is also given in this chapter.

1.3.2 Patient-independent seizure detection

Currently, only a limited amount of research in the literature went to the automated detection of seizures using heart rate information. Therefore, a new patient-independent heart rate-based seizure detection algorithm is proposed and discussed in **Chapter 3**. It automatically detects strong heart rate increases by using a newly proposed preprocessing procedure, and classifies the hereof extracted features as either epileptic or non-epileptic using a supervised weighted support vector machine (SVM). The state-of-the-art ECG-based seizure

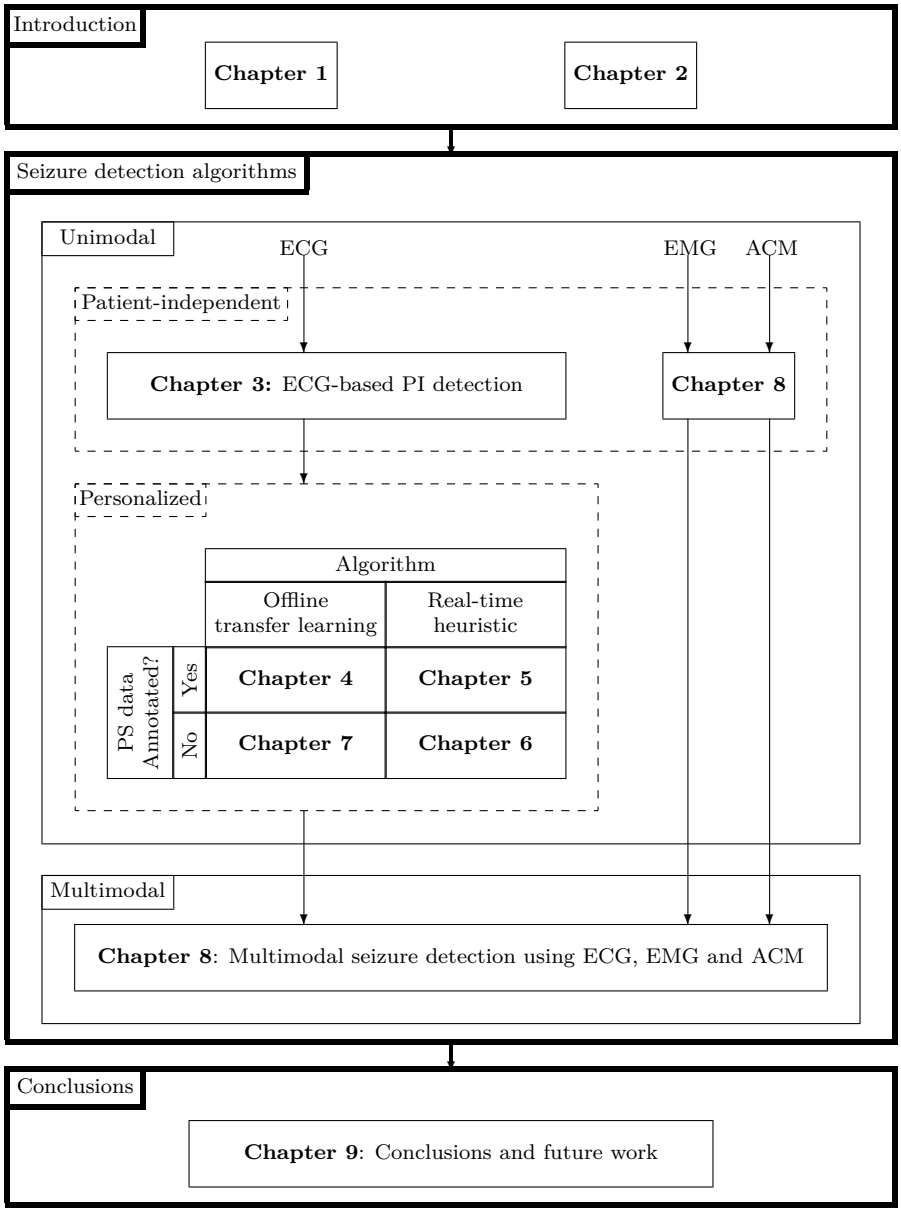


Figure 1.1: Schematic overview of the chapters in this thesis. ECG: electrocardiogram, EMG: electromyogram, ACM: accelerometer, PI: patient-independent, PS: patient-specific.

detection algorithms from the literature were also implemented and evaluated on the same dataset. The proposed algorithm proved to outperform the algorithms from the literature by reducing the average false alarm rate with 50%. This contribution is published in [44].

1.3.3 Personalized seizure detection approaches

Supervised personalization

Chapters 4-7 describe different approaches for personalization of heart rate-based seizure detection. Each proposed algorithm is usable in practice in a different use case as summarized in Figure 1.2. First, a supervised personalization algorithm is proposed in **Chapter 4**. It personalizes the seizure detection algorithm discussed in Chapter 3 using a newly proposed supervised weighted SVM transfer learning approach, which is able to solve the large class imbalance in the applied dataset. This way, it can be personalized offline with a limited amount of annotated patient-specific data for day and night monitoring of refractory patients. Accurately annotated data can originate for example from the prechirurgical evaluation in the hospital in this case. It showed to lead to 35% less false alarms while maintaining a similar sensitivity as the reference patient-independent approach by using only a couple of days of patient-specific data.

Accurately annotated data can only originate from the hospital, but only a limited amount of data can be acquired in that case. Therefore a method is described in **Chapter 5** that can adapt in real-time to the patient's heart rate characteristics by using data that can be acquired outside the hospital. The initial patient-independent SVM classifier is updated in real-time using a new heuristic SVM updating procedure which adapts the classifier quickly and robustly to the characteristics of the new patient's data. In a home environment, no accurate annotations from clinicians are available, so the annotations are assumed to be made by the patients themselves. As the seizure annotations made by patients might be inaccurate, an additional new procedure based on the principles of transfer learning was developed to counter potential incorrect patient feedback. The real-time heuristic adaptation allows an immediate performance increase, leading to a continuous increase in performance rather than a periodic performance increase. The results show a decrease in false alarms of 30%, and the method is robust against potential incorrect patient feedback. This work is published in [41].

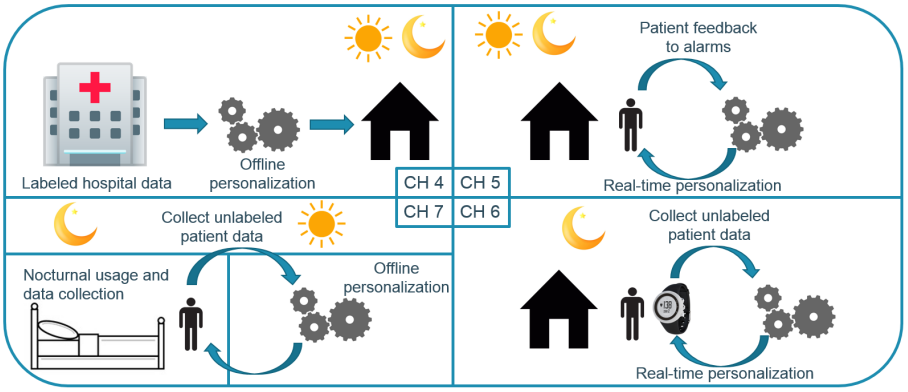


Figure 1.2: Overview of the different use cases for personalization investigated in this thesis. Upper left (Chapter 4): Using accurately annotated data from the hospital to offline personalize seizure detection using supervised transfer learning (applicable day and night). Upper right (Chapter 5): Using real-time patient feedback to previously generated alarms to adapt to the patient-specific characteristics using a heuristic support vector machine updating procedure (applicable day and night). Lower right (Chapter 6): Real-time personalization without requiring seizure annotations and user input using a simplified classification procedure (only applicable during the night). Lower left (Chapter 7): Offline personalization (can be done during the day) using unsupervised transfer learning without requiring seizure annotations and user feedback (seizure detection only applicable during the night).

Unsupervised personalization

After proposing two supervised methods for personalization, also two unsupervised approaches are discussed, which means no seizure annotations are required for the new patient's data. This is useful to avoid requiring possibly inaccurate annotations or having only a very limited amount of accurately annotated patient data. These unsupervised approaches also do not require patient-specific seizure data in order to personalize the algorithm, which is important for patients with a low seizure frequency. **Chapter 6** proposes a new heuristic unsupervised method that in real-time adjusts to the patients' characteristics for nocturnal seizure detection in pediatric patients with epilepsy. This is done by fitting a new low-complexity classifier that allows to adapt in real-time to patient's heart rate characteristics after only a couple of hours of monitoring. It leads to a decrease in false alarms of 57% compared to the reference patient-independent classifier. This work is published in [47].

Chapter 7 discusses an offline personalization method that uses a newly proposed one-class classification method for personalization of nocturnal seizure detection. A new transfer learning approach for SVM is proposed to personalize the reference patient-independent classifier, this time without requiring seizure annotations of the patient-specific data. It is tested on a large dataset of mainly adult patients containing weeks of nocturnal data per patient. An overall false alarm rate decrease of 33% is found compared to the reference patient-independent performance, with performance convergence found after two weeks of personal data. A precursor of this work is described in [45].

1.3.4 Multimodal seizure detection

A detailed comparison of different unimodal patient-independent seizure detection algorithms using accelerometry, electrocardiography and electromyography is discussed in **Chapter 8**. The modalities are evaluated for different sensor placements (seven in total). It is the first time in the literature that these three modalities are compared with each other on one dataset for nocturnal tonic-clonic seizure detection in pediatric patients. The different unimodal algorithms are also combined to form a multimodal seizure detector. This is done by a new combination method that takes into account the time differences between ictal activations in the different modalities. The proposed procedure for combining different unimodal approaches show to lead to decreases in false alarm rate from 75% up to over 90% by adding only one sensor from a different modality to an unimodal algorithm. It is also illustrated that personalization of one modality (using the approach proposed in Chapter 6) also leads to a strong increase in multimodal performance. A precursor of this work was described in [46].

Finally, **Chapter 9** summarizes all findings of this thesis and suggests directions for future research in this topic.

1.3.5 Appendices

Appendix A describes in detail the different features that are extracted during the newly proposed preprocessing procedure in Chapter 3. Both heart rate features from the literature and features specifically for the new preprocessing procedure are extracted. Also the features extracted from the accelerometry and electromyography in Chapter 8 are described in this chapter.

A new algorithm for heart beat extraction from noisy multimodal data is proposed in **Appendix B**. This algorithm was sent in for submission to

the Physionet/Computing in Cardiology Challenge 2014. The algorithm automatically detects signals from the multimodal dataset which contains information that could help with the detection of heart beats (e.g. ECG and blood pressure signals, but also signals contaminated with noise by such a signal) based on the similarity of the power spectral densities. The heart beat locations obtained from the different signals of interest are then combined using a new combination procedure based on heart beat location estimation, Hjorth's mobility and majority voting. The proposed algorithm won the second phase of the competition and received the third place during the final phase of the competition. This work is presented in [40].

1.4 Collaborations

This PhD research was conducted within the Biomedical Data Processing (Biomed) research group, STADIUS, Department on Electrical Engineering (ESAT), KU Leuven, under the supervision of Prof. Sabine Van Huffel.

Different clinical partners were involved with this thesis. This research was possible due to a strong collaboration with UZ Leuven, both the Department of Child Neurology (Prof. Lieven Lagae, Prof. Katrien Jansen) and the Department of Neurology (Prof. Wim Van Paesschen). Also a strong collaboration with the UZ Antwerp and Pulderbos Revalidation Center (Prof. Berten Ceulemans, dr. Anouk Van de Vel) was maintained. These clinical partners helped to understand the obtained results throughout the thesis from a clinical point of view.

In addition, a collaboration with the SeizeIT consortium was set up, which allowed to gather extra insights and data for extra validation of the algorithm proposed in Chapter 4.

Also international collaborations were established during this PhD. A collaboration was set up with DTU (Prof. Helge Sørensen) and Zealand University Hospital (Prof. Troels Kjær) during a three month research visit in Denmark. This led to the work presented in Chapter 5. In addition, a collaboration was set up with Kempenhaeghe research center and University of Technology Eindhoven (TU/E) (Prof. Johan Arends) within the Dutch Tele-Epilepsy Consortium, which led to the results and findings described in Chapter 7.

I would like to thank all collaborators, which allowed me to gain insights from different research and industry fields and allowed me to work with a large amount of data from a wide variety of patient groups and different clinical

centra. Also thank you to FWO (former IWT) for granting me an IWT SBO PhD scholarship.

Chapter 2

Epilepsy, the heart and epilepsy monitoring solutions

Before going into the details of the thesis, first an introduction is required to the different topics that are under research in this thesis. As the thesis topic suggests, the focus of this research is on the detection of epileptic seizures in a home environment. Some of the main aspects of epilepsy such as definitions, seizure classification and treatments are discussed in Section 2.1. As most of the proposed algorithms discuss heart rate-based seizure detection, also a short introduction is given to heart analysis and the changes in heart characteristics during epileptic seizures in Section 2.2 and 2.3. An overview of possible (both existing and possible future) technological solutions for monitoring epilepsy patients in their home environment is discussed in Section 2.4.

2.1 Introduction to epilepsy

2.1.1 Definitions and diagnose

Epilepsy is one of the most common neurological diseases, affecting around 1% of the worldwide population [89]. It is characterized by sudden disorderly discharging of the brain's nerve cells. These discharges are called epileptic seizures. Usually, seizures are self-limiting, occur unprovoked and happen without a prior warning. In some cases, however, seizures can be provoked by the occurrence of flickering lights, stress, alcohol and sleep deprivation [168].

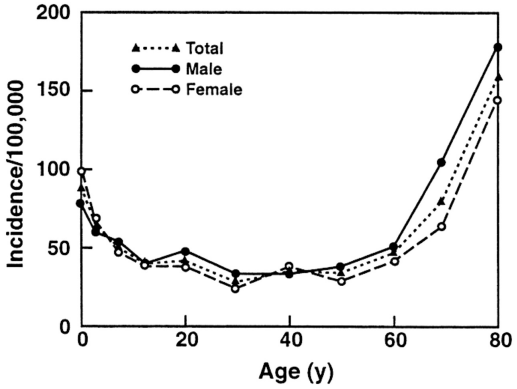


Figure 2.1: Effect of age on epilepsy prevalence. Taken from [85].

Seizures can range in duration from a couple of seconds to a few minutes. They can lead to a number of different symptoms, including changes in autonomic function, spasms and loss of consciousness.

Epilepsy is a disease of the brain defined by any of the following conditions [58]:

- At least two unproved seizures occurred more than 24 h apart;
- One unprovoked (or reflex) seizure and a probability of further seizures similar to the general recurrence risk (at least 60%) after two unprovoked seizures, occurring over the next 10 years;
- Diagnosis of an epilepsy syndrome.

Epilepsy can be caused by different factors. They can be summarized under the VITAMINS acronym [182]: vascular (e.g. deprivation of oxygen at birth, stroke), infections (e.g. meningitis), traumatic, autoimmune (e.g. central nervous system vasculitis), metabolic (e.g. alcohol/drugs), idiopathic (genetic), neoplasm (e.g. brain tumor) or structural (e.g. brain malformation) causes. The effect of these causes can be seen in the prevalence of epilepsy (see Figure 2.1). Some of the causes (e.g. brain malfunction and deprivation of oxygen at birth) cause epilepsy from birth or early age on, but might disappear while the patient is growing up. Other causes tend to occur more often in elderly people, causing again an increase in prevalence in this population. Therefore these patient groups are potential interesting targets for monitoring applications as will be discussed in Section 2.4.

Video-Electroencephalography (vEEG) monitoring is the gold standard for epilepsy diagnosis, for which the patient visits the hospital for a short or long-

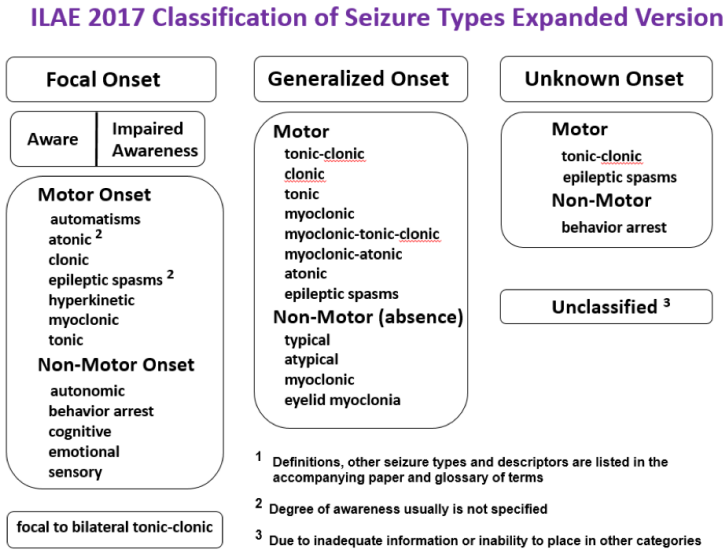


Figure 2.2: New ILAE seizure classification guidelines 2017. Taken from [61].

term monitoring. This monitoring allows to evaluate the patient both ictal (during a seizure) and interictal (between different seizures) both in terms of brain activity (visible on the electroencephalogram (EEG)) and behavior (visible on video). Both modalities thus give great input in the diagnosis, even when no seizures occur during the monitoring application (e.g. interictal spike analysis). This information can then be extended by including also information from the neurological history of the patient, neuroimaging tests (i.e. MRI scan) and physical examination of the patient.

2.1.2 Seizure classification

ILAE 2017

A wide variation of epilepsy and seizure types exists. The most common way of seizure type classification is defined by the International League Against Epilepsy (ILAE) (version 2017), which is summarized in Figure 2.2.

This classification guideline is divided in three big groups, which are defined based on the onset of the seizure. In focal onset seizures, the seizures originate within networks limited to one hemisphere in the brain [61, 158]. Generalized

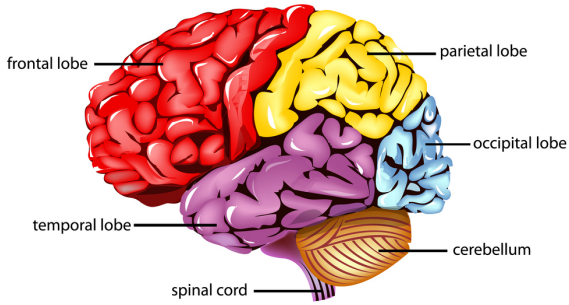


Figure 2.3: Overview of the different lobes in the brain. Taken from [134].

onset seizures however manifest in bilaterally distributed networks, affecting both hemispheres. In case the onset is not known (i.e. not clearly visible during vEEG monitoring), the seizure onset is classified as unknown. In focal seizures, it is also mentioned whether there was impaired awareness or not ('focal aware' seizures) during the seizures.

For all three seizure onset types, there is a further distinction between whether there is ictal motor response or not. Which type of motor or non-motor behavior occurs, depends on the type of epilepsy. For focal seizures, the particular region that is affected during the seizure has a big impact on which behavior can be found during a seizure (see Figure 2.3). Temporal lobe seizures for example tend to more likely lead to autonomic changes, whereas frontal lobe seizures often lead to hyperkinetic seizures.

The most often used types of motoric and non-motoric behaviors that are relevant for this thesis are:

- Myoclonic seizures: Sudden short jerk-like movements ($\leq$0.5 s) [82]. It is a twitch-like contraction which typically includes limbs and shoulders. Awareness is mostly not impaired in myoclonic seizures, and can potentially occur in clusters.
- Tonic: Increased tension in one or multiple muscles for at least 2 s, often resulting in a change of posture. Mostly awareness is preserved in tonic seizures.
- Clonic: Series of myoclonic contractions of agonist and antagonist muscles. Contractions often occur with a frequency between 0.2-5 per second. Typically the whole body is involved in these seizures, and movements can

not be stopped by restraining or repositioning the limbs. Usually there is impaired awareness in clonic seizures.

- Tonic-clonic: Tonic-clonic seizures (also known as grand mal seizures), are a combination of a tonic and a clonic seizure. This means that the involved muscles first stiffen (tonic phase), and afterwards start to jerk (clonic phase). Patients typically lose consciousness during tonic-clonic seizures.
- Hyperkinetic: Seizures with typically violent and repetitive movements. The movements are often normal movements, but abnormal given the situation (e.g. pedaling in bed).
- Automatisms: Seizures are impaired by performing one or more brief unconscious behaviors. Some example of automatisms are lip smacking, swallowing, chewing and rubbing fingers/hands. Automatisms often occur in seizures originating from the temporal lobe.

Old ILAE classification before 2017

As can be seen above, the new ILAE classification guidelines were suggested in 2017, during the thesis [30]. The majority of the studies in this thesis are therefore described with respect to the old ILAE classification guidelines. Without going into detail, the majority of definitions stay the same between both classification standards within the scope of this thesis except for the renaming of some definitions [60]. To summarize, following old terms can be translated into the following new terms in the rest of this thesis (Chapter 3, 5 and 6 apply the old terminology):

- Partial seizures: Focal seizures
- Simple partial seizures: Focal aware seizures
- Complex partial seizures: Focal impaired awareness seizures
- Secondary generalized tonic-clonic seizures: Focal to bilateral tonic-clonic seizures

2.1.3 Treatment

The required treatment that is given to a patient depends on the type of seizure and epilepsy syndrome that are diagnosed. Unfortunately, no cure exists for

epilepsy, and the available treatments mainly focus on the suppression of the seizures. Certain treatments give better responses than other for certain types and should therefore be tried out first.

Anti-epileptic drugs (AED) allow to suppress the seizures in around 70% of the patients [25]. Multiple trials with different AEDs can be tested (including mixtures of AEDs) to try to find the ideal dosage and mixture for each individual patient. If however no such optimal mixture of AEDs leads to sufficient seizure suppression, other treatments should be evaluated. If a patient does not become (and stay) seizure-free after at least 2 AED trials, the patient is called to have refractory epilepsy.

One of the alternatives to AEDs is surgery. In this case, it will be evaluated whether the removal of a particular part of the brain could lead to seizure-freedom. This part of the brain that causes the seizures is called the epileptogenic zone. The patient is then monitored for several days in the hospital for presurgical evaluation to further analyze the patient's ictal and interictal brain behavior. Then it can be decided whether surgery could be beneficial for the patient or not. This is the case if the epileptogenic zone is sufficiently small and not interfering with any other brain function.

For patients for whom the above-mentioned treatments were not helpful, also the ketogenic diet could be useful for seizure frequency suppression [131]. The ketogenic diet mainly consists of increased intake of fat and lowered consumption of carbohydrates. The increased intake of fat leads to more fat being used for energy consumption, which leads to an increased production of ketones. More ketones often lead to an improved seizure control [204]. Also the modified Atkins diet could lead to better suppression of seizures in some patients [110, 120].

A relatively new treatment is vagus nerve stimulation (VNS). In VNS, a device is placed in the left chest, and is connected to the left vagus nerve in the neck. The device gives mild periodic pulses to the vagus nerve, causing different areas of the brain to be activated. The periodicity and the strength of the pulses can be adjusted by the clinician in order to set it to an optimal value for each individual patient. Different studies already have investigated the impact of VNS, resulting in at least 50% reduction in seizure frequency in 50-60% of the patients [54, 56]. Remaining seizures also proved to be shorter on average when using VNS compared to prior VNS usage [84].

Deep brain stimulation (DBS) is another possible treatment for refractory epilepsy patients. Unlike VNS, which sends pulses to the brain through the vagus nerve, DBS directly stimulates the brain in response to specific electrographic behavior. DBS is for example triggered if a specific pattern (i.e. epileptic waves) is detected in the patient's electrocorticogram (ECoG). Different studies already

evaluated the usage of DBS, giving similar results as the ones discussed for VNS [62, 129, 171].

2.1.4 Mortality in epilepsy

Epilepsy patients have an increased mortality risk compared to the normal population of a factor of 2 to 3 [68]. These can be mainly related to the following categories [182]:

- Death caused by the underlying cause of the disease (e.g. brain tumor).
- Death caused by co-morbidities such as co-existent neurological compromise.
- Death caused by the disease directly, including status epilepticus and accidents during/after seizure occurrence or suicide caused by depression.
- Death with unknown cause, called sudden unexplained death in epilepsy (SUDEP).

A patient is in status epilepticus if the patient is having a seizure of at least 5 minutes duration, or a series of seizures without being able to recover from these seizures. In status epilepticus, mechanisms that are responsible for seizure termination fail, causing the patient not to be able to recover from them. It is a possible life-threatening situation if the duration of the condition takes too long.

The unexpected character of SUDEP causes a lot of anxiety in epilepsy patients and their relatives, as this could happen at any given time, with reported incidence rates being as high as 2.3 per 1000 patient-years. For refractory epilepsy patients, this incidence rate can even be as high as 9.3 per 1000 patient-years [167].

In 33%-100% of SUDEP cases, death was preceded by (mostly convulsive) seizures. Different possible risk factors for SUDEP were identified in the literature, including amongst other generalized tonic-clonic seizures, nocturnal seizures and symptomatic epilepsy with partial seizures [87, 155]. It was also found that often cardiac and respiratory dysfunctions were proposed as potential pathophysiological mechanisms for SUDEP, including central respiratory depression, higher baseline vagal tones, tachycardia and tacypnea [101, 132, 141, 147, 155, 172].

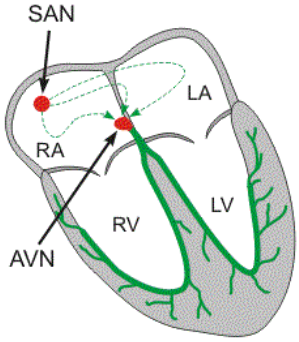


Figure 2.4: Overview of the physiology of the heart. RA: right atrium, LA: left atrium, RV: right ventricular, LV: left ventricular, SAN: sinoatrial node, AVN: atrioventricular node. Taken from [108].

2.2 The heart

2.2.1 Physiology of the heart

The physiology of the heart is illustrated in Figure 2.4. The heart can be divided in two pumps, one left and one right. They provide enough pressure to pump oxygenated blood throughout the entire body. Each chamber has an atrium and a ventricle. Deoxygenated blood from the rest of the body enters the heart in the right atrium, which pumps it towards the right ventricle. The right ventricle forwards the blood to the lungs in order to get oxygenated, which is then send towards the left atrium. The left atrium sends the blood to the left ventricle, which pumps the blood to the rest of the body. This heart cycle can be divided in two parts: the diastole phase (the filling of the heart with blood) and the systole phase (the contraction of the heart to pump the blood outside of the heart).

This heart cycle is monitored by the sinoatrial node (SAN), which initializes the above mentioned heart cycle. It gives an electrical impulse towards the atrioventricular node (AVN), and is later forwarded towards the heart ventricles. The rate at which the heart cycles occur (thus the heart rate) is determined by the SAN, which can be seen as the pacemaker of the heart. The SAN contains cells that allow self-excitation, meaning that it allows for an intrinsic firing rate. It is however normally constantly modified by the autonomic nervous system (see Section 2.2.3).

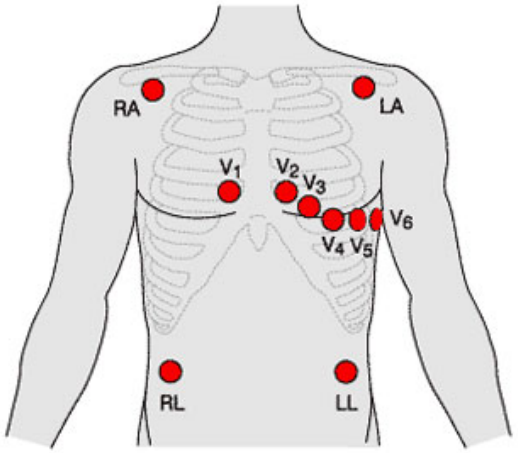


Figure 2.5: Placement of sensors for a standard 12-lead ECG. Taken from [203].

2.2.2 Electrocardiography

The electrical impulses sent by the SAN travel through the heart due to the electrical characteristics of the heart cells. This electrical activity in the heart can be captured by using bipolar sensors on the skin which measure the voltage difference between the different nodes. The corresponding signal is called the electrocardiogram (ECG). A different sensor placement allows to get a different cardiac information. A 12-lead ECG sensor placement is most commonly used for clinical practice, as it allows to give a detailed overview of the cardiac activity (see Figure 2.5). A lead is defined as the signal found by evaluating the difference between two electrodes. For most topics, not that much detail is however needed, and sufficient information can be obtained by using only one lead, which allows an easier acquisition. Often, the so-called lead II (= LL - RA in Figure 2.5) is chosen as it contains most information for medical diagnosis. In this thesis, only lead II ECG signals are used.

Figure 2.6 shows a short segment of a lead II ECG signal. Due to the nearly periodic nature of a healthy heart, the signal consists of quasi-periodic repetition of different characteristic waves, which can all be linked to the original actions made in the heart. Each heart cycle starts with an electrical stimulus from the SAN, which propagates through the atria, causing contraction and depolarization (P wave). When the stimulus reaches the AVN, it is delayed for less than 100 ms, causing an iso-electric segment in the ECG (PQ interval). Next, the impulse

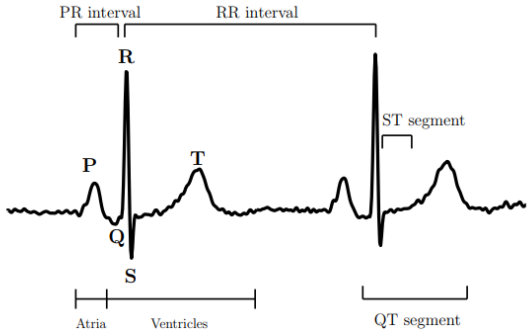


Figure 2.6: Example of an ECG segment to illustrate the different ECG morphologies. Taken from [194].

reaches the ventricles, leading to a strong contraction (QRS complex). The depolarization of the ventricles then again leads to an iso-electric segment (ST interval), after which the ventricles relax and repolarize (T wave). Because the ECG contains information on the (de)polarization of all atria and ventricles, it is ideal to detect abnormalities in the heart.

The heart rate can be extracted from the ECG by computing the differences between subsequent R peaks (RR interval). The immediate heart rate (HR) is then defined as

$$HR = \frac{60}{RR} \tag{2.1}$$

with RR the RR-interval expressed in seconds. A wide variety of possible solutions using different techniques has already been proposed in the literature to automatically extract the heart rate from the ECG [109]. These algorithms automatically detect the R peaks in the ECG signal. Most of these algorithms are based on filtering [2, 143] and wavelet methods [116]. The resulting signal of instantaneous heart rate measurements is defined as the tachogram.

2.2.3 The autonomic nervous system

The autonomic nervous system (ANS) is part of the central nervous system (CNS). It is responsible for the regulation of the heart and respiratory rate amongst others. It quickly modulates these functions and several organs in order to maintain homeostatis if changes occur during external stimuli. The ANS is regulated by multiple centers in the brain, including the hypothalamus, brain stem, spinal cord and cerebral cortex. The ANS can be further divided in two

parts: the sympathetic nervous system and the parasympathetic nervous system. The sympathetic nervous system reacts to stressful situations which require immediate response. An increased sympathetic activation leads to an increased rhythmicity in the heart rate. The parasympathetic nervous system on the other hand helps to recover after an event by slowing down the respiratory and heart rate. An increase in the parasympathetic system leads to a decrease in heart rate. Both systems are normally in balance (the so-called sympathovagal balance), but can get out of balance after an autonomic dysfunction. Patients with such autonomic dysfunctions might have problems to overcome stressful situations, which could lead to life-threatening situations.

2.2.4 Heart rate variability

An option to study the effects of the ANS on the heart rate is by doing heart rate variability (HRV) analysis. In such an analysis, the changes of the heart rate in time are analyzed, which can be done in both time and frequency domain [23]. In the frequency domain of the RR-interval series, it has been shown that in the low-frequency domain (0.04 Hz-0.15 Hz) both sympathetic and parasympathetic activity can be found, whereas mainly parasympathetic activity is found in the high-frequency domain (0.15 Hz-0.4 Hz) [13]. The impact of the respiration on the heart rate, called respiratory sinus arrhythmia (RSA), is also mostly found in the high-frequency domain of the spectrum. The frequency of RSA can, however, decrease below 0.15 Hz, causing an alteration of the low-frequency spectrum.

Typical examples of time-domain HRV features are the average heart rate and standard deviation of successive differences (SDSD), but also more advanced time-domain HRV features can be extracted (see Appendix A for the definition of the used HRV features in this thesis). Time-domain features are often relatively easy to extract and interpret, and already give a good insight on the person's general health and heart functioning. This can be important for the patient, as some of these parameters for example are powerful markers of heart failure [111]. Several studies also showed that these HRV parameters are also altered in several diseases such as apnea and epilepsy (see Section 2.3), indicating influence of the disease on the functioning of the heart [98, 146, 192]. It therefore also allows HRV parameters to be used for automated detection of for example epileptic seizures and apneic events [37, 44, 99].

Heart rate characteristics strongly vary between different subjects. Subject age, activity, lifestyle, drug intake and diseases for example can affect these heart rate characteristics [27]. It can also fluctuate throughout the day (the

so-called circadian rhythm), so that the nocturnal short-term HRV parameters for example is not identical to the awake short-term HRV parameters [123].

2.3 Epilepsy and the autonomic nervous system

As shortly mentioned in the previous section, there is an influence of epilepsy on the heart rate behavior. Literature already showed that in 73-92% of seizures, ictal heart rate increases can be found, leading to tachycardia (heart rate higher than 100 bpm) in around 73-90% of the seizures [50, 98, 114, 140, 166, 212]. These percentages are found to be higher in patients with temporal lobe epilepsy and generalized tonic-clonic seizures. Often the ictal heart rate increases are also stronger in amplitude for these types of epilepsy [96, 114, 138, 140, 169, 170]. Focal impaired awareness seizures result on average in stronger heart rate increases than focal aware seizures originating from the temporal lobe [83, 170]. In rare events (1.4% of seizures), ictal bradycardia (heart rate lower than 50 bpm) is found, often following an ictal heart rate increase [114].

These increases in heart rate can mostly be interpreted as strong activations in sympathetic activity, sometimes accompanied by an increase in parasympathetic activity [49]. These changes lead to a decreased high-frequency and increased low-frequency spectrum of the RR interval sequence. The ictal heart rate changes typically follow a specific pattern, with the most common found pattern shown in Figure 2.7. The most typical pattern consists mainly of three phases, which include a phase of linear heart rate increase, a phase of heart rate stability and an exponential decrease in heart rate [122]. In [114], the increase is rather described as a parabolic increase rather than a linear increase. Several different variations of the ictal heart rate patterns can however be found as illustrated in Figure 2.8 [142, 166]. Despite the fact that often a large patient-dependency of this ictal heart rate pattern exists, there is still variation between the ictal heart rate patterns of seizures within a patient in both shape and amplitude [48, 114, 151, 166].

These ictal heart rate changes can often be seen as changes caused by the ANS. It is believed that the ANS is triggered if part of the brain that helps to regulate the ANS (e.g. the amygdala) is affected by the seizure during seizure propagation [57, 159]. This is for example illustrated by the fact that temporal lobe seizures more frequently lead to a change in the ANS compared to other types of focal seizures [138]. The temporal lobe is located closely to the parts in the brain that regulate the ANS, thus make it more likely that seizures located in this lobe will also alter the ANS.

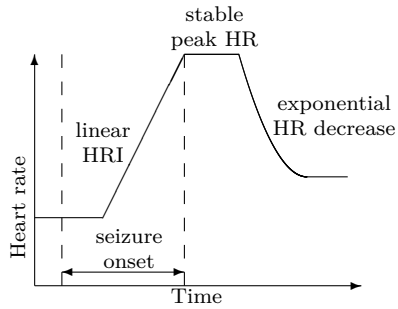


Figure 2.7: Illustration of the most found ictal heart rate pattern. The dashed lines indicate the interval where typically the seizure onset is located with respect to this heart rate pattern.

Often, the above-mentioned changes in the ANS are not only detected during a seizure, but even before the seizure onset as found in the (scalp) EEG. This can occur if a seizure is manifesting deep in the brain, making it not yet clearly visible on scalp EEG, but already affecting the ANS. A review study on the pre-ictal heart rate changes was presented in [21]. It was found that on average, 36% of seizures had pre-ictal heart rate increases, whereas this was 44% for patients with temporal lobe epilepsy. The median onset of pre-ictal heart rate increases was 10.7 s prior to the seizure onset. Similar as for ictal heart rate changes, the pre-ictal changes are more pronounced for temporal lobe seizures and generalized tonic-clonic seizures [21, 138].

Not only simple increases and decreases in the heart rate can be found during seizures. Around 40% of patients have ictal ECG abnormalities such as asystoles, atrial premature depolarizations and atrial fibrillation [133, 212]. Also other morphological changes such as QT lengthening, ST depression, QRS morphology changes and T-wave inversion can occur during a certain amount of epileptic seizures [20, 140, 197].

Also strong changes in the sympathetic system can occasionally be found post-ictally. These are often found as periodic alterations of high amplitude (Mayer waves) in the tachogram, which follow the ictal heart rate pattern discussed above [4, 124] (see Figure 2.8(a)). These postictal autonomic changes have shown to be correlated to postictal EEG suppression [149]. Prolonged EEG suppression could indicate an increased chance of later SUDEP, so it is interesting to keep track of these post-ictal changes, as they allow to indicate how well a patient recovers from a seizure [115].

The ictal effects on ANS by seizures are not limited to ictal changes in the heart rate. Also other symptoms can rise due to changes in the ANS, including

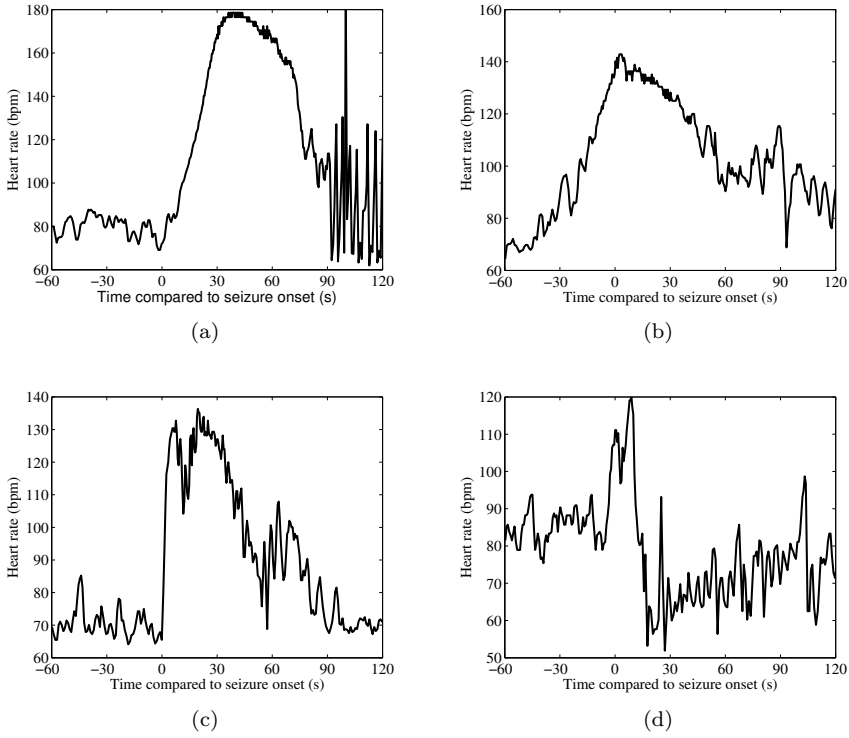


Figure 2.8: Examples of four ictal heart rate patterns from four different patients. (a) Pattern including monotonically increasing heart rate, heart rate stabilization and fast decreasing heart rate. An example of strong post-ictal heart rate changes is illustrated in this figure. (b) A similar pattern as in (a), but with high frequency alterations during the heart rate increase and a much less steeper decrease in heart rate. The heart rate increase also starts pre-ictally, reaching the maximal heart rate around seizure onset. (c) The heart rate increases linearly towards a peak heart rate, then slightly decreases and increases again towards the same peak heart rate. (d) Similar pattern as in (c), but the first reached heart rate peak is much lower than the maximal heart rate peak achieved in the second increase. The heart rate decreases nearly to bradycardia after the heart rate increase.

respiration symptoms (e.g. hyperventilation), gastrointestinal symptoms (e.g. epigastric auras) and urinary problems [8].

2.4 Long-term monitoring of epilepsy patients

The quality of life (QoL) of refractory epilepsy patients is low due to the unpredictable nature of the seizures they continue to experience. They fear to leave the house in case a seizure would occur, which leaves them very vulnerable if no one knows where the patient is, leading to dangerous situations. Not only the QoL of the patients is strongly affected by epilepsy, also the patients' relatives and friends are strongly affected by the disease. Parents of pediatric patients for example are constantly afraid their child will get a seizure, or even worse, post-ictal complications including for example SUDEP. Monitoring of these epilepsy patients in their home environment could increase the QoL of these people in different ways as will be explained in this section.

Monitoring of epilepsy patients in their home environment can also be helpful for clinicians, allowing to get more accurate and objective information from the patient's current health. This can for example be interesting to study whether a new drug dosage works well for the patient or not. This information could then lead to an increased understanding of the patient's disease, and lead to better treatment selection. It also allows them to get more information on how these seizures occur outside the hospital, as differences can occur between seizures experienced in the hospital and real-life outdoor seizures [73].

This section discusses some potential technical solutions and the underlying algorithms to allow home monitoring of epilepsy patients, both already available and potential future applications. VNS and DBS were already discussed in Section 2.1.3 as they have a strong impact on the occurrence and duration of the seizures directly, rather than mainly being available for improving the QoL of patients and relatives and improved treatment selection.

2.4.1 Automated seizure warning system

The most investigated option in the literature is the use of an automated seizure warning system. In such a system, a seizure is detected automatically in real-time by a seizure detection algorithm. If such a seizure is detected, an alarm is sent to the patient's relatives, friends or caregivers in order to warn them about this event. Such an alarm can be given by either sending a text message, a telephone call or sending alerts through a smartphone app [22]. In some cases,

the alarm also gives the GPS coordinates of the patient to indicate where the patient is located. This way, the person who received the alert can come over to the patient to check whether the patient is fully recovered and do the required actions to help and comfort the patient (post-)ictally. This way, the chance of injuries can be reduced if the caregiver comes sufficiently fast to the patient.

The crucial part of such devices are in this case the algorithms that automatically detect the seizures. Ideally, these algorithms

- detect as many seizures as possible (high sensitivity),
- give as few false alarms as possible (low false alarm rate) and
- detect the seizures as fast as possible (low detection delay).

In automated seizure detection, one typically needs to find a good balance between a high sensitivity (percentage of detected seizures) and low false alarm rate (FAR, number of false alarms, expressed in number of false alarms per hour, night or day). Improving the sensitivity of an algorithm often also results in a larger FAR. A good balance between both metrics depends on the types of seizures that want to be detected by the system. For the strong and dangerous types of seizures (e.g. tonic-clonic seizures), often at least ±80-90% is required [162]. These seizures are often more easily detected by automated seizure detection than more subtle seizures. For clinically less dangerous seizures (i.e. focal seizures with automatisms), a lower sensitivity of 70-80% is still acceptable, as missing this type of seizure is less critical.

A fast detection of seizures allows the caregivers to help the patient quicker, reducing the chance of injuries or further development of the seizure. For the seizure detection algorithms, however, there is again a trade-off that needs to be made. Real-time algorithms are typically more accurate (in both sensitivity and FAR) if they are allotted more time to make the decision. The amount of acceptable detection delay depends on the envisioned application. If the caregiver is for example within walking distance (i.e. in the same house as the patient), a faster detection allows the caregiver to come to the patient during a seizure, avoiding potential (further) injuries [80]. If the caregiver is further away from the patient, an extra delay of 10-30 s compared to the traveling time is negligible as the caregiver will probably only arrive postictally anyway. In this last scenario, the effect of false alarms is also larger then in the first solution, so a higher accuracy might be required for such a system. Often systems only focus on detection of seizures with a certain minimal duration, varying between 10 and 60 s in the literature [80, 126, 139, 181].

For real-time applications using wearable devices, it is also important to keep in mind that these devices have a limited computational power and battery lifetime.

It is important for usability that this battery lifetime is sufficiently long, so that they do not have to recharge the device too often [144]. Algorithms requiring a limited computational power also allow smaller devices for the monitoring application, which is important in terms of wearability of the system. Therefore, the algorithms for seizure detection should not be made too complex in order to be usable in practice. An overview of the different requirements that have to been taken into account for designing a seizure warning system is described in further detail in [179].

The required seizure detection performance depends on the end users of the device, this could be for example the patients' relatives or residential care givers. For overall seizure detection performance, typically at least 70-80% sensitivity is required, and a positive predictive value of close to 50% is asked for [162, 186]. Typically, a maximal detection delay of 30-60s is preferred. Stronger demands are typically asked for patients with stronger types of seizures (e.g. generalized tonic-clonic seizures). However, these strong demands are rarely accomplished in current available devices except for devices limited to the detection of a subgroup of a very restricted seizure type, leading to skepticism in a large group of the patient population [22].

Automated detection of epileptic seizures is based on the interpretation of one or more (biomedical) signals. Which type of seizures can be detected depends on the chosen signal. Some devices for acquiring certain biomedical signals are also more comfortable for wearing in a long-term application than others, which is also important to consider for this application. The following section gives an overview of the possible modalities that can be used for automated seizure detection. More detailed overviews on the available devices and proposed algorithms can also be found in [15, 103, 152, 154, 177, 183, 184].

2.4.2 Modalities for automated seizure detection

EEG

Most research in automated seizure detection in the literature went to the analysis of EEG signals [1, 176], with the intention of being used in a hospital environment during monitoring for epilepsy diagnosis or pre-surgical evaluation. EEG typically leads to relatively good performance over all seizure types as it analyzes the signals that are also used by the expert clinicians for diagnosis and annotation. Due to the large amount of signals over the different parts of the brain, it allows to include both spatial and temporal information of the brain, leading to a good seizure detection performance (see Figure 2.9) [93]. Using full EEG in a home environment for long-term monitoring is however not comfortable

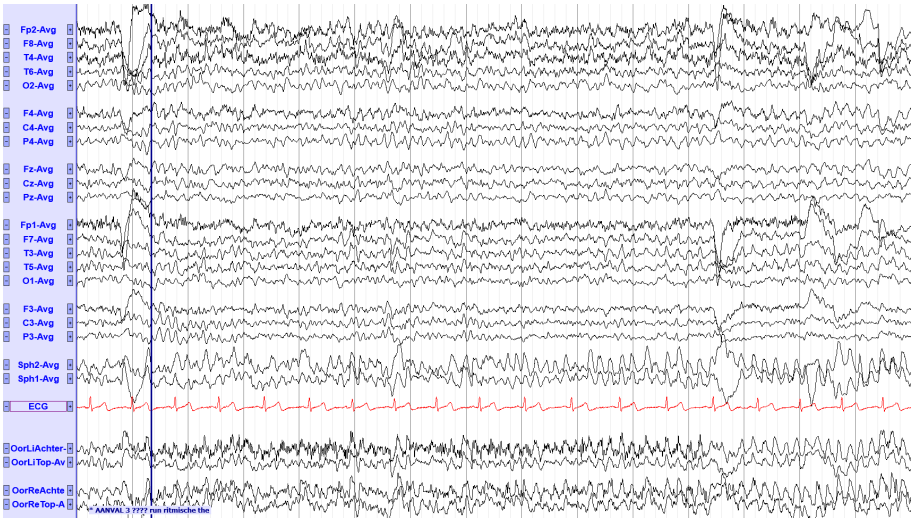


Figure 2.9: Example of EEG signals during a seizure.

and stigmatizing for the patient and is therefore nearly never suggested as a solution for real-time detection in a home environment [12, 22, 91, 162].

To resolve the comfort issue, more wearable variations of a full EEG system are proposed. Mobile EEG are more wearable EEG caps, typically containing only a limited amount of EEG electrodes to increase the wearability for long-term monitoring. Mobile EEG however still remains less wearable than most of the alternatives for seizure detection, and most patients still experience it to be stigmatizing to wear it in public.

A less stigmatizing solution is to use (in)ear-EEG. In this case, only EEG electrodes are placed either around the ear, or EEG signals are extracted from within the ear [78, 210, 211]. The device then looks similar to a hearing-aid device, which is less stigmatizing than the other EEG solutions. The drawback is that only the EEG signals around the ear can be acquired, making it only useful for generalized and focal temporal lobe seizures [190].

In the above mentioned options, scalp EEG signals are used, for which EEG electrodes are placed on the scalp of the patients. In some cases, however, no ictal EEG changes can be found as the seizures originate too deep in the brain, not showing any changes on the surface EEG. Intracranial EEG (also called electrocorticography, ECoG) electrodes can be placed within the brain in order to get more accurate brain activity information. These electrodes are often used during pre-chirurgical evaluation in the hospital and as part of DBS (see Section

2.1.3). Placement is however very invasive, and usage of such electrodes only for automated seizure warning seems too much effort compared to the possible gain (compared to for example VNS and DBS and possible seizure prediction).

ECG & PPG

As discussed in Section 2.3, (pre-)ictal changes can be found in the heart rate in most seizures. Especially in patients with focal impaired awareness seizures originating from the temporal lobe and generalized tonic-clonic seizures have a high prevalence and amplitude in ictal heart rate changes. Heart rate-based seizure detection will therefore give the best performance for these types of seizures. The detection of non-motor seizure is currently not investigated thoroughly in literature, where focus mostly went to the detection of generalized tonic-clonic seizures [16, 187]. For these seizures, heart rate-based seizure detection is of much added value compared to most alternatives, as it also allows to detect a big portion of non-motoric seizures (except absence seizures).

Not only the ECG, but also photoplethysmography (PPG) can be used for heart rate extraction (and thus also heart rate-based seizure detection). PPG monitors the changes in blood volume in tissue (e.g. finger, ear, arm, ...) by using reflecting lights. The obtained PPG signal also has periodic characteristics due to the occurrence of the heart beats (see Figure 2.10). Therefore it can also be used for automated heart rate extraction, and thus also for heart rate-based seizure detection. Due to the low-frequency characteristics of PPG, it is less sensitive to high frequency noise compared to ECG. PPG is however more sensitive to motion artifacts caused by bad contacts to the skin surface (see for example Figure 2.10). A good PPG sensor placement is crucial in this case in order to get a robust PPG signal and heart rate extraction [180]. Due to the low-frequency PPG characteristics, a portion of high-frequency heart rate variability can be lost, which could have a negative effect on the performance of the used seizure detection algorithm. PPG acquisition is less obtrusively than ECG, but often still results in less robust heart rate extraction than ECG [191].

Although these algorithms are highly usable for a more wide variety of seizures than other modalities, the biggest downside of heart rate-based seizure detection is the limited differentiability between epileptic and non-epileptic heart rate changes during the day [48]. Therefore, it is difficult to get a good performance using these methods, often leading to too high false alarm rates to be actually used unimodally in practice. Heart rate-based tonic-clonic seizure detection leads to a lower performance than algorithms using for example accelerometers and electromyogram in the literature. Also the strong variability in ictal heart rate changes between patients causes an extra problem if generic (patient-

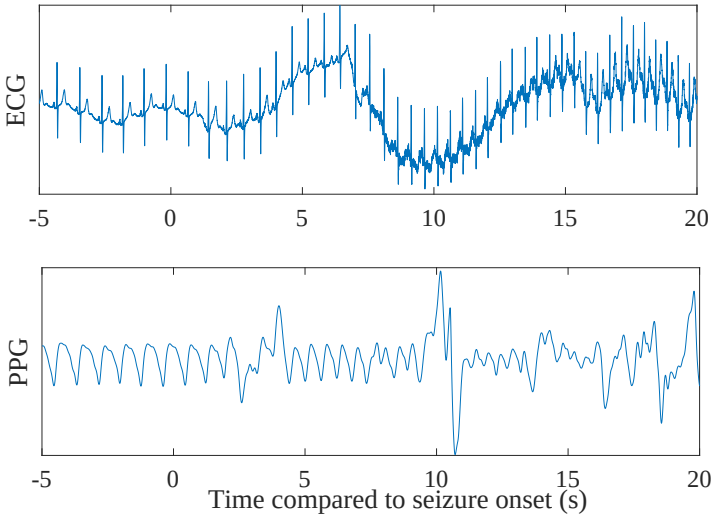


Figure 2.10: Example of ictal ECG and PPG activity acquired synchronously. PPG is obtained in this example from the wrist. Epileptic activity can be recognized by the increased frequency of the heart cycles.

independent) seizure detection algorithms are required. Due to the usability of heart rate features in a wide variety of seizure types, it is however an ideal modality to add in a multimodal setup.

More information can be extracted from the ECG than only the heart rate. Also morphological changes and ECG derived respiration (EDR) can be used for automated seizure detection to improve the performance of ECG-based seizure detection [197, 209].

The above-mentioned methods can not only be used as part of an automated seizure warning system. Algorithms based on ECG or heart rate also allow to make a closed-loop VNS system, which not only periodically activates the vagus nerve, but also automatically whenever a seizure is detected [19, 201].

Accelerometry

Accelerometers measure the acceleration of the device along three axes. They are ideal for the detection of strong (atypical) movements, as for example can be found in tonic-clonic, clonic and hyperkinetic focal seizures. These movements

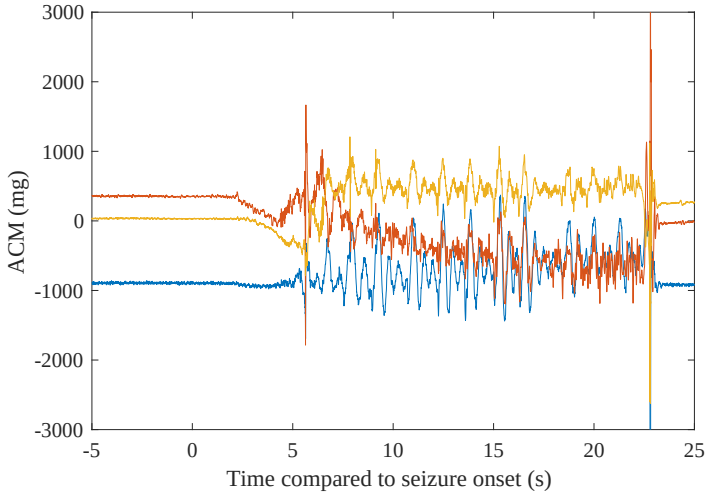


Figure 2.11: Example of changes in the accelerometry along the different axes during a generalized tonic-clonic seizure. The accelerometer signal is acquired from the wrist in this example. The clonic phase of the seizure starts around 5 s after seizure onset.

are typically strong in amplitude and often contain rhythmic patterns (see Figure 2.11 for an example) [9, 125, 126, 145]. Myoclonic seizures can also be detected, but the performance is rather low for these type of seizures due to their short duration [136]. Accelerometers are typically attached to the wrists of the patient, but are sometimes also placed on the upper arms, the ankles or the torso of the patient.

Electromyography

The electromyogram (EMG) measures the muscle tension in function of time. It is therefore suited for the detection of seizures including muscle stiffening, such as tonic and tonic-clonic seizures (see Figure 2.12) [33]. Typically either the deltoid or biceps muscle is analyzed for automated seizure detection in literature, resulting in a device that is worn on the upper arm or near the shoulder [31, 81, 112, 126, 174]. It is of added value compared to accelerometers for tonic-clonic seizure detection as it detects the first part of the seizure (i.e. the tonic phase), leading to a faster detection than accelerometry-based seizure

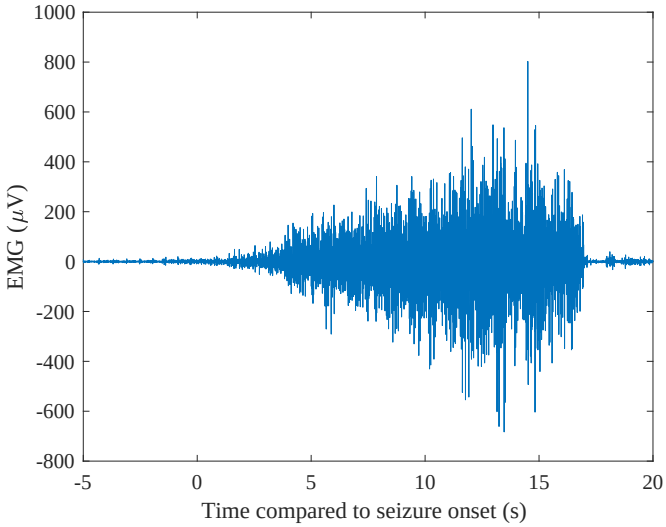


Figure 2.12: Example of changes in a biceps EMG signal during a generalized tonic seizure.

detection. The required equipment for obtaining EMG signals is currently however typically more bulky than accelerometer sensors [10, 11, 126].

Electrodermal activity

Electrodermal activity (EDA) measures the conductance of the skin (also called galvanic skin response (GSR)). EDA can for example detect if a person is sweating, which is an action controlled by the sympathetic nervous system. Therefore, it can also be used for the detection of seizures as it also allows to detect sympathetic changes during epileptic seizures. Changes in EDA are however much clearer in generalized tonic-clonic seizures compared to focal impaired awareness seizures, and this difference is much clearer than can be found in the ictal heart rate changes.

It could therefore be used for similar patient groups as the ones discussed for ECG/PPG, but sympathetic changes in EDA typically occur rather post-ictally than (pre-)ictally. Therefore they might lead to a delayed detection compared to ECG/PPG-based seizure detection algorithms, but are easier differentiable from non-epileptic activity than in ECG/PPG, allowing a better accuracy performance. This is due to the fact that EDA is assumed not to be affected by

the parasympathetic system, which allows to fully analyze the ictal sympathetic changes without interference of other factors [139, 150]. Therefore they might be preferred over ECG/PPG in an offline procedure such as a seizure diary (see Section 2.4.3). EDA can be obtained by using for example a wristband [55].

Audio

Audio signals can also be used for seizure detection [6, 24, 65]. The algorithms often try to detect noises originating from screaming, falling objects, bed movement or autonomic laughing or weeping. Audio signals are cheap and very easy to use as they do not have to be applied on the patient, which especially makes it interesting to use in pediatric patients. The biggest downside of audio signals is that they result in relative poor performance due to noises from the surroundings interfering with the audio signal. Also only a limited amount of patients make specific sounds during epileptic seizures, but a high seizure detection performance can be obtained for these patients [6]. Also privacy issues can be a problem for some patients.

Video

Automated video-based seizure detection has been mainly of interest in motor seizure detection. Video is contactless, but implicates that the patient is restricted in his/her movements to wherever the video can see the patient, causing it often being used for nocturnal monitoring [34, 70, 104]. Video analysis also becomes difficult when something or somewhere is blocking the camera, for example the patient is lying under a blanket in bed. Thermal or thermographic cameras and optical flow analysis can however solve this issue, and can also allow to extract the heart rate and estimation from the video images [69, 70, 130]. Thermal cameras are however costly. Also privacy can also be a problem for seizure detection using video images.

Other modalities

Also other modalities can be used for seizure detection, but are typically not differentiable enough to allow unimodal seizure detection. These include amongst others blood pressure (BP), respiration, gyroscope, pressure sensors in the bed, temperature, eye movements and oxygen saturation [9, 26, 28, 55, 65, 83, 97]. They can however be of added value in multimodal systems.

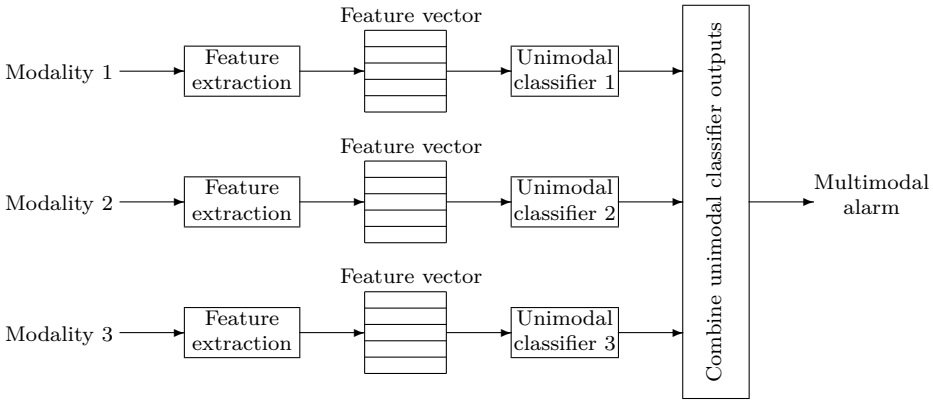
Multimodal systems

Multimodal systems are systems that make use of multiple sensors for seizure detection. Two main options are possible for multimodal seizure detection: late integration and early integration. In late integration, first a decision is made per used modality/sensor (see Figure 2.13(a)). These unimodal decisions are then combined in order to obtain one multimodal output, which then decides whether an alarm is sent or not. In the literature, typically straight-forward multimodal combinations are used such as the AND-operator, OR-operator or majority voting [66, 76, 181]. The AND-operator and majority voting require most or all modalities to be active during a seizure. It is particularly useful if the system focuses on the detection of a particular seizure type in which all or most acquired modalities are assumed to be active, but lacks generalization to other seizure types. The OR-operator allows more generalization, so that it can be used in a system for the detection of multiple seizure types. The unimodal decision steps should however be trained stricter (i.e. lead to a lower FAR) than the ones used when the AND-operator is used in order to avoid a too large FAR in this multimodal setup.

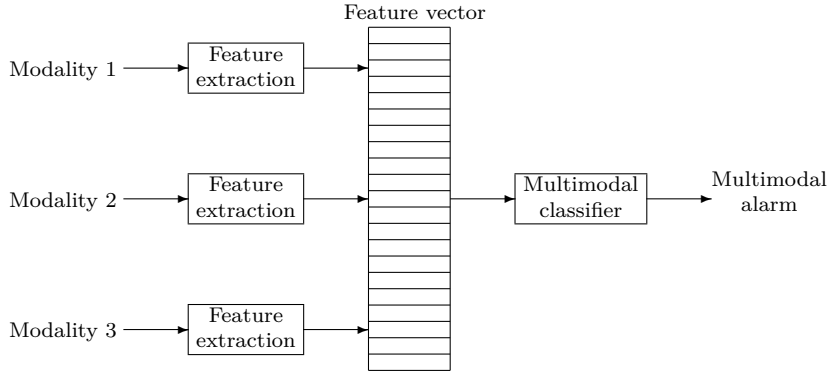
In early integration (see Figure 2.13(b)), there are no intermediate unimodal decision steps. The features are extracted from the different modalities and combined in one general feature vector. The multimodal decision is then based on these features directly. This is of added value if interaction exists between features from the different modalities, which can be used for improved seizure detection performance. One drawback of this approach is the difference in activation times that can often be found in different modalities. Accelerometers and EMG are for example both active during a generalized tonic-clonic seizure, but often at different times within the seizure. This lack of modality synchronization makes it often very challenging to implement this methodology in practice for most modality combinations. Therefore mostly late integration is used in the literature.

Ideally, sensors from different modalities are combined in such a multimodal setup. This is done in order to find a minimal overlap in causes of false alarms from the different modalities. Every modality has its typical type of false alarms, but they often do not overlap in different modalities. False alarms in accelerometry are often caused by periodic movements, false alarms in EEG by chewing artifacts and false alarms in heart rate by stress or arousals. Adding different modalities therefore typically results in drastic decreases in false alarms if the above mentioned AND-operator combination is used.

The most common example is the detection of tonic-clonic seizure, which can be detected with accelerometers, EMG, heart rate, EDA and EEG, and



(a) Illustration of a late integration approach.



(b) Illustration of an early integration approach.

Figure 2.13: Comparison of the two typical options for multimodal seizure detection. (a) Illustration of late integration, in which feature extraction and classification occur per individual modality. Afterwards, the unimodal results are combined into 1 decision output. (b) Illustration of early integration, in which the extracted features from different modalities are combined in one feature vector. This feature vector is the input for a multimodal classifier, which decides on whether there is epileptic activity present or not. Note that in both procedures the classifier steps can be replaced by any kind of decision steps (e.g. rule-based approaches).

nearly every combination of these modalities will result in good performance [9, 32, 113, 126, 139, 181]. For more subtle seizures, accelerometers and EMG are not useful, and a combination of heart or EDA with EEG would be the best option [190]. For absence seizures, however, only EEG seems to lead to a good performance, so a multimodal approach might not be of added value in this case.

In some cases, more than two sensors/modalities are combined [28, 32]. This could lead to further improved performance, but endangers the wearability of the system. In general, modalities could be added to a system as long as the benefit in performance is enough compared to the loss in wearability. If possible, different modalities should be embedded in one device, even though that makes them more vulnerable to noise affecting all modalities at once (losing some added value of the extra modality).

2.4.3 Seizure logging system

Another possibility for usage in the home environment is an automated seizure logging system. In this case, detected seizures are automatically detected by a seizure detection algorithm, and logged in a seizure log file. This logbook can then be used by the clinician in order to have a more objective view on the amount of seizures a patient had since his/her last visit, as patients typically tend to give incorrect logbooks when maintained by themselves [53, 92, 107]. By using such an automated seizure logbook, the clinician will have a better idea on whether the patient is indeed improving or not, for example after the drug dosage has been altered.

Similar approaches and modalities can be used as mentioned above, with the main difference that the algorithms do not have to run in real-time for this application. Seizure detection can be done offline in a server, which allows to use more complex methods for offline seizure detection than often used in real-time detection. As there is no need to detect the seizures fast, it also allows for a more accurate algorithm in this type of setup.

Some commercially available seizure warning systems also provide a seizure logbook, but this is typically for only a small subset of potential seizure types [55].

2.4.4 Seizure prediction system

In seizure detection, seizures are detected while they are occurring. In seizure prediction, the aim is to predict that a seizure will occur within a certain amount

of time (typically between 5-15 and 60 minutes prior to the seizure onset). In such a system, the patient can then put himself/herself in a safe position to protect him/her from the upcoming seizure, or warn a caregiver that a seizure is coming in the near future so the patient is not alone at the time of the seizure.

Seizure prediction algorithms mainly use EEG or ECoG signals [67, 189, 206]. Sometimes already subtle changes in the brain can be found prior to a seizure, which allows seizure prediction. Usage of ECoG electrodes typically leads to better results than scalp EEG electrodes, but are much more invasive for usage. Most users on the other hand do not want to use scalp EEG electrodes in home monitoring, not even for seizure prediction [162]. A few studies claimed to be able to predict seizures based on heart rate analysis [64, 106]. Despite the promising results that are found in multiple studies, most of the proposed prediction algorithms have shown to have reproducibility problems and lack prospective validation studies [63, 67].

2.4.5 SUDEP monitoring and prevention

In an ideal case, also the occurrence of SUDEP can be detected, predicted and/or prevented. Three of the most common causes of death in epilepsy (status epilepticus, SUDEP following a seizure and traumatic accidents) could be detected by well-performing alarm systems, reducing the risk of death in epilepsy in general by allowing patients not to be alone in dangerous situations.

As mentioned in Section 2.1.4, some of the potential risk factors for potential future SUDEP can be detected with dedicated wearable devices (such as ECG sensors) [156]. In case of a prediction of a potential future alarming situation, a patient could then either try to consult a clinician or go to the emergency (depending on the severity of the warning). However, currently still not enough is known about SUDEP to make algorithms for this kind of system. It might however be that in the future, when much more long-term data is recorded from patients who experienced SUDEP (potentially from patients wearing seizure alarm systems), new biomarkers will be found for both short- and long-term detection and prediction of SUDEP.

Chapter 3

Patient-independent ECG-based seizure detection

This chapter introduces a novel real-time method for patient-independent seizure detection using only single-lead ECG. The algorithm extracts the heart rate from the ECG with a novel R peak detection algorithm. The algorithm then detects strong heart rate increases which are possibly caused by epileptic activity by using a new preprocessing procedure based on heart rate gradient analysis. Features are extracted from these heart rate increases and classified using a weighted support vector machine classifier. This classifier is trained to be patient-independent, which means that no data of the evaluated patient is used during the training of the algorithm. Different state-of-the-art patient-independent algorithms from the literature will be implemented and tested on the same dataset as the proposed algorithm in order to compare them with the proposed algorithm.

The first part of this chapter gives a short introduction to ECG-based seizure detection principles and literature. The dataset used in this study is summarized in Section 3.2. The proposed patient-independent seizure detection algorithm is described in Section 3.3. The results of the proposed algorithm and the implementations of several algorithms from the literature are shown and discussed in Section 3.4. This study is also described in [44].

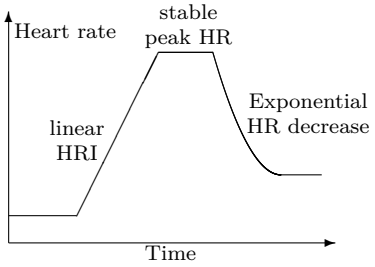


Figure 3.1: Illustration of ictal heart rate changes during temporal lobe seizures.

3.1 Introduction to ECG-based seizure detection

As discussed in Chapter 2, wearable devices can be used to improve the quality of life of refractory epilepsy patients. Three of the mentioned solutions - an automated warning system [183], seizure diary and closed-loop vagus nerve stimulation [201] - , however, require the availability of an automated seizure detection algorithm. In the past, mostly EEG was used for automated seizure detection, but it is not comfortable for the patient to wear a full EEG device for long-term monitoring outside the hospital [162]. Therefore the automated seizure detection should be done using more easier acquirable signals.

One possibility is to use the electrocardiogram (ECG). Literature showed that seizures in patients with temporal lobe epilepsy (TLE) and tonic-clonic (TC) seizures are often accompanied with a strong ictal heart rate increase (HRI) [96, 114, 212]. These HRIs often follow a pattern (see Figure 3.1 for the most common ictal pattern), in which the heart rate (HR) first increases linearly, then stays stable at a certain peak heart rate, and finally decreases exponentially [114, 188]. Variations of this ictal HR behavior can however occur as well (see Figure 2.8) [142, 166].

In literature, mainly three types of solutions for seizure detection are used. In a first solution, the complete epileptic heart rate pattern is automatically detected using fitting approaches on the different phases of the ictal heart rate changes [122, 212]. Another solution is to only detect the strong heart rate increase, typically by assuming that this increase is either atypically fast (high average increase over time) or strong in increase compared to the baseline heart rate (large increase in bpm) [142, 178, 188]. A last solution is to use only heart rate variability (HRV) features in order to detect the seizures automatically, typically searching for a high value in the low frequency band (0.04-0.15 Hz) and a low value in the high frequency value (0.15-0.4 Hz) [23, 51, 64, 77, 99, 100].

Despite the strong patient-dependency of HR features, the patient-independent

(PI) approach is currently preferred for online seizure detection in practice. In a patient-independent approach, no patient-specific data (of the patient who is under evaluation) is used in order to train the algorithm. Patient-specific seizure detection is defined as an algorithm that requires prior patient-specific data in order to operate. Straight-forward patient-specific seizure detection is typically not used in practice due to two main drawbacks:

- There should be sufficient (seizure) data available of the specific patient before the algorithm can be used. This could take more than a week for some patients (for patients with a low seizure frequency) in order to get a fully patient-specific algorithm, which would be a problem for short-term applications. In this case, data collected during diagnostic monitoring will not be sufficient in order to make a robust patient-specific algorithm.
- The data of every new tested patient should be correctly annotated in most current algorithms. This is typically done using the gold standard video-EEG monitoring, and requires the cooperation of medical experts for accurate results. However, often not enough seizures occur during hospital visits for a complete patient-specific algorithm.

Most state-of-the-art heart rate-based seizure detection algorithms are based on thresholding approaches. Most of them, however, do not recommend parameter settings for patient-independent use [122, 142, 188]. Ungureanu et al. (2014) [178] uses minimal patient-specific prior information (i.e. patient's age) in order to tune these parameters to that patient, which is still acceptable for the intended patient-independent goal. The algorithm is trained here automatically in order to optimize seizure detection parameters by using support vector machines (SVM) in a patient-independent way. Most of the algorithms in the literature try to detect the seizures as fast as possible, already a couple of seconds after the seizure onset. Studies however show that a longer-term analysis of the HRI might contain important information (such as the peak HR) for automated seizure detection [42, 43]. Therefore, a method is proposed in this chapter that analyzes the entire HRI. Features of HRIs are then extracted and classified using a PI SVM classifier - using leave-one-patient-out crosstesting (LOPO CT) - into seizure and non-seizure HRIs. Although machine learning has been discussed in previous works for offline seizure detection in neonates [51, 77], it has not been investigated for online HR-based seizure detection in adults and children.

Most studies using the ECG for automated seizure detection only use the heart rate as input for seizure detection. It is however also possible to use specific features from the ECG itself for the detection of seizures. Morphological ECG features such as the QRS complex shape can also be used [79, 198], and also the ECG-derived respiration features can be acquired [195]. The big drawback

of these methods it that they are very sensible to noisy ECG signals, making it hard to use them robustly in seizures with motoric involvement. Therefore, only heart rate features are used in this study.

ECG-based seizure detection is mostly of added value in the detection of clinically subtle seizures, such as partial TLE seizures. These seizures can not be detected by using motion-based algorithms using for example accelerometers or electromyogram, whereas for seizures originating in the temporal the highest prevalence of ictal heart rate changes are found. Therefore, the focus here goes in the detection of complex partial seizures in patients with TLE.

3.2 Data Acquisition

The proposed algorithm was developed based on a dataset containing single-lead ECG recordings from 18 patients with refractory temporal lobe epilepsy, who underwent presurgical evaluation at UZ Leuven. All patients gave consent to their participation in this study. Patients were allowed to move freely inside their hospital room during monitoring. Analysis of one of these patients was left out from this study due to poor overall ECG quality. Seizure onsets were annotated by a clinical expert using video-EEG monitoring as part of a presurgical evaluation, independently of this study. In total 127 complex partial and secondary generalized seizures (currently known as focal impaired awareness and focal to bilateral tonic-clonic seizures respectively in the new ILAE classification system) were recorded during more than 918 hours of continuous ECG data. Table 3.1 gives a detailed overview of the used dataset. The ECG was recorded from lead II with a sampling frequency of 250Hz.

3.3 Methodology

The online seizure detection algorithm using heart rate consists of four main stages: real-time R peak detection, significant HRI extraction, feature extraction and seizure detection.

Table 3.1: Overview of the used dataset in this chapter.

patient number	patient age	number of seizures	seizure types †	record duration‡
1	37	4	1	23:20:19
2	49	10	2,3	26:30:50
3	41	9	2	63:03:34
4	27	13	1,2	70:59:30
5	35	5	2	53:04:57
6	15	10	3	24:40:45
7	17	4	2	43:02:34
8	29	11	2,3	47:06:33
9	27	4	2	30:38:16
10	33	1	1	02:06:25
11	54	2	1,2	15:06:09
12	26	7	1,2	148:08:27
13	34	2	2	10:59:55
14	52	3	1,2	45:37:02
15	9	30	2,3	66:57:00
16	38	11	1,2,3	113:47:48
17	46	1	2	133:12:28
total	[9-54]	127	-	918:10:04

† indicates the seizure severity classes found for each patient. Class 1: Seizures with no ictal motoric manifestations. Class 2: Seizures with discrete motoric manifestations. Class 3: Seizures with strong motoric manifestations. See Section 3.4.8 for detailed definitions.

‡ record duration expressed in hh:mm:ss.

3.3.1 Online R peak detection

In a first step, the single-lead ECG is processed in order to obtain the immediate HR measurements. The immediate HR is computed as

$$HR = \frac{60}{RR}, \tag{3.1}$$

with RR the time difference between consecutive R peaks in seconds. The automated online detection of the R peaks is done by using an adapted version of the R peak detection algorithm described in [208]. The algorithm is adapted in order to operate in real-time according to the following steps:

1. The ECG is filtered by a low-pass (cut-off frequency 40 Hz) and a high-pass (5 Hz) Butterworth filter.
2. The first derivative of this filtered ECG signal is computed.

3. Positive maximum peaks P^+ above a threshold $\frac{T^+}{1.5}$ and a negative minimum peak P^- below a threshold $\frac{T^-}{1.5}$ are searched in real-time. If a combination of such a P^+ and P^- are found with a peak-to-peak distance lower than 140ms, an R peak is detected.

4. T^+ is updated online as

$$T^+ = 0.8 * T^+ + 0.2 * P^+ \tag{3.2}$$

and similarly for T^- with P^-.

The different fixed factors in this algorithm were chosen after the analysis of their performance on the MIT-BIH arrythmia database [75, 128]. T^+ and T^- are initialized as the positive/negative peak values from the third largest/smallest maximum/minimum in the first 5 seconds of the recording. The online search-back procedure from [143] was also added to this procedure for improved R peak detection performance.

3.3.2 Extraction of significant heart rate increases

The ictal HR changes need to be detected automatically from the online constructed HR signal. Detecting the complete ictal HR pattern as shown in Figure 3.1 however has some important drawbacks. First, the entire ictal HR pattern can take over two minutes in some cases, which would lead to a too large detection delay for early seizure detection. A second issue is the fact that multiple variants of the shown ictal pattern exist in practice [142, 166]. Therefore, the proposed method is restricted to detecting and analyzing the incremental phase of the ictal HR pattern. This allows a faster detection of the seizures, and contains sufficient information for accurate seizure detection [38, 99, 142, 178, 188].

The proposed preprocessing algorithm (called *HRI-EXTRACT* from now on) detects this linear HRI phase in three stages, which are mainly based on using the HR gradient of the tachogram. The tachogram is first filtered by applying a median filter (with a filter order of 15 HR measurements). This is done in order to remove smaller peaks inside the linear phase due to respiratory sinus arrhythmia or R peak detection errors. This improves the robustness of the significant HRI extraction procedure. The HR gradient ∇HR is computed as the gradient of the linear fit on 10 HR measurements. In a first stage, the algorithm searches for a gradient larger than 1bpm/s. If such a gradient is found, the algorithm assumes that a strong HRI is occurring. In the second stage, the algorithm searches back in time for when this HRI started (t_{start}).

This is done by searching for the last point for which the gradient is negative or near zero. In the third stage, the end of the linear phase (t_{end}) is found by again searching for the first negative gradient. Once the end of the HRI is found (at time t_{end}), several heart rate features are extracted from this HRI phase, and the following rules are imposed (see Figure 3.2 and Appendix A for the definition of the different variables):

$$\begin{cases} \Delta HR > 10\ bpm & \text{(3.3a)} \\ \dfrac{HR_{peak}}{\overline{HR}_{pre}} > 1.1 & \text{(3.3b)} \\ \dfrac{HR_{peak}}{HR_{rest}} > 1.25 & \text{(3.3c)} \\ \dfrac{\Delta HR}{\Delta t_{HRI}} > 0.35\ bpm/s & \text{(3.3d)} \\ HR_{peak} > 60\ bpm & \text{(3.3e)} \end{cases}$$

with HR_{rest} standing for the estimated average HR at rest for the patient under testing, which is computed on-the-run. HR_{rest} is set to $\overline{HR}_{pre}$ if

- $\overline{HR}_{pre} < HR_{rest}$,
- the linear fit on the data of HR_{pre} results in an absolute gradient $< 0.2 bpm/s$ and
- the above mentioned linear fit is found with $R^2 > 0.7$.

These rules were set after analyzing a different dataset in [38]. As this step is used as a preprocessing step, rule values were chosen so that all seizure HRIs were detected in that independent dataset. If the HRI follows the rules from (3.3a-3.3e), it is assumed that it is a significant HRI, and further processing is required. This algorithm is completely identical for all patients and is summarized in Algorithm 3.1.

3.3.3 Feature extraction

If a significant HRI is detected, several HR features are extracted from the segments before and during this HRI. The extracted features can be divided into three big groups.

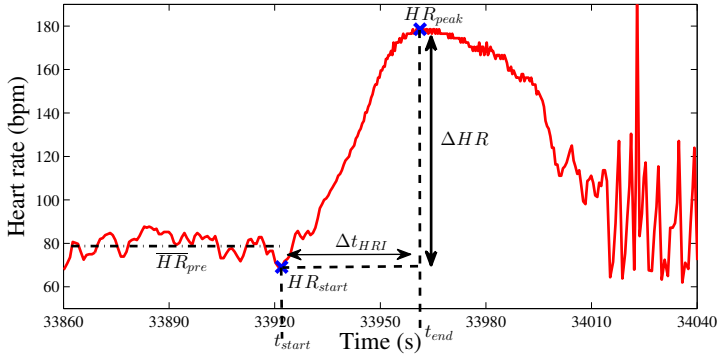


Figure 3.2: Example of an ictal HRI from patient 17. The features required for the extraction of significant HRIs are defined in this figure. $\overline{HR}_{pre}$ indicates the average HR over the 60s before the start of the HRI. HR_{start} indicates the HR at the beginning of the HRI, HR_{peak} the peak HR at the end of the HRI. Δt_{HRI} indicates the length of the HRI (in seconds), ΔHR indicates the amount of increased HR. The EEG seizure onset is located at 33922s in this example.

Algorithm 3.1: *HRI-EXTRACT*

Input: Online HR data
Output: Information of a significantly strong HRI

for each new HR measurement HR_k at time t_{HR_k} **do**
  Filter HR data using median filter
  Compute HR gradient ∇HR_k
  if not currently looking for end of a HRI
    if $\nabla HR_k > 1$ bpm/s
      Look back for last $\Delta HR \leq 0$ bpm/s: start of HRI t_{start}
      Start looking online for t_{end}
    end if
  else if $\nabla HR_k \leq 0$ bpm/s
    Current time is end of HRI $t_{end} = t_{HR_k}$
    if Eq. 3.3a-3.3e apply
      Significant HRI detected!
    end if
    Stop looking for end of HRI
  end if
  Check if HR_{rest} needs an update
end for

The first group contains specific HR features related to the mentioned HRI. These are amongst others HR_{peak}, Δt_{HRI} and the maximal HR gradient ∇HR_{max} during the HRI, which were also discussed in Section 3.3.2. The second group contains heart rate variability (HRV) features extracted from the HRI period. These include linear HRV statistics like the standard deviation of successive differences (SDSD) conform the Task Force recommendations [23]. The last group of features consists of linear HRV features and several HR frequency features extracted from the 60 seconds preceding the HRI. A detailed overview of the extracted features is described in Appendix A.

3.3.4 Feature selection

In the last phase of the algorithm, the features of significant HRIs are classified in order to indicate if these are caused by an epileptic seizure or not. As mentioned in the introduction, a patient-independent classifier is preferred here. The classifier is trained on data from other patients than the tested patient (LOPO CT). The feature selection is also done for each patient separately, using only the data from other patients.

Forward feature selection is used here, incrementally adding features to the pool of selected features. The first two features are chosen simultaneously as a lot of HR features are strongly correlated to each other. Let S be the collection of selected features (initially empty) and the set of N_f features $X = \{X^1, ..., X^{N_f}\}$. The first two selected features are then

$$S^1 = \{[X^{k'}, X^{l'}]\} \tag{3.4}$$

with

$$[k', l'] \quad = \quad \underset{k,l|k\neq l}{\arg\max} F_\beta(SVM([X^k, X^l])) \qquad k, l \quad = \quad 1, ..., N_f \tag{3.5}$$

and

$$F_\beta = (1 + \beta^2) * \frac{Se * PPV}{(\beta^2 * Se) + PPV} \tag{3.6}$$

the F_β-score. Se stands for the sensitivity and PPV stands for the positive predictive value of the support vector machine (SVM) classifier outputs using the selected features $[X^k, X^l]$. The β parameter allows to give β times more importance to sensitivity than to PPV in the computation of the F-score. Additional features were added to S iteratively as

$$S^{m+1} = \{S^m \cup X^{j'}\}, \qquad m \geq 2 \tag{3.7}$$

Table 3.2: Comparison of F_β results for different typical heart rate-based seizure detection results during training. All F_β-scores, PPV and sensitivity (Se) are shown in percentages.

Based on	Se	FP/h	PPV	F_1	F_3	F_5
F_1	74.6	1.9	5.1	9.6	31.6	49.0
F_3	74.6	1.9	5.1	9.6	31.6	49.0
F_5	82.5	2.2	4.8	9.1	31.5	50.9

with

$$j' = \underset{j \notin S^m}{\arg\max} F_\beta(SVM(S^m \cup X^j))\}, \qquad j = 1, ..., N_f \tag{3.8}$$

as long as the addition of this feature $X^{j'}$ leads to an increased F_β-score on the validation set.

The SVM output is found by using LOPO crossvalidation (LOPO CV) on the training set (see next section for classifier details). This means that the training set is split up so that the classifier is trained on every patient with exception of one patient, which is the patient used for validation. The trained classifier is then evaluated on this patient in the validation set. This procedure is repeated so that every patient in the training set is used once for validation. The output for the different validation datasets are then combined to one general sensitivity and PPV output value for the training set. To lower the unbalance in the training sets, the non-seizure samples were limited to the first 100 non-seizure samples of each patient during the training step, similarly as in [126].

Table 3.2 gives an example of the F_β-scores for different β values during feature selection. The F_5-score leads to the extraction of features that result in a more desirable sensitivity (around 80%) and an acceptable false alarm rate ($\pm$2FPs/h) compared to lower β values. Therefore, the F_5-score is used here.

This method of feature selection is chosen here as it is able to identify features that work well in the used generic way. Features that lead to a good classification result in patients for which no data was used during training were chosen, so that the patient-specific impact on the feature selection is reduced as much as possible.

3.3.5 Classification

Let $\{x_i,y_i\}$ be the training data with $x_i \in \mathbb{R}^d$ the input vector containing the selected feature data, and $y_i \in \{-1, +1\}$ the corresponding output values. Let class -1 correspond to seizure samples and class +1 to non-seizure samples.

SVM will then map these input data to a higher dimensional space using a function $\varphi(.)$ [193]. A non-linear mapping φ then makes it possible to make a non-linear separating hyperplane $w^T\varphi(x) + b = 0$, with $w \in \mathbb{R}^d$ the unknown weight vector and $b \in \mathbb{R}$ an unknown constant. The classifier output is then defined as

$$f(x) = sign(w^T\varphi(x) + b). \tag{3.9}$$

The unknown variables w and b can be found by solving the following optimization problem:

$$\begin{aligned} &\min_{w,b,\xi} \frac{1}{2}||w||^2 + C\sum_{i=1}^{N} c_i\xi_i \\ &s.t. \quad \begin{cases} y_i(w^T\varphi(x_i) + b) \geq 1 - \xi_i \\ \xi_i \geq 0 \end{cases} \quad ,\forall i \in [1, N] \end{aligned} \tag{3.10}$$

The term $||w||^2$ is the so-called regularization term of the minimization problem. In separable data, the margin between the hyperplane and the data points closest the hyperplane is defined as $\frac{2}{||w||_2}$. For these data points x_i', it can be seen that $y_i' \left(w^T\varphi(x_i') + b\right) = 1$. In the non-separable case (3.10), the errors ξ_i of data points x_i are defined in comparison to the hyperplanes formed by $y_i' \left(w^T\varphi(x_i') + b\right) = 1$. Maximizing the regularization term leads to a maximization of this margin, leading to a more robust approach which is less vulnerable to overfitting of the training data [193].

Due to the unbalance in the dataset between the datapoints belonging to both classes, a weighted modification of the typical SVM is used [3]. The values c_i are set here to

$$c_i = \begin{cases} \gamma\frac{(N^+ + N^-)}{2N^-} & : y_i = -1 \\ \frac{N^+ + N^-}{2N^+} & : y_i = +1 \end{cases} \tag{3.11}$$

with N^+ and N^- indicating the number of datapoints belonging to classes +1 and -1 respectively. The factor γ in c_i is added in order to increase the impact of misdetecting seizure samples compared to non-seizure samples. γ is set initially to a heuristic value of 3 for the feature selection. Once the features are chosen this way, γ is optimized on the training set by retraining the model for different $\gamma \in [2, 5]$. The γ value leading to the highest F_5-score on the training set is used for training the final SVM model. Both feature selection and γ optimization are treated separately here for computational reasons. Initial tests on 4 patients showed that optimizing both feature selection and γ at the same time led to the

very similar results, but this procedure is computationally much more expensive. Therefore, the separated procedure was used in this study for feature selection and γ optimization.

In practice, however, the dual problem of (3.10) is solved. The dual problem is found by first obtaining the Langrangian of the primal problem (3.10)

$$\mathcal{L}(w,b,\xi;\alpha,\nu) = \frac{1}{2}w^T w - \sum_{i=1}^{N} \alpha_i \left(y_i \left(w^T \varphi\left(x_i\right) + b\right) - 1 + \xi_i\right) - \sum_{i=1}^{N} \nu_i \xi_i \quad (3.12)$$

with $\alpha_i, \nu_i \geq 0$ the multipliers. The solution is then characterized by the saddle point, defined by

$$\max_{\alpha,\nu} \min_{w,b,\xi} \mathcal{L}(w,b,\xi;\alpha,\nu) \quad (3.13)$$

which leads to

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial w} = 0 & \rightarrow w = \sum_{i=1}^{N} \alpha_i y_i \varphi\left(x_i\right) \\ \frac{\partial \mathcal{L}}{\partial b} = 0 & \rightarrow \sum_{i=1}^{N} \alpha_i y_i = 0 \\ \frac{\partial \mathcal{L}}{\partial \xi_i} = 0 & \rightarrow C c_i = \alpha_i + \nu_i, \forall i \in [1,N] \end{cases} \quad (3.14)$$

Replacing the equations from (3.14) in equation (3.12) leads to the dual problem:

$$\begin{aligned} &\min_{\alpha} \frac{1}{2} \sum_{i,j=1}^{N} y_i y_j \alpha_i \alpha_j K\left(x_i, x_j\right) - \sum_{i=1}^{N} \alpha_i \\ &s.t. \quad \begin{cases} \sum_{i=1}^{N} \alpha_i y_i = 0 \\ 0 \leq \alpha_i \leq C c_i \quad , \forall i \in [1,N] \end{cases} \end{aligned} \quad (3.15)$$

with $K(x_i, x) = \varphi\left(x_i\right)^T \varphi(x)$ the kernel function [193]. A gaussian (also called RBF) kernel is used here with kernel parameter σ, which is defined as

$$K(x_i, x_j) = exp\left(\frac{-||x_i - x_j||_2^2}{\sigma^2}\right). \quad (3.16)$$

Solving (3.15) leads to the required values of α_i, but parameter b still needs to be defined. This can be done by using the Karush-Kuhn-Tucker (KKT) conditions, which state that in an optimal solution, the product of the dual variables and the constraints should be zero. The KKT conditions yield [173]

$$\begin{cases} \dfrac{\partial \mathcal{L}}{\partial w} = 0 \rightarrow w = \displaystyle\sum_{i=1}^{N} \alpha_i y_i \varphi\left(x_i\right) & \\ \dfrac{\partial \mathcal{L}}{\partial b} = 0 \rightarrow \displaystyle\sum_{i=1}^{N} y_i \alpha_i = 0 & \\ \dfrac{\partial \mathcal{L}}{\partial \xi_i} = 0 \rightarrow \alpha_i + \nu_i = C c_i & ,\forall i \in [1, N] \\ \alpha_i \left(y_i \left(w^T \varphi(x_i) + b\right) - 1 + \xi_i\right) = 0 & ,\forall i \in [1, N] \\ \beta_i \xi_i = 0 & ,\forall i \in [1, N] \\ \nu_i \geq 0 & ,\forall i \in [1, N] \\ \beta_i \geq 0 & ,\forall i \in [1, N] \end{cases} \tag{3.17}$$

From (3.17) it can be seen that for a solution $\alpha_i^* \in (0, Cc_i)$, $\xi_i^* = 0$ so that

$$y_i \left(w^T \varphi\left(x_i\right) + b\right) - 1 = 0 \qquad \text{for } \alpha_i \in (0, Cc_i) \tag{3.18}$$

Parameter b can then be estimated by taking the average of the values found for the above equation for all $0 < \alpha_i < Cc_i$.

The resulting classifier in the dual formulation then becomes

$$\begin{aligned} y(x) &= sign\left(\sum_{i=1}^{N} \alpha_i y_i K(x_i, x) + b\right) \\ &= sign\left(\sum_{i=1}^{\#SV} \alpha_i y_i K(x_i, x) + b\right) \end{aligned} \tag{3.19}$$

which indicates that only the data points for which $\alpha_i > 0$ need to be evaluated. The data points for which $\alpha_i > 0$ are called support vectors (SV). Due to the sparsity properties of SVMs, this number of SV is only a small fraction of the total dataset, which greatly facilitates the computation of (3.19).

The entire procedure for seizure detection (called *HRI-SVM*) is summarized in Algorithm 3.2 and Figure 3.3.

Algorithm 3.2: Procedure for online seizure detection (*HRI-SVM*)

Input: Online HR data of patient P_p, offline HRI data of other patients $pat_{pool} \setminus P_p$
Output: Alarm in case of an assumed seizure for patient P_p

Offline preprocessing: feature selection

Let pat_{pool} be the collection of patients
Let $feat_{pool}$ be the collection of features
Initialize $feat_{sel}$ as empty and $F_5^{prev} = 0$
while adding features to $feat_{sel}$ leads to improved performance in patients $pat_{pool} \setminus P_p$ **do**
  for each feature $feat_f \in feat_{pool} \setminus \{feat_{sel}\}$ **do**
    $feat_{test} = feat_{sel} \cup feat_f$
    for each patient $P_i \in pat_{pool} \setminus \{P_p\}$ **do**
      Train classifier on HRIs of patients $pat_{pool} \setminus \{P_p, P_i\}$ using features $feat_{test}$ with γ=3
      Test classifier on HRIs of patient P_i
    end for
    Evaluate F_5-score F_5^f on patients in $pat_{pool} \setminus \{P_p\}$
  end for
  if $max(F_5^f) > F_5^{prev}$
    $feat_{sel} = feat_{sel} \cup feat_{\arg\max_f F_5^f}$
    $F_5^{prev} = max(F_5^f)$
  end if
end while

Offline preprocessing: γ optimization

$\gamma_{max} = 3$
for $\gamma_c \in [2, 5]$ **do**
  for each patient $P_i \in pat_{pool} \setminus \{P_p\}$ **do**
    Train classifier on HRIs of patients $pat_{pool} \setminus \{P_p, P_i\}$ using features $feat_{sel}$ with $\gamma = \gamma_c$
    Test classifier on HRIs of patient P_i
  end for
  Evaluate F_5-score F_5^a on patients in $pat_{pool} \setminus \{P_p\}$
  if $max(F_5^a) > F_5^{prev}$
    $F_5^{prev} = max(F_5^a)$
    $\gamma_{max} = \gamma_c$
  end if
end for
Train final classifier using $feat_{sel}$ on data of $pat_{pool} \setminus P_p$ with $\gamma = \gamma_{max}$

Real-time analysis of HR data from patient P_p

for each new HR measurement HR_k at time t_{HR_k} **do**
  Extract significant HRIs (see Algorithm 3.1)
  if the end of a significant HRI is detected
    Extract selected features from the significant HRI
    Classify using patient-independent classifier with $feat_{sel}$
    if class output of classifier = 1
      Alarm!
    end if
  end if
end for

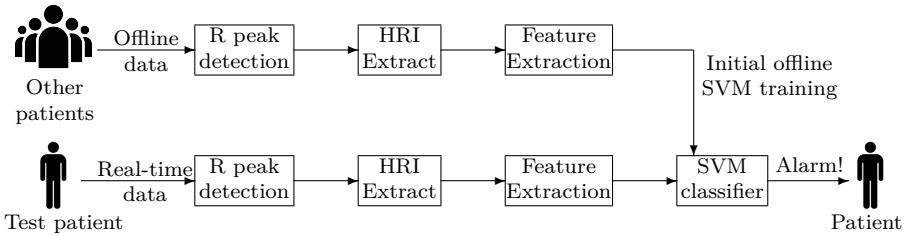


Figure 3.3: Overview of the proposed patient-independent seizure detection procedure followed in this chapter.

3.3.6 Evaluation criteria

Online automated seizure detection algorithms are typically analyzed using four metrics. The sensitivity gives the percentage of seizures that are detected by the algorithm. The number of false alarms are evaluated by both the number of false positives per hour (FP/h) and the PPV. The fourth metric is the detection delay, which is the delay between the seizure onset and the moment the seizure is detected by the algorithm. The seizure onset is taken as the scalp EEG onset if this is available. In some cases (for patients 2, 6 and 8), no clear EEG onsets were found. For these cases the clinical onsets are used as reference onset. Seizures were detected correctly when an alarm was given between 30 s before and 90 s after the seizure onset. These boundaries are chosen as none of the seizures showed any significant HRIs before 30 s prior to the EEG onset, which is in consensus with the findings of [88]. HRIs typically end before 90 s after the EEG onset, and detecting a seizure 1.5 minutes after the EEG onset is considered to be too late for real-time seizure detection. Alarms within 1 minute of each other are counted as one alarm [142]. General measurements of the entire dataset can be divided into overall performance (computed on the total number of seizures) and patient-averaged performance (the average of the performances of each patient). Unless specifically mentioned otherwise, overall performance is meant for the evaluation of the general performance in the rest of this chapter.

3.3.7 Tested HR-based seizure detectors from the literature

Other patient-independent seizure detection algorithms from the literature were also implemented and tested on the used dataset [142, 178, 188]. Also an adapted version of the mentioned algorithm, using a linear SVM with the features found in [38] (∇HR_{max}, Δt_{HRI} in #RR intervals and the pre-ictal spectral entropy) is evaluated. The parameters for the algorithms from [142] and [188] (from now on

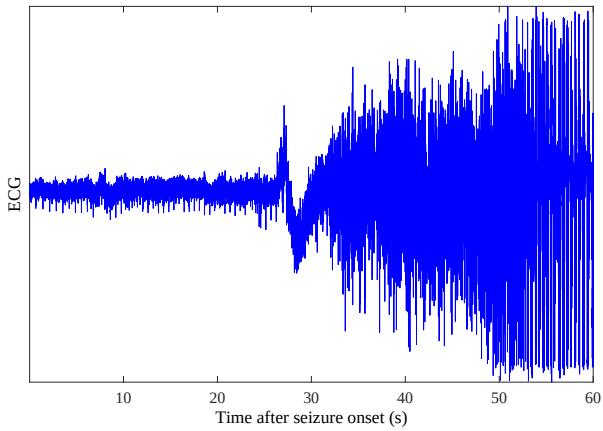


Figure 3.4: Example of ictal noise in the ECG.

referred to as *Osorio* and *vanElmpt* respectively) are set twice so that they result in either a similar sensitivity (marked by *(se)*) or similar FP/h rate (marked by (*fph*)) as the proposed algorithm. This allows a better comparison between these algorithms and the proposed algorithm. This procedure is not required for the algorithm from [178] as the optimal patient-independent parameter settings were already discussed in the corresponding paper. The algorithms were faithfully replicated based on the algorithm descriptions from the corresponding papers and were intensively tested for their correctness. The same R peak detection algorithm is however used for all algorithms.

3.4 Results and discussion

3.4.1 *HRI-EXTRACT*

The *HRI-EXTRACT* algorithm detected 116 of the 127 seizures (91.3%). Some seizures were missed either due to the absence of a strong ictal HRI or due to the presence of artifacts in the ECG during the ictal period. The non-seizure HRIs that were detected (at average 6.0 non-seizure HRIs per hour), were typically other sympathetic activations caused by non-seizure stimuli like arousals, or false detections due to R peak detection errors caused by strong ECG noise (see Figure 3.4 for an example of ictal ECG noise). This preprocessing algorithm thus worked well for the extraction of significant HRIs.

Table 3.3: Results for fHRI-SVM in case a fixed γ is used for all patients during SVM training.

γ	1	2	2.5	3	4	5
Se (%)	69.3	81.1	84.3	85.0	85.0	84.3
FP/h	1.32	1.91	2.11	2.26	2.48	2.75
F_5-score (%)	51.1	53.2	53.3	52.4	50.5	47.9

3.4.2 Feature selection & γ optimization

Feature selection typically resulted in the selection of $\overline{HR}_{pre}$(selected in 17/17 LOPO CV runs), HR_{peak} (17/17), HR_{start} (17/17) and the pre-ictal SDSD (15/17). The combination of these four features was the most common combination during feature selection (15/17). All these features are easily extractable in real-time with a very limited computational load, making them easy implementable in wearable devices. It took 1785 s to analyze 918 hours of HR data on a desktop pc with 24 GB RAM and a 3.6 GHz processor using Matlab®.

The γ parameter in Eq. 3.11 was optimized automatically during the training procedure. A value of 2 was chosen the most (12/17 LOPO CV runs), 2.5 was chosen 4 times, 3 only once. If the same γ was used for every patient, $\gamma \in [2, 3]$ indeed resulted in the best performance (see Table 3.3).

All results in this chapter are obtained by using a Gaussian kernel K during SVM training. After evaluating the results for the Gaussian kernel, the training of the SVM classifiers was also redone for other kernels using the features mentioned above. No other kernel however led to an increase in performance compared to the Gaussian kernel.

3.4.3 Fastened algorithm *fHRI-SVM*

A faster seizure detection is possible with a small adjustment to the originally proposed algorithm. As discussed in Section 3.4.2, $\overline{HR}_{pre}$, pre-ictal SDSD, HR_{start} and HR_{peak} are selected as the used features in 15/17 LOPO CT runs. The first 3 of these features can all be extracted when ∇HR_{max} is found, leading to the possibility to enhance the proposed algorithm. The algorithm can be changed so that once ∇HR_{max} is detected, each new HR measurement HR_k is assumed to be the end of the HRI (so that $HR_{peak} = HR_k$). Equations (3.3a)-(3.3e) can then already be evaluated, and if this newly formed quadruple of features leads to a seizure HRI (using the original classifier), an alarm can go off much faster. If this does not hold, the procedure keeps iterating until the

real end of the HRI is found. This does not lead to many extra false alarms compared to *HRI-SVM* because:

- The filtered HR signal is an increasing signal during the HRI. The parameters evaluated in (3.3a)-(3.3c) and (3.3e) will therefore only increase or stay unchanged during the HRI, and so will not cause extra false alarms.
- The parameter in Eq. 3.3d can decrease if the HR gradient decreases significantly. This can happen near the end of the HRI as some HR gradients are nearly zero near the end of the HRI for a long period. The impact of this is however limited as it typically only occurs in very strong HRIs that are mainly caused by seizures.
- The used SVM classifier typically led to a classification boundary which could be simplified to

$$class = \begin{cases} -1 & : HR_{peak} \leq thr \\ +1 & : HR_{peak} > thr \\ -1 & : HR_{peak} >= \pm\ 200bpm \end{cases}$$

  with thr the threshold value as defined by the classification boundary given the other feature values. In this case, no HRIs that were classified as non-seizure in the original algorithm can be classified as a seizure in this approach except if they had a HR_{peak} larger than $\pm 200bpm$. Such a high HR is however only reached in this database due to grouped R peak detection errors. This caused some extra false alarms with this adjusted algorithm.

The pseudocode for the fastened version (called *fHRI-SVM* from now on) is shown in Algorithm 3.3, which replaces lines 31-40 in Algorithm 3.2. Unless specifically mentioned, this algorithm and its results are discussed from now on.

3.4.4 Results for *fHRI-SVM*

Table 3.4 gives an overview of the obtained results for each patient individually by using the *fHRI-SVM* algorithm. The results of all proposed algorithms and the algorithms from the literature are shown in Table 3.5. An overall sensitivity of 81.9% (range 50.0-100.0%) is obtained, with on average 1.97 FP/hour (range 0.88-3.52 FP/h). The seizures were detected on average 17.8 s after the seizure onset (range -18.2-54.3 s). The used SVM classifier showed a sensitivity of 88.8%, with a specificity of 67.9% in the *HRI-SVM* algorithm.

Algorithm 3.3: Pseudocode for *fHRI-SVM*.

Input: Online HR data, patient-independent classifier SVM
Output: Information of a significantly strong HRI

for each HR measurement HR_k at time t_{HR_k} **do**
  Filter HR data using median filter
  Compute HR gradient ∇HR_k
  if not currently looking for end of a HRI
    if $\nabla HR_k > 1$ bpm/s
      Look back for last $\Delta HR \leq 0$ bpm/s: start of HRI t_{start}
      Extract pre-ictal HR SDSD, $\overline{HR}_{pre}$ and HR_{start}
      Start looking online for t_{end}
    end if
  else
    Try $t_{end} = t_{HR_k}$ so that $HR_{peak} = HR_k$
    if Eq. 3.3a-3.3e apply and SVM(extracted features)=1
      Significant HRI detected!
    end if
    if real end of HRI is found
      stop looking for end of HRI
    end if
  end if
  Check if HR_{rest} needs an update
end for

Table 3.4: Results per patient for *fHRI-SVM*.

Patient	Se (%)	FP/h	PPV(%)	$\overline{delay}$ (s)
1	75.0	1.84	6.52	-16.7
2	80.0	1.14	21.05	17.8
3	100.0	3.52	3.90	9.5
4	61.5	1.72	6.15	54.3
5	60.0	2.96	1.87	32.2
6	80.0	1.58	17.02	9.0
7	100.0	3.21	2.82	4.9
8	90.9	1.91	10.00	12.3
9	100.0	2.71	4.60	18.1
10	100.0	2.37	16.67	-18.2
11	100.0	1.13	10.53	-2.6
12	100.0	2.00	2.31	50.6
13	50.0	2.18	4.00	4.6
14	66.7	0.88	4.76	9.6
15	86.7	2.02	16.15	10.3
16	63.6	1.26	4.67	33.3
17	100.0	1.70	0.44	2.1
Overall	81.9	1.97	5.43	17.8
Pat.-av.	83.2	2.01	7.85	13.6

Table 3.5: Comparison of the results for the different evaluated algorithms. Bold numbers indicate the algorithms with similar sensitivity or FP/h rate as the proposed algorithm. Values for sensitivity (Se) and positive predictive value (PPV) are given in percentages, detection delay is expressed in seconds.

	Overall				Patient-averaged			
Method	Se	FP/h	PPV	$\overline{delay}$	Se	FP/h	PPV	$\overline{delay}$
HRI-EXTRACT (Alg. 3.1)	91.3	6.00	2.06	25.26	91.7	6.43	2.81	22.08
HRI-SVM (Alg. 3.2)	**81.1**	**1.93**	5.50	26.54	81.4	1.94	7.98	24.94
fHRI-SVM (Alg. 3.3)	**81.9**	**1.97**	5.43	17.83	83.2	2.01	7.85	13.61
De Cooman et al. [38]	**81.1**	2.93	3.69	18.28	85.2	3.59	4.77	14.03
Ungureanu et al. [178]	55.1	4.56	1.65	19.23	55.5	4.20	5.07	23.33
Osorio(se) [142]	**80.3**	4.19	2.59	15.86	81.9	4.31	3.63	5.89
Osorio(fph) [142]	69.3	**1.98**	4.61	17.07	73.2	2.19	8.49	8.42
vanElmpt(se) [188]	**80.3**	6.69	1.63	10.82	81.7	6.88	3.43	10.63
vanElmpt(fph) [188]	44.9	**2.01**	2.99	16.09	42.4	2.38	6.87	13.71

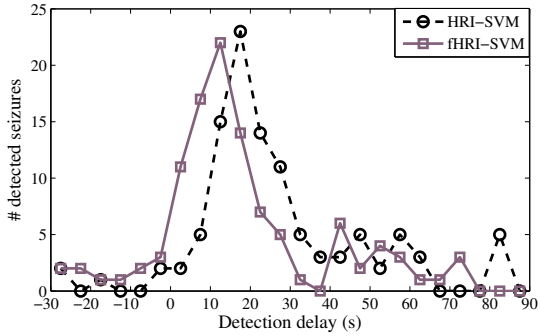


Figure 3.5: Histogram of the detection delay between the seizure onset and the alarm time for the proposed seizure detection algorithms.

3.4.5 Comparison between *HRI-SVM* and *fHRI-SVM*

By using the *fHRI-SVM* algorithm, the detection delay dropped from 26.54 s to 17.83 s compared to the original *HRI-SVM* algorithm. Figure 3.5 shows the detection delay histograms for both *HRI-SVM* and *fHRI-SVM*. An overall shift towards lower delays can be noticed, causing less large delays (>60 s). 83 of the detected seizures (79.8%) were detected within 30s of the seizure onset using *fHRI-SVM*, and 10 seizures (9.6%) were detected before the seizure onset.

The number of false alarms increased slightly to 1.97 FP/h using *fHRI-SVM*, but also the sensitivity increased to 81.9%. The slight increase in detections

compared to the results of *HRI-SVM* is due to the 2 possible causes described in Section 3.4.3, which happen only rarely as can be seen in the limited increase in false alarms.

3.4.6 Comparison between *fHRI-SVM* and the literature

The proposed algorithm showed to give the best combination of sensitivity and PPV on the mentioned dataset compared to the tested algorithms of the literature. If around 2 FP/h were requested, this caused a significant sensitivity drop in both Osorio(fph) and vanElmpt(fph). Ungureanu et al. (2014) resulted in both significant lower sensitivity and PPV compared to the proposed algorithm. The main reason for this is that this algorithm is focused on detecting seizures with high HR values (> 110 *bpm*), and showed problems in the detection of seizures that did not lead to tachycardia (defined here as $HR_{peak} > 100$ *bpm* and $\Delta HR > 10$ *bpm*). This is illustrated in Figure 3.6, which shows the sensitivity of all algorithms for both seizures leading to tachycardia (86/127 seizures) or not (41/127 seizures). All algorithms were able to detect at least 60% of the seizures with tachycardia, but bigger differences were found for the non-tachycardia seizures. The drop in sensitivity for the fph-based versions compared to the se-based versions of Osorio and vanElmpt were mainly caused by a big drop in sensitivity for the non-tachycardia seizures. The proposed method was able to detect 91.9% of the tachycardia seizures and 61.0% of the non-tachycardia seizures.

Figure 3.7(a) shows the boxplots for the inter-patient sensitivity variability for the different tested algorithms. The proposed algorithm always led to a sensitivity of at least 50% in this dataset. All other algorithms from literature showed a greater sensitivity variability, leading to 0% sensitivity for some patients in Ungureanu et al. and vanElmpt(fph).

Larger differences were found in the number of false positives when comparing the proposed algorithm to the literature. The proposed method led towards at least 2.2 FP/hour less than the other algorithms with comparable sensitivity, a decrease of at least 50% of false alarms. Figure 3.7(b) shows that the proposed method also resulted in less inter-patient variability in the number of false alarms. The proposed algorithm thus not only worked best overall for these statistics, but it also resulted in the least inter-patient variability for both FP/hour and sensitivity, leading to reasonable results for all patients. PPV values (see Figure 3.7(c)) vary more between patients for all algorithms, but this will be discussed in Section 3.4.7. Figure 3.8 shows the results for all evaluation measurements per patient for all tested algorithms.

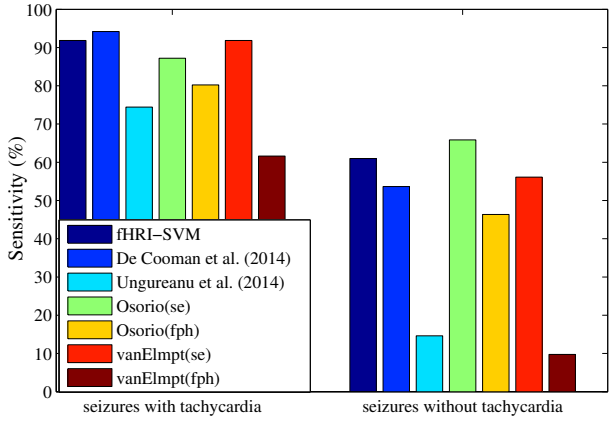
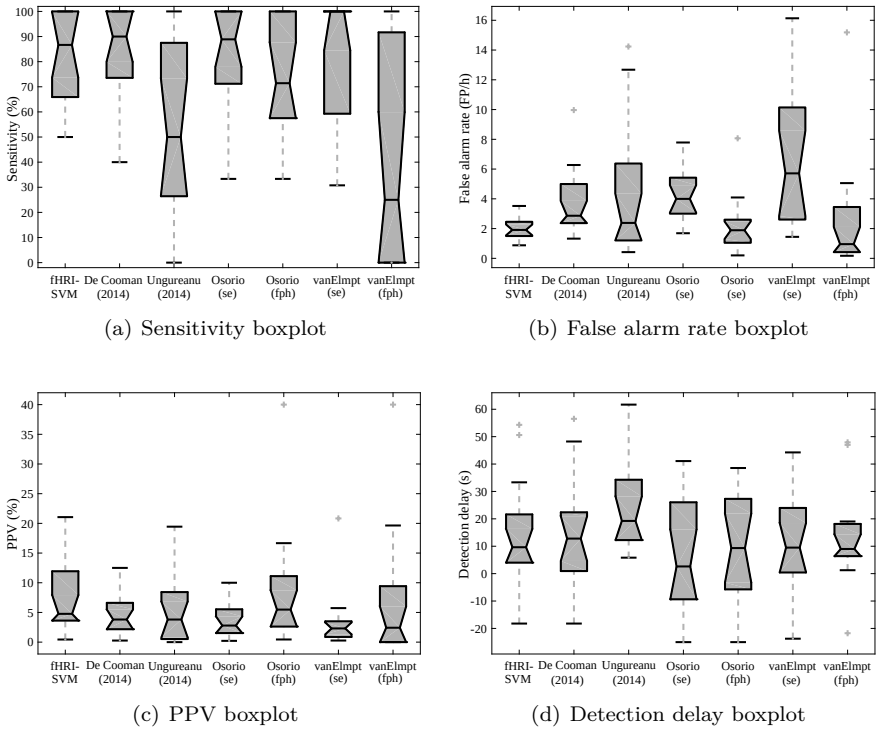


Figure 3.6: Sensitivity performance for seizures with and without tachycardia for the different tested algorithms.

Figure 3.9 shows the sensitivity and FP/h performance for the fHRI-SVM algorithm and the algorithm from [38] in case the SVM bias term b is altered. Also the results for different parameter values for the algorithms from [142] and [188] on this dataset are shown. It can be seen that the fHRI-SVM algorithm always led to the best sensitivity-FP/h combination (if Se>50%). Some possible explanations why these algorithms from the literature show significantly lower results on this dataset, as discussed in the corresponding papers, are:

- A much larger patient group has been analyzed than in [188] (3 patients) and [178] (5 patients).
- Both day and night data were used, while [178] only discussed detection during the night.
- No preselection on the patients is done based on their ictal HR or motoric responses, nor on their seizure frequency [178, 188]. Only a preselection towards TLE patients is done here, which contained a wide range of seizures and is the largest group of epilepsy patients. Sensitivity results of [142] were only computed on clinical seizures.
- No restriction was done on the patient movements [142] (although they are still limited to the hospital rooms).
- No manual correction of R peak detection errors was performed [188].



(a) Sensitivity boxplot

(b) False alarm rate boxplot

(c) PPV boxplot

(d) Detection delay boxplot

Figure 3.7: Boxplots showing the inter-patient variability in sensitivity (a), false alarm rate (b), PPV (c) and detection delay(d) for the different tested algorithms.

When looking at the detection delay, the *HRI-SVM* algorithm detected the seizures later than the other algorithms (see Table 3.5). *HRI-SVM* detected the seizures on patient-average after 24.9 s, while for Osorio(se) this was on average 5.9 s after the seizure onset. The main reason for this, is the fact that *HRI-SVM* is required to wait for the HRI to stop, whereas other algorithms can already trigger an alarm sooner. It seems however that detecting the seizures too soon also causes a large amount of extra false alarms. Figure 3.10 illustrates this for a non-seizure and a seizure HRI coming from the same patient during sleep. The start of both HRIs are similar, hence making it hard to detect this seizure very rapidly without detecting the non-seizure HRI. The only difference seems to be that the seizure HRI maintains the strong HR acceleration longer, thus leading to a higher peak HR. *fHRI-SVM* tries to alarm as soon as it is sufficiently sure that the HRI is caused by a seizure. More than 50% less false

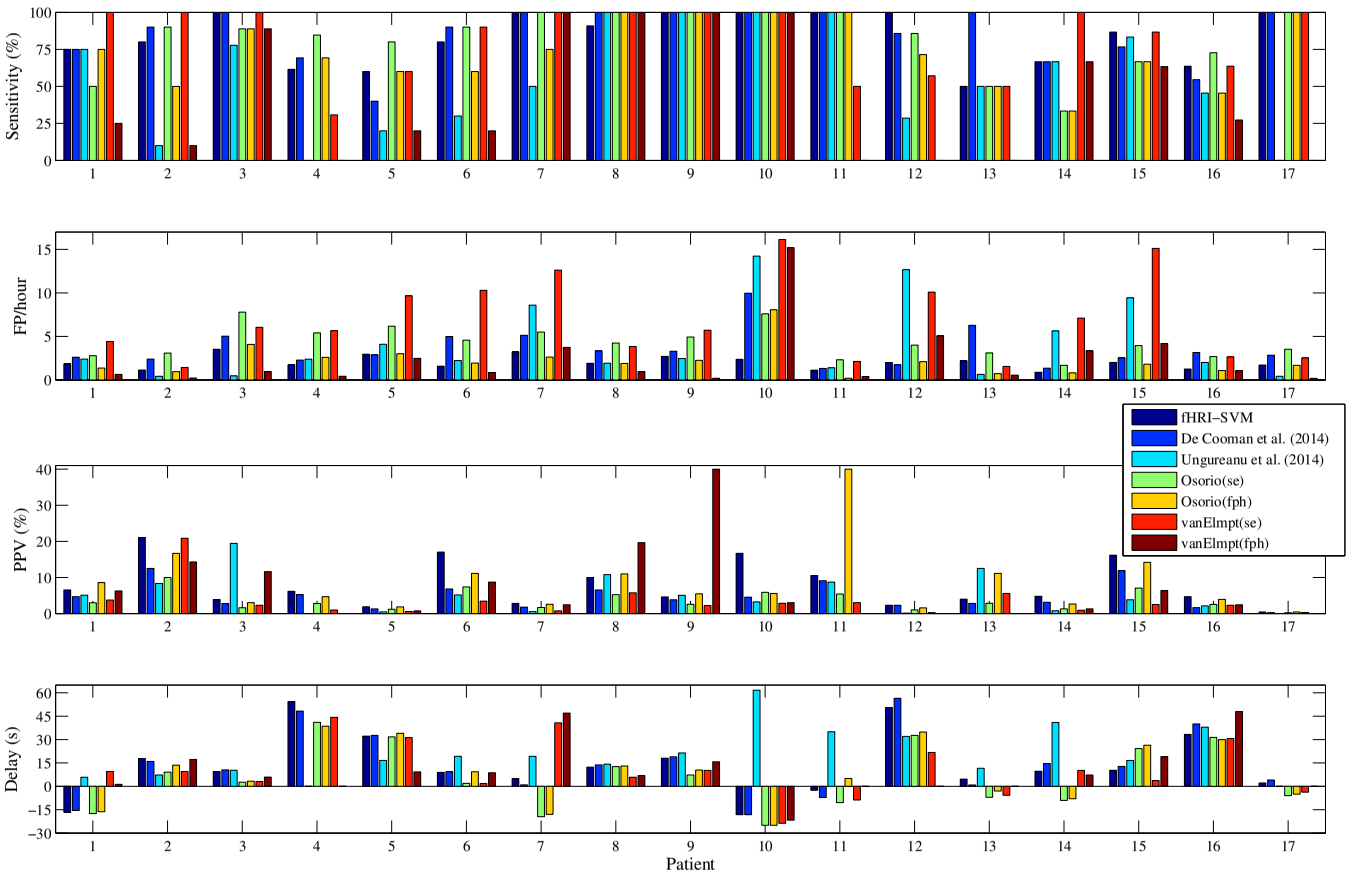


Figure 3.8: Patient-specific sensitivity, FP/hour, PPV and detection delay results for the different tested patient-independent seizure detection algorithms.

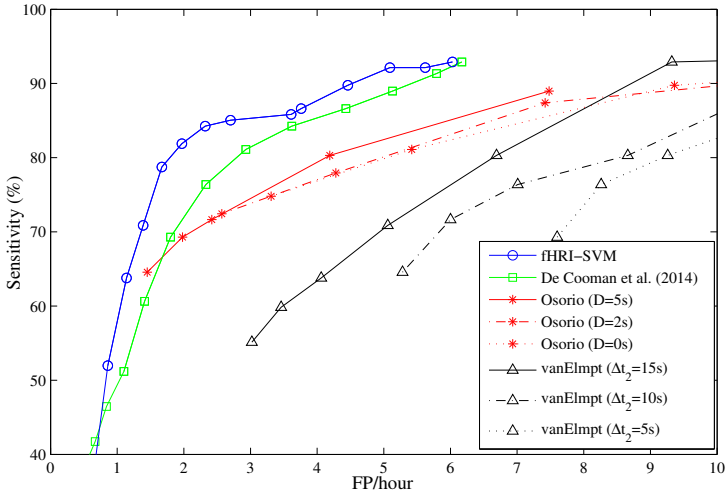


Figure 3.9: Sensitivity versus FP/hour results for the different tested algorithms. For fHRI-SVM and the method from [38] the different performance results were found by altering the SVM bias term b. For the Osorio and vanElmpt algorithm, the parameter values from the papers are altered. For Osorio, each line contains a constant delay term D with varying heart rate threshold values [142]. For vanElmpt, each line contains a constant second time window Δt_2, with changing heart rate threshold values [188]. We refer to the corresponding papers for details of these algorithms and their parameters.

alarms were found compared to the literature in this way, with only a limited extra patient-averaged detection delay (3.0 s patient-averaged compared to vanElmpt(se), 7.7 s compared to Osorio(se)). Part of this extra delay (typical 3-5 s) is caused by the median filter in *HRI-EXTRACT*. The results however show that this step is very helpful in removing R peak detection errors from the heart rate signal, hence getting more robust seizure detection results.

3.4.7 Inter-patient variability

As the results vary strongly for different patients (see Figure 3.7 and 3.8), a closer look is taken on the inter-patient variability for the sensitivity, PPV, FP/h and detection delay using the *fHRI-SVM* algorithm.

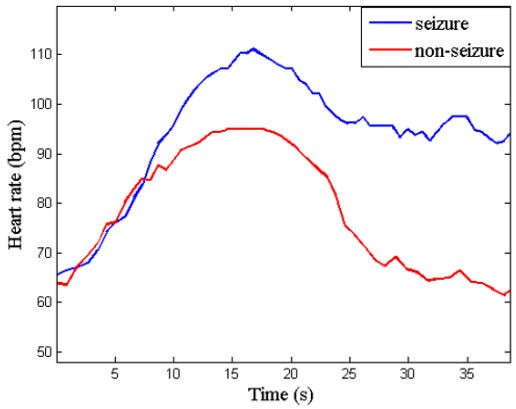


Figure 3.10: Comparison of a seizure and a non-seizure HRI originating from the same patient. Both HRIs were detected during sleep, within only a few hours of each other.

Sensitivity

Only 2 patients did not have a sensitivity higher than 60%. For patient 5, only 3/5 seizures were detected, as the other 2 seizures did not contain strong ictal HRIs ($\Delta HR = 20$ bpm with $\Delta t_{HRI} = 80$ s).

Patient 13 only showed 50% accuracy as one of the two seizures was missed due to a high pre-ictal HR HR_{pre}. Seizures in other patients were typically missed by atypical HRIs for that patient or due to strong ictal ECG noise.

FP/h

The number of FP/h depended mainly on how often each patient had non-epileptic sympathetic activations during the recordings and how strong these sympathetic activations typically were. For patients for which these were typically stronger (e.g. children), non-epileptic sympathetic HRIs (caused by for example arousals or mental stress) led much faster to stronger ΔHR values, causing more FPs. In some patients, the average ECG noise levels (both ictal and inter-ictal) were much higher compared to other patients. This unavoidably led to R peak detection errors and extra false alarms. Although the effect was rather limited for the proposed algorithm, this led to a much higher FP/h rate for for example patient 10 for the algorithms from the literature (see Figure 3.8).

PPV

As described above, the number of FPs mainly depended on the number of non-epileptic sympathetic HRIs, and was independent on the number of seizures. The PPV values were therefore strongly correlated with the seizure frequency. Big negative outliers in PPV values were found in patients with a low seizure frequency (e.g. patient 17).

Detection delay

The found detection delay for *fHRI-SVM* mainly depended on two factors:

- The start of the HRI compared to the EEG onset: Previous studies showed that the time difference between the HR changes and the EEG onset can vary between different seizures [50, 114]. HRIs can start either just before or after the EEG onset. The sooner the HRI starts, the higher the chance to detect the seizure faster.
- The strength of the sympathetic activation: The seizure can only be detected when the HR has increased sufficiently compared to $\overline{HR}_{pre}$. For patients with relative slowly increasing sympathetic activations, it took longer to get to this temporary HR_{peak} required for alarming (see for example patients 4 and 12).

3.4.8 Relationship between clinical phenomena and seizure detection results

For home monitoring applications, it is interesting to see how the different algorithms cope with different types of ictal motoric manifestations. The seizures are therefore divided into three categories:

- Class 1 (27 seizures): Non-motor seizures with no clear motoric manifestations (except for eye-blinking and lip movements)
- Class 2 (45 seizures): Seizures with discrete motoric manifestations (subtle dystonia and discrete motoric agitation)
- Class 3 (55 seizures): Strong convulsive seizures including tonic, clonic and tonic-clonic seizures

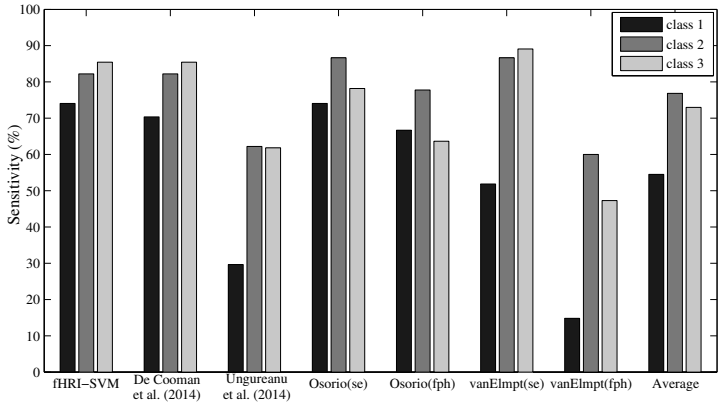


Figure 3.11: Sensitivity results for each class of seizure severity for the different tested algorithms. The average sensitivity over all algorithms for the 3 severity classes is also shown in this figure.

Table 3.1 gives an overview of which seizure classes were found in each patient. The seizures from group 3 are clinically the most severe seizures, which most alarm systems aim to detect due to their high risk of injuries [184]. Patients can however lose their consciousness during most of the class 1 seizures, which might indirectly lead to potential dangerous situations (e.g. walking on stairs), so it is also important to detect those seizures. Figure 3.11 shows the sensitivity for each class of seizures for the evaluated seizure detection algorithms. It can be seen that for most algorithms, the sensitivity is higher for seizures with ictal motoric activity (classes 2 and 3) than for non-motor seizures. The main reason for this is the fact that these non-motor seizures have a much lower tachycardia occurrence (37.0%) compared to classes 2 and 3 (68.9% and 81.8%), making it harder to detect them. Class 2 seizures are found slightly easier by methods from literature than class 3 seizures, which is probably due to the fact that these contain less ECG noise due to more subtle ictal movements. The proposed method was able to detect 74.0%, 82.2% and 85.5% of the seizures from respectively classes 1, 2 and 3.

3.4.9 Applications of HR-based seizure detection

Automated seizure detection algorithms using ECG can be used for several purposes. A first purpose is as part of an alarm system in a home environment. The algorithm can automatically give a warning to caregivers whenever a seizure

is detected by the algorithm. This is a useful system, certainly for children with epilepsy and their parents. The number of FP/hour for these patient-independent algorithms show that work still needs to be done in order to make such a system successful for 24/7 monitoring. The number of FPs during the day are typically too high due to patient activity. Another study on children with epilepsy showed that such a system can however produce better FP rates ($\pm$1 FP/h) when the system is used during the night [42]. In this case, the false alarm rate mainly depends on the sleep quality of the patient. These results could however not be replicated here due to the use of sleep deprivation during monitoring.

A second practical use would be to include such an automated seizure detection in a closed-loop VNS system. VNS is shown to be effective for preventing and possibly shortening ongoing seizures [84]. By automatically detecting the seizures, it can trigger the VNS activation during this seizure. It is important to keep the number of false VNS activations to a minimum to prolong the battery life-time, although continuous periodic activation is provided in current devices. An alarm of the seizure detection algorithm will cause an extra activation within the periodic activation scheme with the intention of trying to shorten the seizure duration by ictally triggering the VNS. The proposed method showed a strong decrease in false alarms for a similar sensitivity compared to the literature, hence avoiding false VNS activations and unnecessary disturbances in the VNS activations. Due to the slightly increased detection delay, it will however take a few seconds longer to stop ongoing seizures.

The system can also be used in the hospitals or revalidation centers. For the investigated seizures from temporal lobe epilepsy patients, over 90% of the seizures were accompanied with an ictal strong HRI. An interesting fact can be seen for the patients in which the seizures were not (clearly) visible in EEG. Automated seizure detection algorithms using the EEG would probably not be able to detect these seizures accurately, whereas the proposed algorithm using ECG was able to correctly detect 83.8% of these seizures (31 seizures in total). This algorithm can thus also be useful in the epilepsy monitoring facilities, possibly as enhancement of online seizure detection applications using EEG [1, 72, 176]. The discussed methods also have the advantage that some of the seizures can be detected by using ECG before the seizure is clearly visible during video-EEG monitoring (as can be seen for example in Figure 3.5).

The proposed algorithm only uses the heart rate signal as input for the automated seizure detection. In this case, the ECG is used to extract the heart rate in real-time. The method however also works with heart rate signals acquired trough other modalities (e.g. PPG), which might lead to a more comfortable device for the patient. PPG however tends to lead to more noisy heart rate signals with a bit of loss on the HRV features, depending on the location of the

PPG sensor [180, 191].

Due to the use of a patient-independent algorithm, the applications can be used from the beginning, without requiring patient-specific data. It does not require manual analysis of the patient's data, and no manual setting of parameters should be done. This approach makes the algorithm more useful for short-term applications. It can also be used as the starting point for long-term applications such as the alarm system or an automated seizure diary, with possible personalization on-the run while more data for patient becomes available (see Chapters 4-7).

3.4.10 Remaining issues

The downside of the approach is however that on a longer time-scale, patient-specific algorithms will lead to a better performance. An intermediate solution might however solve the issues of both patient-independent and patient-specific seizure detection. A possible solution would be to start from a patient-independent algorithm, which can learn the patient-specific characteristics on-the-run (using adaptive learning) when more patient-specific data becomes available. This way, the algorithm will converge to a patient-specific algorithm over time [43]. Another option is to personalize an initial patient-independent algorithm by using a limited amount of (annotated) patient-specific data. Personalization is the focus of the next chapters in this thesis.

Another option to improve the performance for patient-independent algorithms, would be to add extra modalities to the system. Modalities such as accelerometry, electromyogram (EMG) and electrical dermal activity (EDA) have shown to be useful for detecting seizures including movements [7, 11, 31, 35, 126, 150]. This will be illustrated in Chapter 8. Also using ECG morphology features and cardiorespiratory dynamics could possibly enhance the unimodal performance [198], but might lead to an algorithm that is more sensible to ECG noise.

A crucial problem for automated seizure detection using ECG, remains the high noise levels in the ECG. Although an R peak detection algorithm was used that showed similar performance as online state-of-the-art algorithms, still a lot of R peak detections errors were found, both ictal and interictal. These obviously influence the performance of the detector and cause a significant part of the missed seizures and false alarms [43]. Nevertheless, this seems to be unavoidable when acquiring the ECG for every day use. The proposed method tried to cope with this limitation by using the median filter on the tachogram, which filters out a lot of noisy HR measurements, but also includes an extra detection delay to the system (typically $\pm$3-5 s).

3.5 Conclusion

A patient-independent online algorithm using only single-lead ECG was proposed for the automated detection of complex partial and secondary generalized seizures. The method showed a significant performance increase in sensitivity and PPV compared to other patient-independent methods from the literature, but also resulted in a slightly larger detection delay. The selected features from feature extraction showed that the required features are easy to compute in real-time, which makes the proposed algorithm suitable for wearable devices for direct usage in practice. The performance of the proposed unimodal patient-independent seizure detection algorithm should be further improved through either efficient personalization or adding extra modalities to the system. The proposed method of preprocessing are used throughout Chapter 4-7, and the trained classifier procedure leads to the initial classifiers used in Chapter 4-7 as reference classifier to facilitate personalized seizure detection.

Chapter 4

Supervised transfer learning for personalization of ECG-based seizure detection

A patient-independent algorithm for real-time heart rate-based seizure detection was proposed in the previous chapter. Ictal heart rate characteristics are however patient-dependent, which means that the patient-independent algorithm leads to suboptimal results. An efficient procedure for offline personalization of heart rate-based seizure detection is proposed in this chapter. A new supervised weighted transfer learning approach is proposed, which allows to personalize the reference patient-independent classifier from Chapter 3 with a limited amount of annotated patient-specific data.

The dataset used to evaluate the proposed transfer learning approach is presented in Section 4.2. The method for personalizing heart rate-based seizure detection through transfer learning is described in Section 4.3. The results for the proposed algorithm and comparisons with other personalization options are shown and discussed in Section 4.4 and 4.5.

4.1 Introduction to personalization

In the previous chapter, a patient-independent approach was proposed for heart rate-based epileptic seizure detection. The benefit of this approach is that it

does not require to have data from each new subject to be available before using the algorithm, which is useful when it is used in practice as part of a seizure warning system. The previous chapter however showed that there is a lot of variation of the results between the different patients. This is due to the fact that the heart rate characteristics (both epileptic and non-epileptic) are patient-dependent. Personalized seizure detection algorithms could thus improve the performance.

In an ideal case, sufficient patient-specific data is available in order to train a fully patient-specific classifier using only patient-specific data. Such a classifier should however need a large amount of data for that patient, and mainly to collect a sufficient amount of seizures might be a problem for patients with a low seizure frequency. Often only reliable annotated data is available from the hospital, for example during presurgical evaluation. In this case, only data from a couple of days is available. The question is how seizure detection can be personalized with such a limited amount of data.

Several options are possible for personalizing seizure detection algorithms. A first option is the manual setting of some parameter thresholds per patient [19, 142, 188]. This requires manual screening of previous data, and only works well if the parameters are easily understandable for a clinician. Simple thresholding approaches are however too easy to grasp the complexity of the problem. Automated personalization is therefore advised, but normally requires a large amount of patient-specific data in order to find a robust personalization [29]. Often only a limited amount of accurately annotated patient-specific data is available, and thus often a lot of complex approaches are not useful for making fully patient-specific classifiers. An alternative is to put a threshold based on non-epileptic segments from that patient [47, 99], which makes the assumption that epileptic heart rate activity can be seen as an outlier to normal heart rate activity. This is however only the case during the night [48], which makes this approach only sufficiently successful for nocturnal monitoring. Supervised approaches are thus still required for personalizing full day monitoring applications.

A more robust and optimal solution can be found through transfer learning. In transfer learning, the solution to a classification problem is found by using the solution from a similar problem as a starting point. This way, less data for the new problem is required in order to get a robust solution. Transfer learning allows for a better starting performance, faster learning curve and potential better maximal performance as illustrated in Figure 4.1. In this chapter, a patient-specific model is trained through transfer learning by using a patient-independent classifier (trained similarly as in Chapter 3) as a reference model. Therefore, only a limited amount of patient-specific data is required in order to get a robust patient-specific model, obtained with a relative limited complexity.

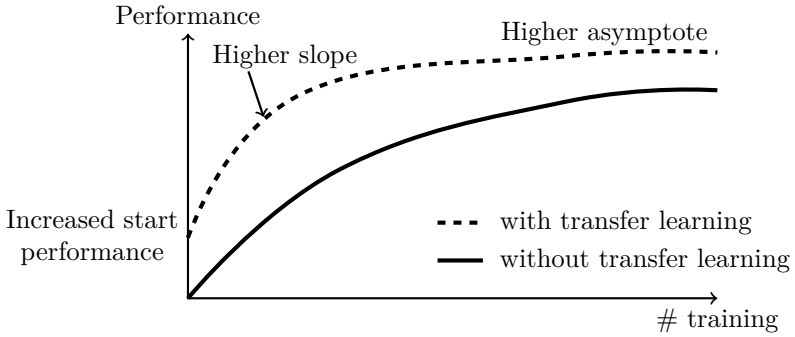


Figure 4.1: Illustration of the effect of transfer learning on the performance compared to traditional machine learning approaches.

4.2 Data Acquisition

The dataset contains recordings from refractory epilepsy patients, who underwent presurgical evaluation at the University Hospitals Leuven (UZ Leuven), Belgium and had during the evaluation at least five FIAS originating from the temporal lobe. The patients were recorded with 10-20 scalp EEG with 1 bipolar ECG channel with a sampling frequency of 250 Hz. The ECG signal was continuously unreadable due to strong noise during 4.6 (patient 13), 8.8 (patient 16) and 7.4 (patient 18) hours of data. Those segments were removed from the analysis. The remaining dataset consists of 24 patients with 2172 hours of data. In total, 228 seizures were recorded (see Table 4.1).

A clinical expert annotated the seizure onsets and ends with the use of video-EEG data, without considering the ECG. Afterwards, a medical doctor validated the annotations. The seizure duration (SD) is defined as the time between EEG seizure onset and EEG seizure end. However, during some seizures, the end could not be determined. The fraction of the determined seizures (Determined S) is also indicated in Table 4.1. The ethical committee of the UZ Leuven approved the study. All patients signed the informed consent for their participation in this study.

4.3 Methodology

The proposed seizure detection algorithm is based on the algorithm described in Chapter 3, and it works for both the patient-independent and the personalized

Table 4.1: An overview of the dataset used in this chapter. RD: Recording Duration, SD: Seizure Duration, B: Bihemispheric, R: right, L: left, Nc: not clear, Te: Temporal, Fr: Frontal, Oc: Occipital.

Patient	**Seizures**	**RD (h)**	**Hemisphere**	**Origin**	**Age**	**Gender**	**Mean SD (s) [Determined S]**	**Range SD (s)**
1	10	26	B	Te	49	M	31 [33%]	[24 39]
2	9	63	L	Fr-Te	41	F	13 [11%]	[13 13]
3	13	71	R	Te	27	M	71 [100%]	[28 96]
4	10	25	B	Te	15	M	19 [30%]	[14 26]
5	11	47	R	Te	29	F	50 [36%]	[40 60]
6	7	148	L	Te	26	M	63 [100%]	[32 116]
7	30	67	R	Fr-Te	9	M	50 [20%]	[17 90]
8	11	114	L	Te	38	M	39 [55%]	[19 75]
9	8	64	L	Te (7), Oc (1)	28	M	23 [63%]	[11 31]
10	7	111	EEG not readable		35	M	126 [29%]	[69 183]
11	6	100	B (5), R (1)	Te	67	F	26 [100%]	[21 31]
12	9	91	R	Fr-Te	24	F	47 [89%]	[33 85]
13	8	109	R	Te	32	M	46 [100%]	[25 61]
14	5	100	L	Te	19	F	56 [40%]	[29 83]
15	13	110	EEG not readable		49	M	16 [94%]	[8 30]
16	5	96	L	Te	45	M	64 [100%]	[11 95]
17	7	102	Nc (4), R (3)	Nc (4), Te (3)	18	F	33 [14%]	[33 33]
18	5	84	L (3), R (2)	Te	62	M	125 [100%]	[89 187]
19	15	113	Nc		40	F	23 [93%]	[6 52]
20	5	113	L	Te	41	F	74 [100%]	[58 83]
21	8	103	L (3), R (5)	Te (3), Fr-Te (5)	43	M	67 [100%]	[29 99]
22	12	115	L (4), R (3), B (5)	Te	35	M	27 [83%]	[17 38]
23	6	101	R	Te	23	F	80 [100%]	[51 108]
24	8	99	R	Te (1), Oc-Te (5)	24	F	42 [100%]	[17 84]
Total	228	2172			[9-67]	14M/10F	50 ± 30 [70 ± 34%]	[6 187]

version. The main difference between both versions lies within the classifier training. Both approaches will be described below.

4.3.1 Preprocessing

Single-lead ECG is used as input for the proposed seizure detection algorithm. First, the heart rate is extracted from the ECG using an approach that uses a real-time R peak detection algorithm. It detects the R peaks based on the derivative signal and an adaptive threshold. Next, strong heart rate increases (HRI) caused by sympathetic activations are detected by automated slope analysis on the tachogram. HRI extraction is performed on a filtered tachogram, using a median filter with an order of 15 heart beats. If a heart rate gradient is larger than 1 bpm/s, a strong HRI is assumed to be detected. The beginning and end of this HRI is found by analyzing when this heart rate gradient becomes negative again. The HRI is then assumed to be a strong HRI if certain thresholds on the length of the HRI, the achieved peak heart rate during the HRI and the (percentual) increase in heart rate during the HRI are exceeded [44] (see Section 3.3.2).

4.3.2 Feature extraction

Features are extracted whenever such a strong HRI is detected in the previous step. In Chapter 3, it was shown that 4 features extracted from the HRI led to optimal patient-independent results: the peak heart rate, the heart rate at the start of the HRI, the baseline heart rate (extracted from the minute before the HRI) and the standard deviation of the baseline heart rate period. As the primary goal of this study is to lead to optimal patient-specific results with a limited amount of patient-specific data, only the first two features are used in this study. The reason for this is that most of the information is already contained in these two features, as discussed in Chapter 3. Adding more features to the system requires more training data for robust personalization through transfer learning. Choosing these two features leads to an optimal balance between performance and limited requirement of patient-specific data. These features are then classified with either the patient-independent (PI) classifier or the patient-specific (PS) transfer learning classifier.

4.3.3 Patient-independent classification

The same procedure for patient-independent classification is used here as in Chapter 3 according to the following steps. Let $\{x_i, y_i\}$ be the training data points extracted from patients different than the ones used for testing the algorithm, with $x_i \in \mathbb{R}^2$ the data samples and $y_i \in \{-1, +1\}$ the corresponding labels. Let class -1 correspond to seizure samples and class +1 to non-seizure samples. Support vector machines (SVM) will map data points to a higher dimension using a (non-)linear transformation $\varphi(x)$, so that the data points can be separated in this space by the hyperplane $w^T\varphi(x) + b$, with w the unknown weight vector and b an unknown constant.

The solution for weighted SVM can be found by solving the following optimization problem (see Section 3.3.5)

$$\begin{aligned} &\min_{w,b,\xi} \frac{1}{2}||w||^2 + C\sum_{i=1}^{N} c_i\xi_i \\ &s.t. \quad \begin{cases} y_i(w^T\varphi(x_i) + b) \geq 1 - \xi_i \\ \xi_i \geq 0 \end{cases} \quad ,\forall i \in [1, N] \end{aligned} \tag{4.1}$$

with

$$c_i = \begin{cases} \gamma\frac{(N^+ + N^-)}{2N^-} & : y_i = -1 \\ \frac{N^+ + N^-}{2N^+} & : y_i = +1 \end{cases} \tag{4.2}$$

Parameter γ gives more importance to the correct classification of seizure samples compared to non-seizure samples during classifier training. In practice, the dual problem is solved, which is defined as

$$\begin{aligned} &min_\alpha = \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} y_i y_j \alpha_i \alpha_j K(x_i, x_j) - \sum_{i=1}^{N}\alpha_i \\ &s.t. \quad \begin{cases} \sum_{i=1}^{N} \alpha_i y_i = 0 \\ 0 \leq \alpha_i \leq C c_i \end{cases} \quad ,\forall i \in [1, N] \end{aligned} \tag{4.3}$$

with $\alpha_i, \nu_i \geq 0$ the multipliers and $K(x_i, x_j) = \varphi(x_i)^T\varphi(x_j)$ the kernel function. The values of the hyperparameter C, γ and the RBF kernel parameter σ are taken as described in Chapter 3. The classifier is trained using the leave-one-patient-out (LOPO) approach, training the classifier on all patients except the one used for evaluating the algorithm.

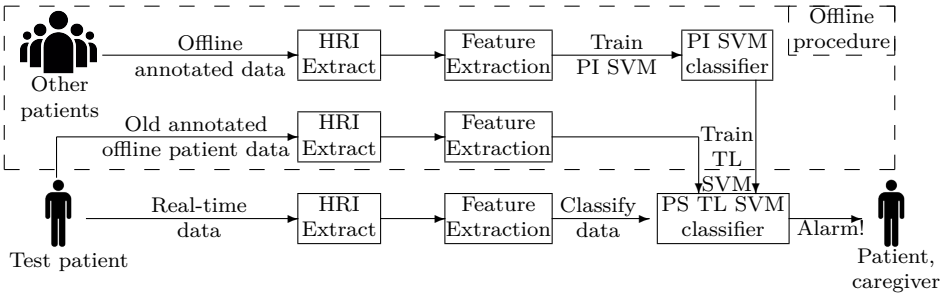


Figure 4.2: Overview of the proposed supervised transfer learning approach for personalized heart rate-based seizure detection. PS TL SVM: patient-specific transfer learning support vector machine; PI SVM: patient-independent support vector machine.

4.3.4 Personalized classification through transfer learning

The goal of this study is to personalize the seizure detection classification in order to get an optimal patient performance with a limited amount of patient-specific data. One solution for this is to use transfer learning. Transfer learning (TL) allows to train a new classifier for a problem with a limited amount of data by using a reference classifier that solves a similar problem. In this case, the previously trained patient-independent classifier discussed in Section 4.3.3 is used as the reference classifier. This way, the classifier can be personalized with a limited amount of patient-specific data by using the knowledge already incorporated in the patient-independent classifier. An overview of the proposed procedure for personalizing the heart rate-based seizure detection is given in Figure 4.2.

One way of applying transfer learning to SVM is to find a weight vector which is of sufficient resemblance compared to the one originating from the reference model, while also minimizing the misclassification error for that specific patient. The following minimization problem to create a support vector machine for the patient-specific data points $\widetilde{x}_k$ with corresponding labels $\widetilde{y}_k$ is solved [207]

$$\begin{aligned} &\min_{\tilde{w},\tilde{b},\tilde{\xi}} \frac{1}{2}||\tilde{w}-w||^2 + D\sum_{k=1}^{M} \tilde{c_k}\tilde{\xi_k} \\ s.t. \quad &\begin{cases} \tilde{y_k}(\tilde{w}^T\varphi(\tilde{x}_k)+\tilde{b}) \geq 1-\tilde{\xi}_k \\ \tilde{\xi}_k \geq 0 \end{cases} \quad ,\forall k \in [1,M] \end{aligned} \tag{4.4}$$

so that the new personalized classifier is defined by $\widetilde{w}$ and $\widetilde{b}$, whereas w is the weight vector obtained from the patient-independent classifier trained using

(4.1) with the patient-independent data x_i, defined as

$$w = \sum_{i=1}^{N} \alpha_i y_i \varphi(x_i) \tag{4.5}$$

by the original SVM optimization problem. The same kernel K used in the reference classifier is used here. The parameters D and $\tilde{c}_k$ are similar as C and c_i in the original patient-independent SVM problem:

$$\tilde{c_k} = \begin{cases} \widetilde{\gamma}\frac{(M^+ + M^-)}{2M^-} & : \tilde{y}_k = -1 \\ \frac{M^+ + M^-}{2M^+} & : \tilde{y}_k = +1 \end{cases},$$

with M^+ and M^- indicating the number of patient-specific non-seizure and seizure data points. Hyperparameter D allows to balance between minimizing the errors for the patient-specific data points and minimizing the difference compared to the reference patient-independent classifier (defined by w). Parameter $\widetilde{\gamma}$ is experimentally initialized to a value of 1.5. The initial value for hyperparameter D depends on the amount of seizures available in the training set, and is set to 0.1 for patients with less than 10 seizures and set to 100 for patients with more than 10 seizures. These values are based on the findings of previous studies [45]. Nevertheless, these parameters should ideally be optimized per patient, based on their validation performance, but it is very challenging to get robust hyperparameter optimization results due to the low amount of patient training data. Therefore, a more heuristic method is applied, which lowers the values of $\widetilde{\gamma}$ (linear decrease of 0.25) and D (exponential decrease by factor 1/10) if the resulting classifier leads to 50% more false alarms than the available patient-independent approach. This reduction is repeated until the FAR is dropped below this threshold or minimal values for $\widetilde{\gamma}$ (=0.25) and D (=0.01) are reached.

An optimal solution is found for the saddle point in the Langragian

$$\begin{aligned} &\max_{\tilde{\alpha},\tilde{\beta}} \min_{\tilde{w},\tilde{b},\tilde{\xi}} \mathcal{L}(\widetilde{w},\tilde{b},\widetilde{\xi};\widetilde{\alpha},\widetilde{\beta}) \\ &= \max_{\tilde{\alpha},\tilde{\beta}} \min_{\tilde{w},\tilde{b},\tilde{\xi}} \frac{1}{2}||\widetilde{w} - w||^2 + D\sum_{k=1}^{M} \widetilde{c}_k \widetilde{\xi}_k \\ &\qquad - \sum_{k=1}^{M} \widetilde{\alpha}_k \left(\tilde{y}_k \left(\widetilde{w}^T \varphi(\widetilde{x}_k) + \tilde{b}\right) - 1 + \widetilde{\xi}_k\right) - \sum_{k=1}^{M} \widetilde{\beta}_k \widetilde{\xi}_k \end{aligned} \tag{4.6}$$

with $\widetilde{\alpha}_k$, $\widetilde{\beta}_k \geq 0$ the multipliers. This then leads to

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \widetilde{w}} = 0 & \rightarrow \tilde{w} = w + \sum_{k=1}^{M} \widetilde{\alpha}_k \tilde{y}_k \varphi\left(\widetilde{x}_k\right) \\ \frac{\partial \mathcal{L}}{\partial \widetilde{b}} = 0 & \rightarrow \sum_{k=1}^{M} \tilde{y}_k \widetilde{\alpha}_k = 0 \\ \frac{\partial \mathcal{L}}{\partial \tilde{\xi}_k} = 0 & \rightarrow \widetilde{\alpha}_k + \widetilde{\beta}_k = D\tilde{c}_k, \forall k \in [1, M] \end{cases} \tag{4.7}$$

with w the weight vector from the reference patient-independent classifier from (4.1), such that

$$\begin{aligned} \widetilde{w} &= w + \sum_{k=1}^{M} \tilde{\alpha}_k \tilde{y}_k \varphi(\tilde{x}_k) \\ &= \sum_{i=1}^{N} \alpha_i y_i \varphi(x_i) + \sum_{k=1}^{M} \tilde{\alpha}_k \tilde{y}_k \varphi(\tilde{x}_k) \end{aligned} \tag{4.8}$$

which indeed indicates that the patient-specific $\tilde{w}$ is a combination of patient-independent and patient-specific information.

The dual problem of (4.4) is defined as

$$\begin{aligned} \min_{\widetilde{\alpha}} &= \sum_{i=1}^{N} \sum_{k=1}^{M} y_i \widetilde{y}_k \alpha_i \widetilde{\alpha}_k K\left(x_i, \widetilde{x}_k\right) - \sum_{k=1}^{M} \widetilde{\alpha}_k + \frac{1}{2} \sum_{k=1}^{M} \sum_{l=1}^{M} \widetilde{y}_k \widetilde{y}_l \widetilde{\alpha}_k \widetilde{\alpha}_l K\left(\widetilde{x}_k, \widetilde{x}_l\right) \\ s.t. \quad & \begin{cases} \sum_{k=1}^{M} \widetilde{\alpha}_k \widetilde{y}_k = 0 \\ 0 \leq \widetilde{\alpha}_k \leq D\tilde{c}_k \quad , \forall k \in [1, M] \end{cases} . \end{aligned} \tag{4.9}$$

A new data point $\tilde{x}_n$ is then classified using

$$y(\tilde{x}_n) = sign\left(\sum_{k=1}^{M} \tilde{\alpha}_k \tilde{y}_k K(\tilde{x}_k, \tilde{x}_n) + \sum_{k=1}^{N} \alpha_k y_k K(x_i, \tilde{x}_n) + \tilde{b} \right) \quad . \tag{4.10}$$

Solving (4.9) leads to the required values of $\widetilde{\alpha}_k$, but parameter $\tilde{b}$ still needs to be defined. This can again done by using the KKT conditions, which state that in an optimal solution, the product of the dual variables and the constraints should be zero. The KKT conditions thus yield

$$\begin{cases} \dfrac{\partial \mathcal{L}}{\partial \tilde{w}} = 0 \rightarrow \tilde{w} = w + \displaystyle\sum_{k=1}^{M} \widetilde{\alpha}_k \widetilde{y}_k \varphi\left(\widetilde{x}_k\right) & \\ \dfrac{\partial \mathcal{L}}{\partial \tilde{b}} = 0 \rightarrow \displaystyle\sum_{k=1}^{M} \tilde{y}_k \widetilde{\alpha}_k = 0 & \\ \dfrac{\partial \mathcal{L}}{\partial \tilde{\xi}_k} = 0 \rightarrow \widetilde{\alpha}_k + \widetilde{\beta}_k = D\tilde{c}_k & ,\forall k \in [1, M] \\ \widetilde{\alpha}_k \left(\tilde{y}_k \left(\widetilde{w}^T \varphi(\widetilde{x}_k) + \tilde{b} \right) - 1 + \widetilde{\xi}_k \right) = 0 & ,\forall k \in [1, M] \\ \widetilde{\beta}_k \widetilde{\xi}_k = 0 & ,\forall k \in [1, M] \\ \widetilde{\alpha}_k \geq 0 & ,\forall k \in [1, M] \\ \widetilde{\beta}_k \geq 0 & ,\forall k \in [1, M] \end{cases} \tag{4.11}$$

Similar as for binary SVM without transfer learning (see Chapter 3), these conditions yield that for a solution $\widetilde{\alpha}_k^* \in (0, D\tilde{c}_k)$, $\widetilde{\xi}_k^* = 0$ so that

$$\tilde{y}_k \left(\widetilde{w}^T \varphi(\widetilde{x}_k) + \tilde{b} \right) - 1 = 0 \qquad \text{for } \widetilde{\alpha}_k \in (0, D\tilde{c}_k) \tag{4.12}$$

Parameter $\tilde{b}$ can then be estimated by taking the average of the values found for the above equation for all $0 < \widetilde{\alpha}_k < D\tilde{c}_k$.

The TL classifier is trained and tested using a 5-fold crosstesting scheme on the patient-specific data, in which 4 folds are used for training and 1 for testing. This is then repeated 5 times so that each fold is used once as test set.

4.3.5 Alternative automatic personalization solutions

The proposed transfer learning approach is also compared to two different alternatives for personalization. The first alternative includes a fully PS approach which is trained with only PS data using the SVM classifier defined by (4.1). The other alternative is a so-called mixed model, in which both PI and PS data are used for training an SVM classifier defined by (4.1), but changing the values of c_i in (4.2) into c_i^{MIX}:

$$c_i^{MIX} = \begin{cases} s_i \gamma \frac{(N^+ + N^-)}{2N^-} & : y_i = -1 \\ s_i \frac{N^+ + N^-}{2N^+} & : y_i = +1 \end{cases} \tag{4.13}$$

with

$$s_i = \begin{cases} 4 & : i \in PSdata \\ 1 & : i \notin PSdata \end{cases} \tag{4.14}$$

such that misclassification of PS data is more critical during training than misclassification of non-PS data. The value 4 is chosen as recommended in [41].

4.3.6 Algorithm evaluation

In order to compare the different seizure detection algorithms, five metrics were used to evaluate the seizure detection performance similarly as in Chapter 3. The sensitivity (Se), false alarm rate (FAR, expressed in false positives/hour, FP/h), positive predictive value (PPV), F_3 and detection delay were calculated. A seizure is detected when an alarm was given 30 s prior to the seizure onset and 90 s after the seizure onset [44]. False alarms within 1 min of each other were counted as one alarm.

Average measures over the entire dataset can be expressed as patient average performance (P-Av), which is the average of the performance of each patient, or overall average (T-Av), computed on the total number of seizures or recording duration. The first average measure is used in this chapter, unless specifically mentioned otherwise. To prove the significant differences between the algorithms, paired t-tests were performed with a significance level of 0.05. All results were obtained retrospectively in a simulation which replicated a real-time setting.

4.3.7 Impact of number of seizures in training

One of the advantages of transfer learning is that it allows to train a new classifier with a relative limited amount of data by using a reference classifier. In one simulation, it is evaluated how many seizures are needed in order to gain sufficient added value in ECG-based seizure detection performance. In this simulation, only a certain amount of seizures from the patient-specific training set are used in training (using the crosstesting scheme described in Section 4.3.4). These were chosen randomly in 100 simulations per selected amount of seizures (tested for 0-4 seizures). This simulation is done for the proposed transfer learning approach and the alternatives for personalization. Testing was performed on the same test sets used in the original evaluation tests.

Table 4.2: Results for the patient-independent (PI), fully patient-specific (PS), mixed (Mix) and transfer learning (TL) approach. Both patient average (P-Av) and overall average (T-Av) are shown in the table.

	Se (%)		FAR (FP/h)		F_3-score	
	P-Av	T-Av	P-Av	T-Av	P-Av	T-Av
PI	**70.75**	**72.81**	3.10	3.01	0.204	0.190
PS	56.07	57.46	2.22	2.27	0.230	0.184
MIX	69.70	71.49	2.66	2.57	0.241	0.209
TL	68.93	71.05	**2.00**	**2.06**	**0.283**	**0.242**

4.4 Results

The preprocessing procedure discussed in Section 4.3.1 detected 84.71% of the seizures, which gives an indication of the amount of seizures with ictal heart rate increases. In Table 4.2, the patient (P-Av) and overall average (T-Av) of Se, FAR and F_3-score are shown for the patient-independent (PI), fully patient-specific (PS), mixed (MIX) and transfer learning (TL) approaches. Figure 4.3 shows the results of the proposed TL approach and the reference PI approach for each patient. The PI algorithm results in an average Se of 70.75% with 3.10 FP/hour and an F_3-score of 0.204. By adapting the model to the patient characteristics with the TL algorithm, a similar Se is observed (68.93%) with 35% less false positives (2.00 FP/hour). The average F_3-score is increased to 0.283. The TL approach did not only result in an average decreased FAR, but also decreased FAR variability over the different patients (see Figure 4.4). A similar effect can be seen for the F_3-score, although there is no decreased F_3-score variability in the TL approach. The alternative mixed approach results in a similar Se, but with on average 0.66 FP/h more than the proposed TL approach. The fully PS approach resulted in a decreased Se, with a slightly increased FAR (2.22 FP/h) than the TL approach.

By performing a two-sided paired t-test, the sensitivity of the PI and MIX algorithm are not statistically different from the TL algorithm (p=0.53 (PI vs. TL) and p=0.8205 (MIX vs. TL)). However, the FAR of the PI and MIX algorithm are statistically different from the TL algorithm ($p< 10^{-3}$ (PI vs. TL), p=0.028 (MIX vs. TL)). The fully PS approach has a significantly lower Se than the TL algorithm (p=0.019), whereas the FARs are not significantly different (p=0.46). The F_3-score of the PI, fully PS and MIX algorithm are statically lower than the TL algorithm (($p< 10^{-3}$ (PI vs TL), p=0.019 (PS vs TL), p=0.0091 (MIX vs TL)). This shows that the proposed transfer learning method is indeed statistically better than the other evaluated approaches for personalization. The average detection delay was 21 s.

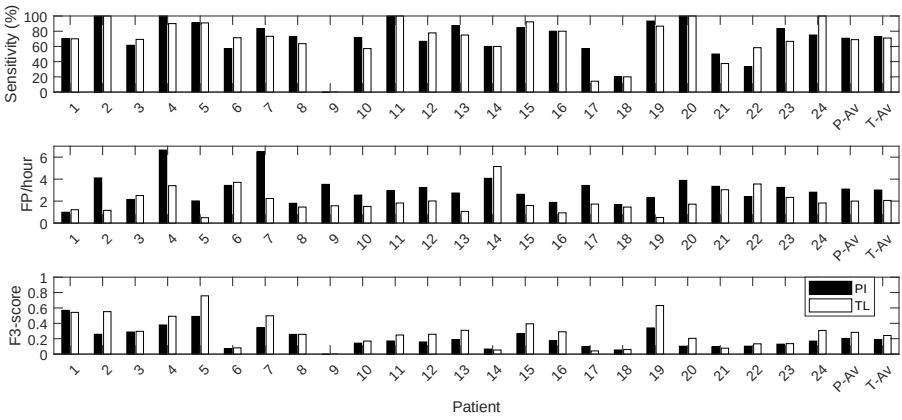


Figure 4.3: Sensitivity, FAR and F_3-score per patient with patient average (P-Av) and overall average (T-Av) performances for the patient-independent (PI) and transfer learning (TL) algorithm

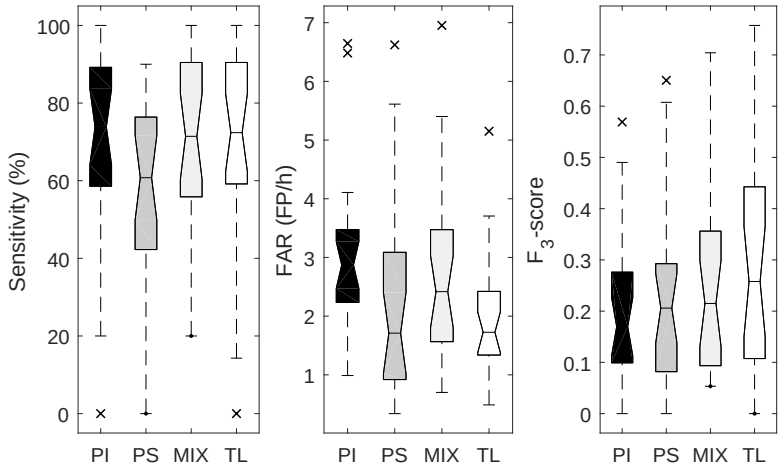


Figure 4.4: Boxplots of sensitivity, FAR and F_3-score for the patient-independent (PI), fully patient-specific (PS), mixed (MIX) and transfer learning (TL) algorithm.

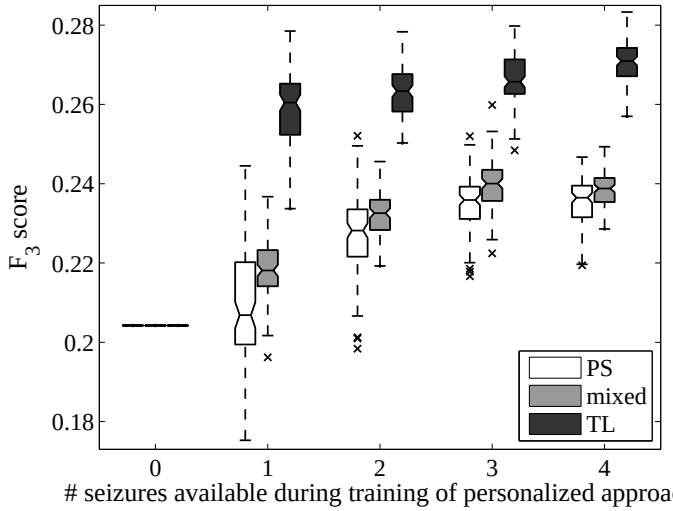


Figure 4.5: Impact of the number of seizures on the F_3-score performance for the full patient-specific (PS) approach, mixed approach and transfer learning (TL) approach. In case 0 seizures are available for training, the performance of the patient-independent algorithm is shown.

As explained in Section 4.3.7, the influence of the number of training seizures on the TL algorithm performance is investigated. Figure 4.5 shows the impact on F_3-score performance in function of the number of seizures available during training. The F_3-score for the TL approach already strongly improves compared to the PI performance if only 1 patient-specific seizure is available in the training set. When 2 seizures are available, the variation between the results decreases, while the average performance increases. The performance further increases by including additional seizures to the training set, while still maintaining a similar variation in results. Figure 4.5 also shows the impact of the number of seizures in the training set on the performance the other personalization methods. These alternatives resulted in clearly lower performances than the proposed transfer learning approach for each number of available seizures during training.

4.5 Discussion

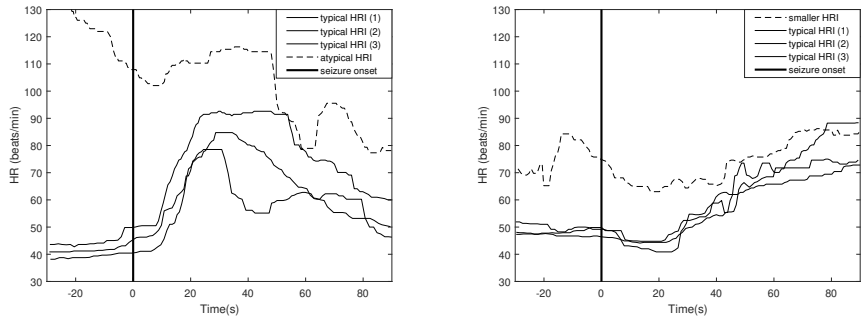
4.5.1 Performance comparison of the PI and TL approach

Table 4.2 and Figure 4.4 show that the mean and median sensitivity of the PI and TL approaches are similar, whereas the FAR decreases and the F_3-score increases. By looking at the patients individually (Figure 4.3), it can be observed that for some patients the TL approach clearly reduced the FAR (e.g. patient 4). Furthermore, for some patients, the sensitivity dropped slightly. This is due to the fact that the model adapts to the patient characteristics. However, some seizures resulted in atypical ictal heart rate behavior for that patient or had a smaller heart rate increase compared to other seizures, which caused these seizures not to be detected with the personalized TL model. An example of atypical ictal heart rate behavior within patient 13 is compared to three typical ictal heart rate changes in Figure 4.6(a). Figure 4.6(b) compares an atypical small heart rate increase for patient 6 a seizure with a three examples of typical heart rate increases for that patient. The proposed TL approach adapts to the majority of seizures, and therefore might lead to a missed detection of these atypical seizures. This small decrease in FAR is however typically accompanied with a strong decrease in FAR, justifying this small sensitivity decrease. The reverse also occurred in some patients, where borderline seizures that were missed by the PI approach are detected by the proposed TL approach (e.g. patient 24).

Not for all patients an increase in performance is observed, for example in patient 1. For those patients, no typical ictal HR increase is observed, or the HRIs are too weak in magnitude, making it very challenging to differentiate them from non-epileptic heart rate activity with the used feature set. If the seizure activity is too similar to non-seizure activity for that patient, the model can not be improved by using personalization. On the other hand, for major seizure types, most often a better differentiation is possible between epileptic and non-epileptic behavior. Overall a better performance increase can be found by personalization in patients with stronger seizures.

Only 84.71% of seizures had ictal HRIs in the analyzed dataset, which is a similar percentage as in the literature [96, 114, 212]. Personalizing the algorithm will not help to detect these seizures without ictal heart rate changes. Also false detections or missed seizures caused by too strong ECG noise can not be avoided by the personalization [191]. Other approaches should be used to further improve the performance, such as further improved noise removal techniques [127].

An example of the obtained classifier boundaries for patient 11 are illustrated



(a) (b)

Figure 4.6: Examples of atypical ictal heart rate behavior within a patient. (a) Example of 3 typical and 1 atypical ictal heart rate patterns for patient 13. (b) Example of a atypical ictal small heart rate increase and 3 normal ictal heart rate increases for patient 6. The vertical line in both figures indicate the seizure onset of the different seizures. The shown tachogram signals are filtered using the median filter discussed in Section 4.3.1.

in Figure 4.7. The seizure and non-seizure data points are shown together with the PI, PS and TL boundary for two different values of hyperparameter D (0.01 and 1), with a fixed $\widetilde{\gamma}$ value of 1.5. The PI boundary gives a good general indication, but lacks adaptation to the patient characteristics. The transfer learning approach adapts to these characteristics with a limited amount of available patient-specific data. By choosing a low value for D (i.e. 0.01) the TL model will be similar to the PI model because a relative low weight is given to the errors obtained for the patient-specific data compared to the first term in (4.4), which quantifies how different the new model is to the original one. With a higher value of D the model adapts more to the patient-specific data (and the corresponding errors $\tilde{\xi}_k$) and shows less similarity to the PI boundary. The PS boundary is less optimal, compared to the TL solution. It can be seen from the TL boundary that it still contains information gathered in the PI classifier (especially for low values of D), indicating the added value of this approach. This way, the FAR is strongly decreased, without affecting the sensitivity.

4.5.2 Comparison of different personalization approaches

Different alternatives for personalization were implemented and evaluated in order to compare with the proposed transfer learning approach (see Table 4.2 and Figure 4.4). The naive approach only uses patient-specific data points for

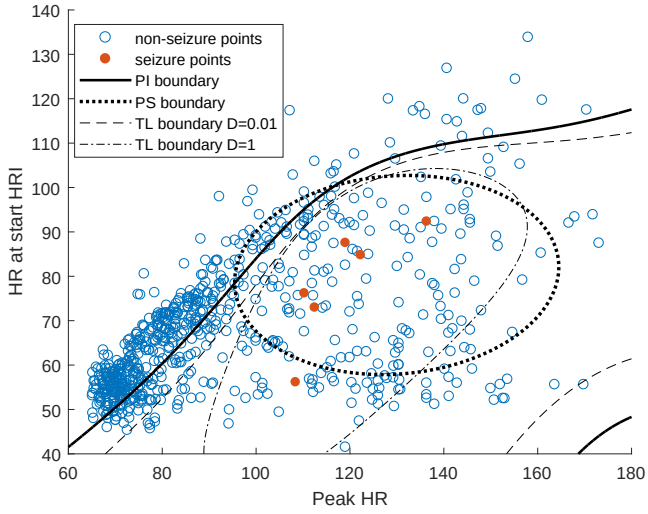


Figure 4.7: Visualization of SVM boundary of the patient-independent (PI), patient-specific (PS) and transfer learning (TL) model for different D values (0.01 and 1).

training a normal SVM classifier. Despite all patients had at least 5 seizures, this was still an insufficient amount of data for most patients in order to get a robust patient-specific classifier. In [28], it was mentioned that at least 6-8 seizures were required in order to make a personal algorithm, and the reason for this was to better include the inter-seizure variability of autonomic changes within a patient. Due to the relative low amount of patient-specific data, it often occurred that the classifier was overtrained on a limited amount of seizure data, not taking into account potential fluctuations between different seizures from a patient. This led to a strong decrease in sensitivity, although the FAR was not so much higher than the proposed TL approach. The TL approach is able to better take this inter-seizure variation into account by holding on to the knowledge described by the reference PI classifier.

The mixed model (MIX) uses a mixture of patient-specific data with data from other patients in the training set of a standard SVM training procedure. It produced more robust results than the fully patient-specific approach without strong negative outliers. However, it generated, on average, around 0.5 FP/h more than the proposed TL approach. Applying transfer learning to the reference classifier allows to better take over the information of the PI classifier and translate it to the patient-specific model, whereas the mixed approach

would try to create a new model without this prior model knowledge. The proposed transfer learning method not only leads to a better performance, but is also trained faster with an optimization problem which contains less data to analyze.

4.5.3 Impact of number of seizures on personalization

Transfer learning allows to train a personalized classifier for heart rate-based seizure detection. It is however important to have an idea of how much data is actually needed for this. In the seizure detection applications, the amount of seizures is typically the restricting factor as some patients might have a low seizure frequency. In Section 4.3.7, a simulation study was described to evaluate the impact of the number of seizures in training on the performance of the personalized approach (see Figure 4.5).

The proposed TL approach already leads to an increase in F_3 performance by only including 1 seizure in the training set. This shows that with only 1 seizure, the algorithm can already be initially personalized. There is a lot of variation in the results, which is caused by which seizure data is used during training. There is also a lot of variation in ictal heart rate changes between different seizures within a patient, so if an atypical seizure is used in training, this will lead to suboptimal results for that patient. By using 2 seizures during training, the F_3 performance increases slightly, and also the variability between results of different simulations decreases. However, from 3 seizures onwards, the variability is no longer strongly decreased with the proposed method. The average performance increases further up until 4 available seizures, and is expected to increase further if more seizures are included in training [28]. From a certain amount of seizures, it is however expected that the fully PS approach would lead to a better performance than the proposed TL method. Nevertheless, it is currently not known how many seizures are required for this with the proposed procedure.

Figure 4.5 also shows the effect of the number of seizures on the alternative personalization approaches. It can be seen that by only having 1 patient-specific seizure in the training data, the TL performance strongly increased, whereas this increase is less evident for the other approaches. For the PS approach, the median performance is increased, but a large fraction of the results were worse than the reference PI performance. This is due to the fact that the PS method is often strongly overfitting on data from 1 seizure, which is not a robust way for training a classifier. The average F_3 scores for the fully PS approach increase by adding more seizures, and the variation on performance decreases. The results however remain lower than those from the proposed TL approach when using

4 seizures. The mixed model is more robust than the fully PS approach for a limited amount of seizures, but it slightly loses its added value when more seizures are added to the training set.

4.5.4 General discussion

Despite the increased performance caused by personalization through transfer learning, still a too high FAR is obtained for usage in practice. Personalization allows to solve some of the issues leading to a too high FAR, but is not able to solve all issues. Some seizures do not contain ictal heart rate changes, and still a large fraction of non-epileptic heart rate changes can not be differentiated from epileptic HRIs with the current techniques. The proposed method should, however, be used as part of a multimodal algorithm, where it is combined with another modality. Accelerometers and EMG sensors could lead to an increased performance for the detection of motor seizures [181, 46]. For the detection of non-motor focal seizures, ear EEG could be used [78]. The advantage of heart rate-based seizure detection is that ECG is often monitored as well during video-EEG monitoring in the hospital, which allows to get accurately annotated heart rate data to personalize the algorithm. The algorithm can then be improved by adding either a patient-independent algorithm using a different modality, or also by applying a similar transfer learning approach to this modality in order to also personalize this algorithm.

The proposed transfer learning approach is a supervised approach, which means that annotated data was required. In practice, these annotations can be made in the hospital during, for example, presurgical evaluation, but they could also be made by the patient or their caregivers/family. Extra procedures should then be added to avoid a too big impact of incorrectly annotated data as patients are not always aware about whether they actually had a seizure or not [41, 59]. Ideally, an unsupervised approach could be used [99], which assumes that epileptic heart rate activity can be seen as an outlier to normal heart rate activity. This is however only the case during the night [48], which makes this approach only sufficiently successful for nocturnal monitoring. Supervised approaches are thus still required for personalizing full day monitoring applications. Due to the supervised approach, still at least one patient-specific seizure is required in order to adapt to the patient characteristics. In patients who have a very low seizure frequency, this might still be a problem for a fast personalization. For these patients, it is then advised to have a large pool of patients, and seizures from patients with similar HRV parameters as the test patient. Then, these seizures could be added to the patient-specific data in order to be able to apply this method.

4.6 Conclusion

Transfer learning allows to personalize ECG-based seizure detection in a fast and robust way by using only a limited amount of annotated patient-specific data. The false alarm rate dropped by 35% compared to the patient-independent approach while maintaining a similar sensitivity on full day recordings. The proposed method can be used as part of a multimodal algorithm in order to increase the performance and close the gap to usage of real-time epileptic seizure warning systems.

Chapter 5

Real-time personalization using patient feedback

In the previous chapter, an offline algorithm for supervised personalization was proposed. From a practical point of view, it is, however, also interesting if this personalization can occur in real-time. This chapter proposes an algorithm that allows real-time personalization using a supervised approach. The initially patient-independent classifier is personalized on-the-run by giving annotations to previously generated alarms for that specific patient. This is done by using a new heuristic updating procedure which allows to update an SVM classifier in real-time with a limited algorithmic complexity compared to normal incremental SVM updates. As in most real-time applications no clinicians are available for accurate annotations, hence, the assumption is made that the patients or the patients' family have to provide the seizure annotations. As these annotations might be inaccurate, this chapter also proposes a new method to limit the effect of potential false annotations based on the transfer learning principles discussed in Chapter 4.

The goal and added value of this chapter is summarized in Section 5.1. The datasets used for this research are summarized in Section 5.2, and the proposed method is described in Section 5.3. The results of the proposed method are shown and discussed in Section 5.4 and 5.5. This method is also described in [41].

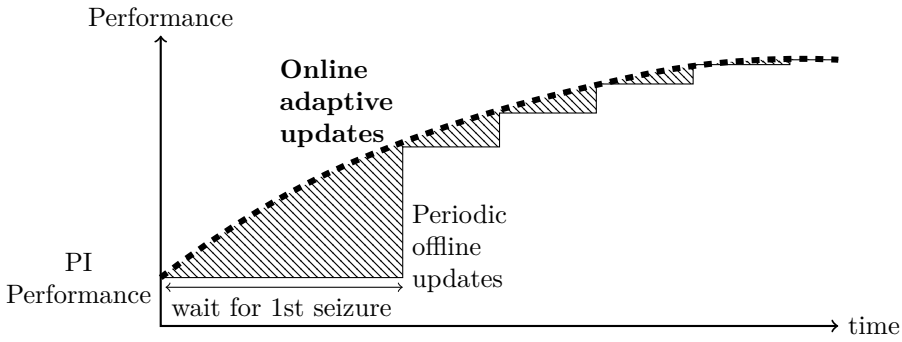


Figure 5.1: Illustration of added value of real-time personalization compared to periodic offline personalization. The online adaptive algorithm allows for a lot of performance increase (illustrated with the marked area) during the different periodic updates, certainly in the beginning of the recordings of each patient.

5.1 Introduction

In the previous chapter, an offline personalized heart rate-based seizure detection algorithm was proposed. It assumes to personalize the algorithm once with a certain amount of patient-specific data, but could also be applied multiple times periodically. The downside of this approach is that it requires the occurrence of at least 1 patient-specific seizure, which might be an issue for patients with a low seizure frequency. Also a lot of time is wasted between periodic updates, whereas the algorithm could immediately be updated with patient-specific data obtained between different updates.

Therefore an algorithm is proposed that allows real-time personalization of heart rate-based seizure detection. The initial patient-independent algorithm immediately gets updated with every new patient-specific data point that is detected by the *HRI-EXTRACT* preprocessing procedure, personalizing the detector on-the-run. This way there is also improvement in performance between the periodic updates, giving a large performance boost during possibly the first days, weeks or even months of patient monitoring. In order to allow real-time personalization updates on wearable devices, the algorithm should be of limited complexity. Therefore a heuristic approach is proposed here, which is easily implementable but on the long-run might lead to suboptimal results compared to the more complex offline personalization procedure proposed in Chapter 3. The added value of a real-time personalization algorithm is illustrated in Figure 5.1.

In this chapter, a supervised method is proposed, so annotations should be available for the new patient-specific data points. Normally annotations are used from the hospital by using the gold vEEG standard. These accurate annotations can only be made by clinicians in a controlled environment, lowering the usability of the system in practice as typically only a limited amount of data can be obtained in such an environment.

Therefore, patient feedback is used here in order to annotate the patient-specific data in real-time. Several currently available seizure warning systems have a prewarning function, which allows the patient to indicate whether a warning might be correct or not before actually warning the relatives. If the patient can annotate it as a false alarm within a certain time limit (i.e. 10 s), it generates no alarm to the relatives, otherwise the relatives are alarmed. This patient feedback can be used as annotations to these alarm data. The algorithm updates itself with this information in real-time, adapting to the patient characteristics over time. As patients might not be aware whether they actually had a seizure or not, extra steps are added to the updating algorithm to increase the robustness against incorrect patient feedback. To our knowledge, this is the first study that discusses real-time adaptive seizure detection using patient feedback during full day monitoring.

5.2 Data Acquisition

Two different datasets are used in this study. The first dataset is the dataset described in Section 3.2, which contains data obtained from the University Hospital of Leuven (UZ Leuven), consisting of 17 patients with TLE and 127 complex partial seizures during in total 918 hours of recording. This dataset is only used for training as will be described in Section 5.3.

The second dataset is the publicly available EPILEPSIAE database, which is used for the evaluation of the different algorithms [94]. It contains data from epilepsy centra from Freiburg (Germany), Coimbra (Portugal) and Paris (France). Every patient has at least 5 clinical seizures during recordings of at least 92 hours per patient. In total 19 patients with TLE are used from the EPILEPSIAE dataset containing scalp EEG data, including in total 158 clinical seizures in 2833.4 hours of single-lead electrocardiogram (ECG) data. Five seizure annotations were not used as the ECG was temporarily turned off just before, during or just after the seizure. ECG signals were measured on the chest with a sampling frequency of 256 Hz. Seizures were annotated by expert neurologists using the gold video-EEG standard, and grouped in five categories: simple partial (SP, currently called focal aware in the new ILAE

classification scheme), complex partial (CP, now focal impaired awareness), secondary generalized (SG, now focal to bilateral tonic-clonic), unclassified (UC) and subclinical (SC, now focal non-motoric) seizures. Only the clinical seizures are examined here, as these are of most interest to detect in a home environment. The EEG seizure onsets are used as reference if available, otherwise the clinical onsets are used. An overview of the used datasets is given in Table 3.1 and 5.1.

5.3 Methodology

Three methodologies are evaluated here: a patient-independent, a patient-specific and an adaptive approach. These approaches will be explained in more detail below. All patient-specific data originates from the EPILEPSIAE database and all patient-independent data originates from the UZ Leuven database. The same preprocessing steps are used for all methodologies, and only the classification procedure is different in each method.

5.3.1 Patient-independent seizure detection

The first method is the patient-independent (PI) algorithm. No patient-specific data from the patient under testing are used in order to train the classifier of this algorithm. Training is done using only data from other patients (from the UZ Leuven dataset). This type of classifier can be used from the beginning of the monitoring period as it does not require prior patient-specific data, making it very useful in practice.

The same preprocessing method used in Chapter 3 is implemented here according to the following steps. The heart rate is extracted from the acquired ECG signal by using a real-time R peak detection algorithm based on a dynamic threshold on the derivative signal. Next, *HRI-EXTRACT* (see Section 3.3.2) extracts strong heart rate increases from the heart rate signal, of which then features are extracted. Unlike in Chapter 3, different features are extracted here: the maximal heart rate gradient ∇HR_{max}, HRI duration Δt_{HRI} and the pre-HRI spectral entropy showed to be the 3 best features to use for automated PI seizure detection with a linear support vector machine (SVM) classifier [38] (see Appendix A for the definitions of these features). This is preferred here to allow a heuristic classifier that can be easily updated in real-time with limited computational power (see Section 5.3.3). Therefore only these features are used for classification.

Table 5.1: Overview of the used evaluation dataset which contains a subset of the EPILEPSIAE database.

patient number	patient code	Gender (M/F)	patient age	total # seizures	# SP‡ seizures	# CP‡ seizures	# SG‡ seizures	# UC‡ seizures	record duration†
1	586	M	32	22	1	17	0	4	161:08:36
2	564	M	47	7	0	1	2	4	178:12:50
3	552	F	17	9	2	4	0	3	159:05:35
4	110	M	41	8	2	4	0	2	143:15:45
5	162	F	46	7	0	5	0	2	177:24:22
6	219	M	47	6	0	5	0	1	93:57:37
7	239	M	36	5	5	0	0	0	118:06:35
8	261	M	65	8	0	8	0	0	135:24:47
9	308	M	28	9	9	0	0	0	159:43:47
10	591	M	47	7	4	3	0	0	164:35:01
11	327	F	62	6	0	6	0	0	162:10:05
12	454	F	41	5	2	3	0	0	118:43:48
13	89	F	67	5	0	4	0	1	166:47:26
14	752	M	13	7	5	0	0	2	159:02:23
15	939	M	50	9	1	6	1	1	237:16:25
16	852	F	54	9	0	6	0	3	94:26:47
17	960	M	58	8	0	7	0	1	151:03:52
18	1095	M	50	10	0	8	0	2	160:04:34
19	467	F	15	6	1	4	1	0	92:53:47
total	-	12M/7F	[15-67]	153	32	91	4	26	2833:24:06

† record duration expressed in hhh:mm:ss.

‡ SP: simple partial, CP: complex partial, SG: secundary generalized, UC: unclassified.

Similarly as in Chapter 3, a weighted SVM is used, which can be found by solving the following optimization problem

$$\begin{aligned} &\min_{w,b,\xi} \frac{1}{2}||w||^2 + C\sum_{i=1}^{N} c_i\xi_i \\ &s.t. \quad \begin{cases} y_i(w^T\varphi(x_i)+b) \geq 1-\xi_i \\ \xi_i \geq 0 \end{cases} \quad ,\forall i \in [1,N] \end{aligned} \tag{5.1}$$

with C the hyperparameter, c_i the weights of each data point x_i with corresponding ground truth y_i and errors ξ_i. The classifier output $\widetilde{y}_j$ of a new data point x_j is then defined as $\widetilde{y}_j = sign(w^T x_j + b)$. The classifier training procedure depends on the different methods described in the following sections. The P-I classifier is trained on the entire UZ Leuven dataset described in Table 3.1.

5.3.2 Patient-specific seizure detection

The (semi-)patient-specific (PS) approach follows a similar procedure as the PI method: The same preprocessing procedure and features are used here as in the PI approach, only a different classifier is used. This classifier indicates what the maximal performance could be for the adaptive approach (see Section 5.3.3).

Because not enough PS data is available in the EPILEPSIAE database to create a fully PS classifier, a mixed model is used. Data from a specific test patient (from the EPILEPSIAE database) is mixed with data from other patients (from the UZ Leuven database) in order to train the classifier. Misclassification of PS data should have more impact on the optimization function, so (5.1) is rewritten as

$$\min_{w,b,\xi} \frac{1}{2}||w||^2 + C\sum_{i=1}^{N} \omega_i c_i\xi_i \quad \text{, with} \quad \omega_i = \begin{cases} \kappa & : i \in PS \\ 1 & : i \notin PS \end{cases} \tag{5.2}$$

where PS is the collection containing the indices of the PS data points in the collected dataset. This way, more importance is given to the PS data points with a factor κ, which is set automatically by using leave-one-patient-out (LOPO) crossvalidation.

Training and testing of the PS classifier is done using 5-fold crosstesting. The 4 folds of PS data for training are mixed with PI data from the UZ Leuven dataset to train the classifier, and is then evaluated on the fifth fold.

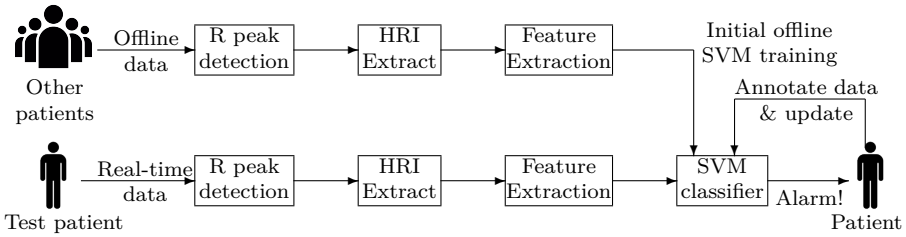


Figure 5.2: Overview of the proposed real-time adaptive seizure detector.

5.3.3 Real-time adaptive seizure detection

The adaptive seizure detection classifier differs from the PS classifier due to real-time adaptation to the patient-specific characteristics, whereas the PS classifier mentioned in Section 5.3.2 is trained offline with PS data. The adaptive algorithm originally starts with the PI classifier described in Section 5.3.1. For each new PS data point (i.e. each HRI detected by *HRI-EXTRACT* extracted in the EPILEPSIAE database) that is detected, the adaptive classifier is immediately updated (see Figure 5.2). Two SVM parameters are updated here: the weight matrix w and the bias term b. From now on, w_t and b_t ($t \geq 1$) indicate the t-th value of w and b after adding data point x_t to the classifier. The initial values w_0 and b_0 are the parameter values obtained by the PI classifier described in Section 5.3.1. Although there exists an incremental way to update an SVM [148], this method is too complex for real-time updating. Therefore, a simplified heuristic updating approach is used, which adapts to the PS data equally fast and with lower complexity. This approach also includes procedures in order to lower the impact of incorrect user feedback, which is required if the system is used in practice.

For each new data point, it is first evaluated whether this data point is possibly caused by noise. This can occur if a noisy ECG segment leads to R peak detection errors. This kind of data should not cause a classifier update as it does not represent the real patient's characteristics. A data point is assumed to be noise-free if less than 10% of all immediate heart rate values in the HRI differ more than 20% compared to the previous immediate heart rate value. This metric thus gives an indication of possible R peak detection errors. Ectopic heartbeats during a HRI happen only rarely during ictal HRIs, so that this procedure did not lead to missing any ictal HRI.

Updating w

If the new data point x_t is assumed to be noise-free, w_t can be updated. In SVM, w can be written in the dual formulation as [193]

$$w = \sum_{i=1}^{N} \alpha_i y_i x_i. \tag{5.3}$$

To lower the complexity of the updating procedure, the primal SVM problem is further evaluated in the proposed approach. That way, the PS α_i (support vector values) are not required, the same principle is used. Let w_0 be the weight vector from the PI classifier. If a PS data point x_t causes a classifier update, w_t^{nn} ("nn" standing for "not normalized") is set to

$$w_t^{nn} = w_{t-1}^{nn} + \tau t x_t y_t \qquad \text{with } \tau_t = \begin{cases} 1 & : y_t = -1 \\ 0.1 & : y_t = +1 \end{cases} \tag{5.4}$$

and $w_0^{nn} = w_0$. The τ parameter replaces the original α parameter. For simplicity, τ_t can only take 2 values, which are chosen to have a similar ratio as the maximum values the α values can get in the PI classifier (see equation (3.15), $\alpha_i \leq c_i$). That way the τ parameters get larger values for the seizure samples as well. Normally only data points causing an error have $\alpha_i \neq 0$ in SVMs. Similarly, the update of w_t only occurs here if $|w_{t-1}^T x_t + b_{t-1}| \leq 1$. This way, also strong outliers will not cause an update of w, increasing the robustness of the SVM updates.

To optimally benefit from the knowledge from the PI classifier, the adaptive classifier should not deviate too much from it. It can be assumed that each PS classifier shows sufficient resemblance with the PI classifier. Therefore, w_t is restricted so that its norm does not change more than σ compared to w_0 ($||w_t - w_0|| < \sigma$):

$$w_t = w_0 + \frac{\sigma}{||w_t^{nn} - w_0||}(w_t^{nn} - w_0). \tag{5.5}$$

This is an assumption that is also used in transfer learning approaches [207], which benefit optimally from a reference classifier (here the PI classifier) in order to get a stable classifier with limited training data. By including it in the adaptive approach, the classifier remains more stable and reliable during all classifier updates. The value of σ is automatically set for each patient using LOPO crossvalidation on all patients of the EPILEPSIAE database except the test patient. If no update occurs after x_t, w_t is set to w_{t-1}.

Updating b

Once w is updated, the bias term b also needs to be updated. b_t can be found by choosing the value that leads to the minimal sum of errors ξ_i, with

$$\xi_i = max(y_i(w_t^T x_i + b'), 0), 1 \leq i \leq N. \tag{5.6}$$

The errors ξ_i are only evaluated for $b' \in [b_{t-1} - \epsilon, b_{t-1} + \epsilon]$ to lower the complexity, preferring a smooth transition from PI towards PS classification. Parameter ϵ should be a small value to maintain a smooth classifier transition (empirically set as 0.1). The errors are only evaluated for a fixed number of data points N, which is a mixture of PI and PS data. Every time a new PS data point is annotated, a PI data point from the same class is removed from both memory and computation. This is done in order to limit the computation and memory requirements of the updating procedure.

Let $E_{b'}$ be the averaged error caused by the constant term b'

$$E_{b'} = \frac{\sum_{i=1|y_i=+1}^{N} \omega_i \xi_i}{\sum_{i=1|y_i=+1}^{N} \omega_i} + 1.5 * \frac{\sum_{i=1|y_i=-1}^{N} \omega_i \xi_i}{\sum_{i=1|y_i=-1}^{N} \omega_i} \tag{5.7}$$

The weights ω_i give extra weight to the PS data points as in (5.2). The class imbalance is removed by using weighted averages of both classes. The experimentally found factor 1.5 gives seizure data more impact on the b update. The new b_t is then set as

$$b_t = \arg\min_{b'} E_{b'} \tag{5.8}$$

After this update, the classification rule for the next data point x_{t+1} then becomes

$$\tilde{y}(x_{t+1}) = sign(w_t^T x_{t+1} + b_t). \tag{5.9}$$

The usage of annotations in the updating procedure has two main drawbacks. First, the annotations made by the user might be incorrect, leading to bad classifier updates. Secondly, short seizures are too hard to detect with heart rate-based seizure detection as they lead to very short HRIs, making it hard to differentiate them from non-epileptic HRIs. They are often also of less importance for real-time seizure detection applications due to their low risk for injuries. Therefore they should not be given the same importance as longer seizures during the update procedure, certainly because the adaptive approach might base its model with just 1 PS seizure. In order to remove the impact

Algorithm 5.1: Procedure for adaptive seizure detection

Input: Online HR data, initial classifier $SVM(w_0, b_0)$ and corresponding patient-independent data
Output: Alarm in case of an assumed seizure
1: Set D to the available patient-independent data
2: Set indices of patient-specific datapoints PS as empty vector
3: **for** each i-th HRI detected with *HRI-EXTRACT* **do**
4: Extract features x_t from the HRI
5: Classify features using $sign(w_{t-1}^T x_t + b_{t-1})$
6: Alarm in case classified into group -1
7: Give annotation y_t to x_t using user feedback
8: Obtain w_t and b_t and update D and PS according to Algorithm 5.2
9: **end for**

of these issues, an extra procedure is able to detect most of these smaller or incorrectly annotated ictal HRIs. If the output $w_t^T x_i + b_t$ (with annotation $y_i = -1$) of x_i (PS data point) is larger than 20% of the outputs from the already collected PS non-seizure data points, ω_i is set to 0 in equation (5.7) in order to remove it from updating, as it is unclear whether it really represents seizure data, or might be an atypical seizure for that patient. In this case, the data point annotated as a seizure closely resembles non-epileptic PS data, and is therefore unlikely to be accurately detectable with the proposed method. As a result, it is removed from the updating procedure. Incorrectly annotated data should not cause a classifier update, and too short seizures leading to too much false alarms after a potential update should not cause a drastic change in the classifier either. The overall procedure for adaptive seizure detection and SVM updating are summarized in Algorithms 5.1 and 5.2.

5.3.4 Feedback procedure

The real-time SVM updates require the availability of the annotations y_t of each PS data point x_t. They are made by the patients or their caregivers for all PS HRIs that are detected by *HRI-EXTRACT*, not only the HRIs that cause an alarm. The feedback procedure is assumed to go as follows:

- An alarm is generated: The user gives feedback about the alarm correctness. If no response is given within a given time period (i.e. 10 s), the alarm is assumed to be correct as the patient is not able to cancel the alarm during a seizure.

Algorithm 5.2: Procedure for SVM updating

Input: New patient-specific data point x_t, corresponding ground truth y_t, previous classifier SVM_{t-1}, old D and PS matrices
Output: Updated classifier SVM_t with w_t and b_t, D and PS

if x_t is assumed to be noise-free
  Insert x_t in D and PS
  if $|w_{t-1}^T x_t + b_{t-1}| \leq 1$
    $w_t^{nn} = w_{t-1}^{nn} + \gamma_t y_t x_t$
    if $||w_t^{nn} - w_0|| < \sigma$
      $w_t = w_t^{nn}$
    else
      $w_t = w_0 + \frac{\sigma}{||w_t^{nn} - w_0||}(w_t^{nn} - w_0)$
    end if
  else
    $w_t = w_{t-1}$, $w_t^{nn} = w_{t-1}^{nn}$
  end if
  Compute $E_{b'}$ on all data from D for all $b' \in [b_{t-1} - \epsilon, b_{t-1} + \epsilon]$ (see (5.7))
  $b_t = \arg\min_{b'} E_{b'}$
else
  $w_t = w_{t-1}$ and $b_t = b_{t-1}$
end if

- A seizure is missed: The user can indicate that a seizure is recently missed. In this case, the algorithm can keep information about the HRIs of the last hour. These HRIs are then marked as potential seizure HRIs. The procedure to remove falsely annotated seizure HRIs eliminates most of the falsely annotated seizure HRIs for SVM updating. If no such ictal HRI is found, no classifier update occurs.

- A non-seizure HRI is correctly ignored: If the procedure for a missed seizure is not activated, the HRI is annotated as non-seizure and the classifier is updated.

5.3.5 Evaluation criteria and additional tests

All classifiers are evaluated only on the EPILEPSIAE database. The same metrics (false alarm rate (FAR), sensitivity (Se), positive predictive value (PPV) and detection delay) are used as in Chapter 3. A seizure is said to be detected if an alarm is generated between 60 s prior to the seizure onset and 90 s after the seizure onset [44]. Different alarms within one minute are clustered as one alarm. The patient feedback is assumed to be 100% accurate unless specifically mentioned otherwise. To evaluate the impact of incorrect feedback, an additional

Table 5.2: Overview of the results of the different seizure detection classifiers.

Method	Se (%)	FP/h	PPV (%)
Preprocessing procedure (*HRI-EXTRACT*)	84.31	3.63	1.24
Patient-independent approach	78.43	1.73	2.39
Patient-specific approach	76.47	1.09	3.67
Adaptive approach	77.12	1.24	3.25

test is performed when a certain percentage of user feedback is incorrect. For each percentage, 40 runs were simulated in which PS data was randomly chosen to be incorrectly annotated with the corresponding probability. In order to prove that the adaptive classifier performs better than the PI classifier, paired t-tests (with corresponding confidence intervals (CI)) are performed on Se, FAR and PPV.

5.4 Results

5.4.1 Performance of the different seizure detection classifiers

Table 5.2 gives an overview of the results of the different real-time seizure detection approaches on the EPILEPSIAE database. 84.31% of the seizures were detected with the preprocessing procedure *HRI-EXTRACT*, which indicates the percentage of seizures with ictal strong HRIs. The PI classifier resulted in a sensitivity of 78.43% with on average 1.73 FP/h. Seizures were detected on average 18s after seizure onset. The PS approach resulted in only 1.09 FP/h with a similar sensitivity performance. By using the adaptive approach, only on average 1.24 FP/h were found, again with a similar sensitivity as the PS approach. This is a drop of nearly 30% of false alarms compared to the PI classifier. Figure 5.3 shows the results per patient for the 3 approaches. The results for the false alarm rate (FAR) ($p<10^{-3}$, CI [-0.81,-0.26]) and PPV ($p<10^{-4}$, CI [0.74%,1.57%]) from the adaptive classifier prove to be significantly better than those of the PI classifier, but not for Se (p=0.38, CI [-3.71%,1.48%]). Both PS and adaptive classifier resulted in an average detection delay of 19s.

The purpose of the adaptive approach is to adapt to the PS characteristics over time. Figure 5.4(a) shows the average FAR per 12 hours for the three approaches for two patients. For patient 1, the adaptive classifier starts with a FAR close to the PI classifier. When more data is collected from the test patient, the results start to converge towards PS performance, which is accomplished around 5 days after the start of the recording. In case of patient 10, a similar

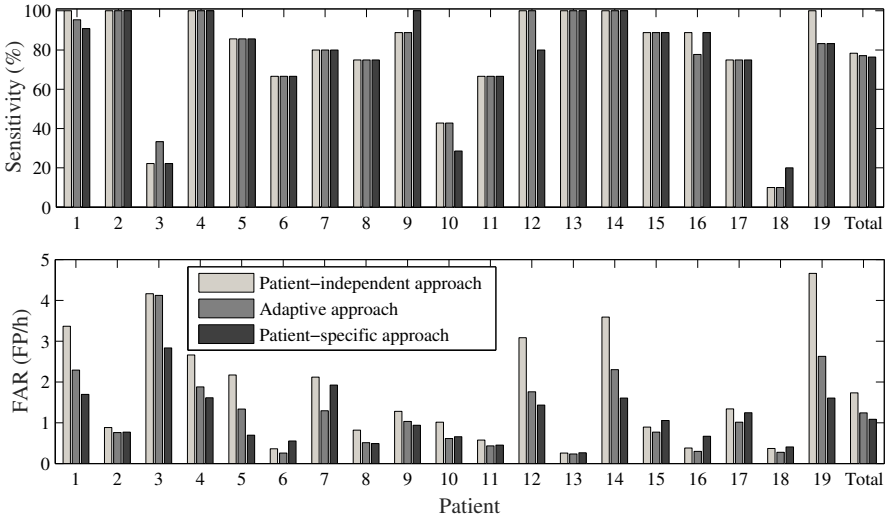


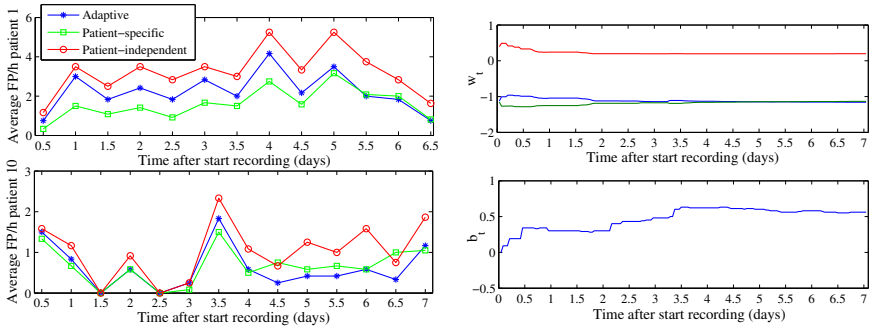
Figure 5.3: Sensitivity and FAR results per patient for the different classifiers.

pattern can be seen, but the adaptive classifier is able to show even a lower FAR than the PS version after 4 days. Figure 5.5 shows a similar effect with the FAR averaged over all patients, illustrating that on average also improvement can be found over time. Figure 5.4(b) shows the SVM parameter convergence for patient 10, which reaches a nearly stable SVM classifier after 4 days of recording. The time of convergence typically depends on when sufficient PS seizure data was collected.

5.4.2 Seizure sensitivity

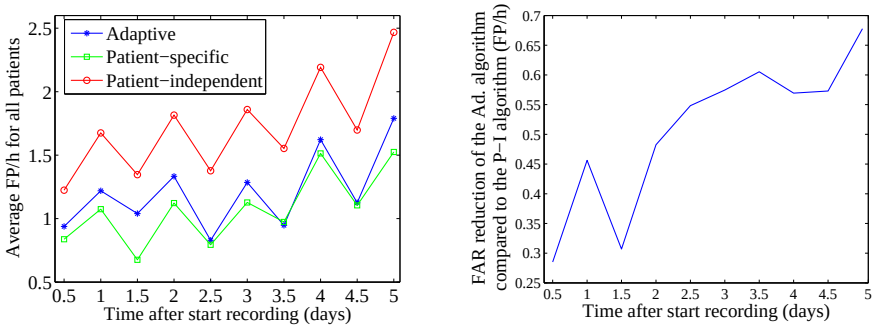
Figure 5.6(a) gives an overview of the percentage of detected seizures for the different seizure types for both *HRI-EXTRACT* and the adaptive seizure detection approach. In the results for *HRI-EXTRACT*, it can be seen that $\pm 84\%$ of the SP, CP and UC seizures have significant ictal HRIs. All 4 generalized seizures show ictal strong HRIs.

The overall sensitivity however dropped by using the adaptive classifier. Mostly extra SP and UC seizures are missed by this procedure, resulting in a bigger difference in sensitivity for the CP and SG seizures compared to the SP and UC seizures.



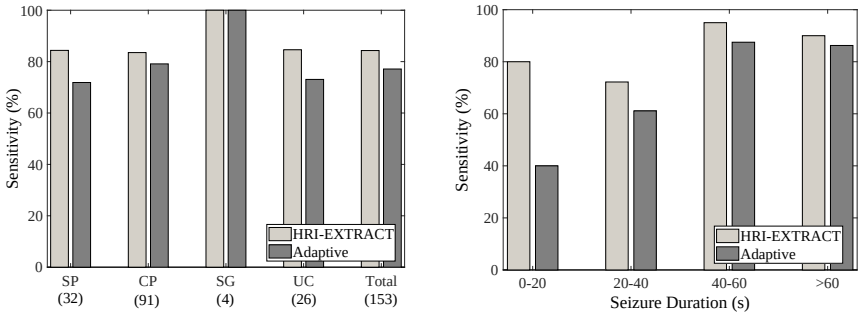
(a) (b)

Figure 5.4: Illustration of convergence for the adaptive classifier. (a) Average FP/h per 12 hours of recording in function of time for patients 1 and 10 for the three different approaches. (b) Example of the SVM parameter convergence for patient 10. The different lines in the w_t figure indicate the terms for the 3 features.



(a) (b)

Figure 5.5: (a)Effect of time after the start of recordings on the false alarm rate (FAR) of the patient-independent (PI), patient-specific (PS) and adaptive (Ad.) algorithm averaged over all patients and per 12 hours for the first 5 days of data of each patient. (b) Reduction in false alarm rate for the adaptive algorithm compared to patient-independent algorithm in function of time after the start of the recording (averaged over all patients). The rhythmic behavior of the FAR results in (a) are caused by synchronized day-night differences in all recordings, causing large differences in FAR rates between two subsequent averages.



(a) (b)

Figure 5.6: Impact factors on seizure sensitivity for *HRI-EXTRACT* and the proposed adaptive seizure detection classifier. (a) Impact of seizure type on sensitivity. The numbers between brackets indicate the number of seizures per type. (b) Impact of seizure duration on sensitivity.

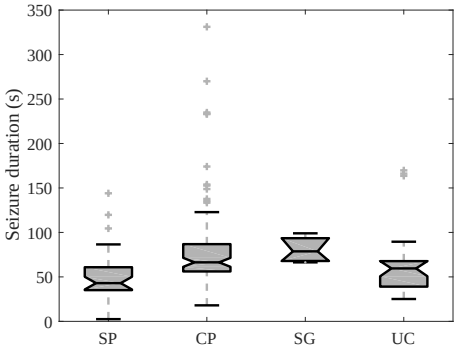


Figure 5.7: Boxplots of the seizure duration for each seizure type analyzed in the dataset.

Seizure duration also has an impact on the sensitivity. Figure 5.6(b) shows the sensitivity of the seizures based on their seizure duration for *HRI-EXTRACT* and the adaptive approach. The longer the seizure lasts, the higher the chance of detecting it with the adaptive classifier. The extra seizures that are missed in comparison with *HRI-EXTRACT* are mostly relatively short seizures ($< 60\ s$).

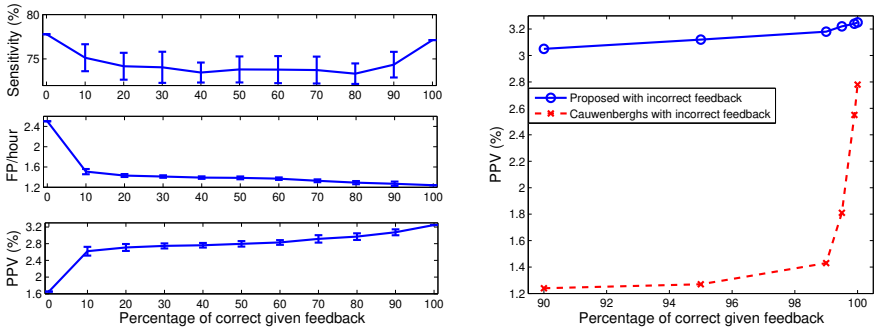
5.4.3 Impact of incorrect user feedback on real-time adaptive seizure detection

For now the assumption was made that all given user feedback was correct. In practice, however, this will not be the case. Figure 5.8(a) shows the effect if only a certain percentage of user feedback is correct. From 10% correct feedback on, the PPV value increases slowly towards optimal PPV performance. The sensitivity does not change much for different feedback percentages, but mainly the FAR changes drastically. However, the overall effect of incorrect user feedback is rather limited. Figure 5.8(b) shows the impact of incorrect user feedback for both the proposed updating procedure and the standard updating procedure [148]. If the standard SVM updating procedure is used with 100% correct feedback, 78.43% sensitivity is achieved with 1.48 FP/h and 2.78% PPV, which is 0.24 FP/h more and 0.47% PPV less than with the proposed updating method. The performance of the standard updating procedure decreases very fast with only 0.1% incorrect user feedback, and with 10% incorrect feedback the PPV is already equally bad as for the preprocessing procedure *HRI-EXTRACT*. The proposed method however only shows a small drop in PPV here, with still a PPV of 3.05% at 10% incorrect user feedback.

5.5 Discussion

5.5.1 Comparison between the different approaches

If the PI classifier is retrained on the UZ Leuven dataset in order to get the same sensitivity as the adaptive classifier on the validation dataset, the PI classifier results in 77.78% sensitivity with 1.63 FP/h on the EPILEPSIAE database. This is an increase of almost 0.4 FP/h with a nearly identical sensitivity compared to the adaptive classifier. The adaptive classifier also shows results close to the assumed maximal PS performance of the adaptive classifier, leading to on average only 0.15 FP/h more than the PS performance. The adaptive approach indeed shows results that come close to PS optimal performance after less than one week of recorded data with around 3-5 seizures. This shows that the used procedure to adapt to the patient characteristics leads to a nearly as efficient classifier as the PS classifier, which is normally adapted offline by a neurologist or nurse. This procedure is thus more time and cost effective than the standard way of personalization. The proposed annotation procedure does not cause any extra burden to the patient if it is added to a system that already uses prewarnings. This adaptive approach just takes into account information that is already available in such systems, and will also lead to a lower FAR in practice.



(a) (b)

Figure 5.8: Impact of incorrect feedback on the performance. (a) Effect of different percentages of incorrect user feedback on sensitivity, FP/hour rate and PPV for the proposed adaptive seizure detection approach. The variation in performance between different runs is shown by the variation bars. (b) Comparison of the PPV values for the proposed updating procedure and the traditional SVM updating procedure in case 0, 0.1, 0.5, 1, 5 and 10% of user feedback is incorrect.

Figure 5.3 shows the results for the different approaches for every patient. In most cases the FAR for the adaptive classifier lies between the FAR of the PI and PS classifiers with similar sensitivity results. For some patients the FAR is even lower for the adaptive classifier than for the PS classifier. The reason for this is that the PS classifier does not have a procedure to remove short or abnormal seizures from the training set, whereas such a procedure was implemented for the adaptive approach. This procedure removed seizures that were unlikely to be detected from the training set, as adding them to the classifier would lead to an underestimation of the SVM boundary, leading to too many false alarms than strictly necessary for that patient. Normally atypical seizure HRIs would only have limited effect during SVM training if sufficient PS seizures are already collected, but in the incremental approach it can occur that only one (abnormal) seizure might be available for a long time. By making this choice based on the collected non-seizure HRI data, this decision can be made much faster with sufficient accuracy. A second cause for this is that by assuming that the adaptive classifier should have a w_t sufficiently similar to the PI w_0, a better classifier can be achieved with the adaptive classifier. In some cases, training the PS classifier led to strong changes in w compared to the PI w_0, which often resulted in a slightly worse performance. An example of this can also be seen in Figure 5.4(a), for which after 4 days the adaptive classifier becomes better than the PS classifier due to overtraining of PS classifier on

several SP seizures which are too hard to detect.

5.5.2 Seizure sensitivity

Figure 5.6 shows that the sensitivity changed for both seizure duration and seizure type. In order for a seizure to be accurately detectable using the heart rate, the seizure should indeed last long enough, so that the seizure HRIs can be sufficiently different from non-seizure HRIs [44, 83]. Both reasons are, however, correlated in this database as shown in Figure 5.7. Both CP and SG seizures were on average longer than SP and UC seizures, making them easier to detect with heart rate-based seizure detection. This is not only applicable between patients, but also within patients. An example of this can be found in patient 10, for which all CP seizures are detected, but all SP seizures were missed.

There is not only variability between seizure types, but also between patients. Most patients show to have sensitivities of 70-100%, but for 2 patients the sensitivity was lower than 40% for all approaches. For patient 3, all ictal HRIs showed to be insufficiently strong compared to inter-ictal HRIs, causing all seizures to be removed from the update procedure as they were marked as too short. Therefore this patient has a similar high FAR as the PI approach. Patient 18 often had ictal bradycardia (sometimes following a small HRI), with an ictal ideoventricular rhythm, which is a potentially life threatening heart condition. Although these seizures were not detected by the algorithm due to the lack of a strong HRI, it shows the benefits of monitoring the heart rate in epileptic patients in a home environment, as heart rate monitoring allows to detect dangerous heart conditions (if additional algorithms are applied) [186].

5.5.3 Impact of incorrect user feedback on performance

Patients might not be aware that they had a seizure, or might not be certain about it. If they do not have somebody accompanying them to indicate whether they indeed had a seizure or not, this could potentially lead to incorrect user feedback.

The impact of incorrect feedback showed to be limited if enough feedback was correct. The robustness against incorrect feedback is obtained by making 2 changes to normal SVM updates. The first one is the use of the restriction $||w_t - w_0||$ during the updates. This avoids the adaptive classifier to become worse than the original PI classifier by assuming that in the PS approach still most of the information of the PI approach applies (i.e. a longer HRI leads

to a higher seizure probability). That way, the classifier still gives a mediocre performance even if all feedback is incorrect.

The second change was the addition of the rule to remove data points x_i with $y_i = -1$ from updating b_t if the output $w_t x_i + b_{t-1}$ was higher than 20% of the non-seizure PS data. This not only limits the impact of feedback errors, but also limits the impact of short seizures, again leading to a better performance. During real-time monitoring, it might be that only 1 seizure is available for adaptation, meaning that this seizure data is crucial for further performance of the adaptive algorithm. If this seizure is a noisy or too short seizure, it strongly affects the entire algorithm, potentially leading to a poor performance for a long time in patients with a low seizure frequency. The impact of annotating a seizure HRI as non-seizure is limited due to the large amount of non-seizure data, leading to a limited change in the classifier update. Figure 5.8(b) shows these changes did not only make the adaptive classifier more robust against incorrect feedback, it also gave better results than the normal SVM updating procedure in case of 100% feedback accuracy.

By incorporating a heuristic updating approach, the complexity of the updating procedure is much lower than for the standard SVM updating procedure described in [148]. Therefore this approach is more usable for real-time updating with limited hardware. Although it was only evaluated here on linear SVM, it can also be applied to other kernels if an estimation of the kernel function is used [205]. That way it could also be applied on the PI classifier used in Chapter 3.

5.5.4 General discussion and possible improvements

Section 5.4.3 shows that although giving correct feedback is important for optimal performance, the adaptive approach was able to cope well with large percentages of incorrect feedback. In cases of doubt, patients giving incorrect feedback will not have a drastic effect on the performance, which would be the case with normal updating procedures. A procedure is added here to automatically check if trying to add this seizure data to the SVM model would result in too many false alarms. This way, the model takes care of removing data caused by too short seizures, as they are very hard to detect using only the heart rate and are less important to detect due to their lower risk for injuries. These procedures make the algorithm more useful in practice in case no video-EEG is available to give true data annotations, which is ideal for usage in a home environment using wearable devices.

The adaptive approach showed to adapt to the PS heart rate characteristics. Each patient has its own heart rate variability, both ictal and interictal, and

incorporating this PS information leads to an improved performance compared to the PI performance. However, the results can still be further improved. One of the drawbacks of the system is that we try to get a PS optimal performance while using features that showed to be good for PI use. Better results might be obtained by also including features that work better in a PS setting, such as for example the achieved ictal peak heart rate or the modified cardiac sympathetic index [43, 100]. It might lead to worse initial performance, but might lead to increased performance once sufficient patient-specific data is required.

Despite the increased performance and usability of the heart rate-based seizure detection approach, the unimodal performance is still too low in order to be used successfully in practice. Typical reasons for false alarms like motion artifacts in the ECG, still cause inaccurate heart rate signals and thus, sometimes inaccurate seizure detection, despite the effort of personalization. In order to accomplish better results, it is required to add additional signal modalities. Depending on the targeted seizure types, either limited-channel EEG [78, 210] or ACM/EMG [31, 126, 150] should be added to the heart rate modality. Implementation of a similar adaptive approach for these modalities could also increase their unimodal performance. A combination of several adaptive approaches on different modalities will lead to a much stronger detection algorithm which is advised to be used in practice.

5.6 Conclusion

An adaptive seizure detection algorithm using only heart rate information was proposed. The initial patient-independent classifier was adapted to the patient-specific characteristics by using new patient-specific data points which are annotated with user feedback. The adaptive approach showed to rapidly increase in performance, while also being robust against potential incorrect feedback. This approach not only led to a better performance compared to patient-independent algorithms, but also led to increased usability of such a system in practice. This algorithm is especially of added value compared to other personalization procedures during the first days/week of monitoring at home when only a limited amount of patient-specific data is available.

Chapter 6

Unsupervised nocturnal seizure detection through real-time personalization

Two supervised approaches for personalizing heart rate-based seizure detection were discussed in the previous two chapters. These approaches require that the patient-specific data points are annotated, which leads to either potential inaccurate annotations from patients or a limited amount of patient-specific hospital data. Ideally, no seizure annotations are required to personalize the seizure detection to facilitate usage in practice. In this chapter, a new real-time adaptive heart rate-based seizure detection algorithm is proposed which does not require annotations from the specific test patient. A heuristic classifier quickly adapts itself to the heart rate characteristics of the patient, typically leading to an initial personalization after already some hours of recording. It is evaluated on nocturnal recordings of pediatric patients with different types of epileptic seizures.

A short introduction to the goal of this chapter is described in Section 6.1. The datasets used for this study are presented in Section 6.2. The proposed real-time adaptive heart rate-based seizure detection algorithm is described in Section 6.3, and the results are shown and discussed in Section 6.4 and 6.5. This study was also published in [47].

6.1 Introduction

In Chapter 4 and 5, two algorithms for supervised personalization were proposed. Personalization already showed to be able to lead to performance increase, but the previous two approaches required the availability of annotated patient-specific data. In Chapter 4, accurate annotations were assumed to be made by an expert clinician during the video-EEG monitoring. In this case, such accurate annotations are only possible for a limited amount of time (2-5 days), and it might be that the patient actually did not have a seizure during these recordings, which makes it impossible to adapt the algorithm to this patient. In Chapter 5, the patient or the patient's relatives were assumed to give these seizure annotations. Although extra precaution was built into the algorithm to make it robust against potential incorrect annotations, they still have a negative effect on the performance. In some cases, it is also not clear how to annotate the data. It might be that the algorithm is designed for a certain type of seizures (e.g. strong motor seizures). Then the question is: Does the patient then need to annotate non-motor seizures as seizures? Or what about seizures of only a couple of seconds duration, which are not likely to be accurately detected by state-of-the-art automated seizure detection algorithms? Despite all these problems, all the patient-specific algorithms from the literature require the availability of annotated patient-specific data [99, 157].

To avoid these issues, an unsupervised approach could help. In unsupervised approaches, no annotations are given to data points. It is assumed that the classifier itself decides which data points are different from others, typically based on the data distribution. A well-known example is clustering, in which data points are grouped together based on their data distribution [95]. Other possibilities are the so-called novelty or outlier detectors. In these problems, a large majority of data points originate from one class, and only a small portion correspond to the other class. These outliers are then automatically defined by the classifier based again on the data distribution and often a minimization problem [118, 119, 161].

This chapter proposes an heuristic method that can be interpreted as an outlier detector. In the case of automated seizure detection, typically a lot of non-seizure data points are collected, with only a rare epileptic event. In novelty detection, the method then characterizes what can be seen as normal behavior for the patient (assuming the majority of data is non-epileptic), and everything that is seen as abnormal is then classified as epileptic behavior. The outlier detector is constructed on-the-run, starting with a patient-independent classifier as defined in Chapter 3, and continues to adapt itself to the patient's heart rate characteristics while being used by the patient. In order for the procedure to work, the epileptic heart rate activities should be able to be seen as outliers

compared to non-seizure heart rate activity. This is however only the case during the night. During the day, epileptic heart rate activity is harder to differentiate from non-epileptic activity [48]. Therefore only nocturnal recordings are evaluated in this study. A precursor of this work, discussing a nocturnal adaptive algorithm using seizure annotations, is described in [43].

6.2 Data acquisition

The data used to evaluate the proposed algorithm were recorded in two clinical centers. A first part of the dataset contains nocturnal data from the Pulderbos Revalidation Center for Children and Youth. 14 pediatric patients with 69 seizures were monitored from bed time until the morning (±7-8 a.m.). In the second dataset, data from another 14 pediatric patients with 38 seizures were obtained from the University Hospital of Leuven. Only the night time parts of these recordings (22 p.m. - 8 a.m.) were used here. Both datasets contain electrocardiogram (ECG) signals with 250 Hz sampling frequency. In total 694.6 hours of data were recorded, and both convulsive and subtle seizures were analyzed, both with partial (temporal and frontal lobe) and generalized onsets. Seizures were annotated by experts using video-EEG as gold standard. Only seizures with a duration of at least 20 s were evaluated here as detection of shorter seizures is very difficult with heart rate-based seizure detection [44, 83]. 68 additional seizures shorter than 20 s from both databases were not taken into account in this study, of which 40 seizures were shorter than 10 s. The obtained data was analyzed in a retrospective study using Matlab®, in which a real-time setting was mimicked. An overview of the used datasets is shown in Table 6.1.

6.3 Methodology

6.3.1 Preprocessing

The proposed adaptive seizure detection algorithm uses as input the real-time tachogram. The preprocessing procedure *HRI-EXTRACT* including R peak detection, extraction of the strong heart rate increases and feature extraction, is similar as described in Section 3.3.2. There are however two changes made in this study. First, because the subjects are now pediatric patients rather than adults as in Chapter 3, 4 and 5, the minimal HR_{peak} to detect a heart rate increase is changed from 60 to 90 in (3.3e). Second, other features are extracted here from the heart rate increase. In [43], it was shown that the maximal peak

Table 6.1: Overview of the datasets used in this chapter.

Dataset	Patient	Age	Type of epilepsy†	Type of seizures‡	Duration recordings*	# seizures
Pulderbos dataset	1	7	FLE	HK	70:42:28	13
	2	6	FLE	HK	79:51:23	6
	3	8	GE	TC	22:23:50	5
	4	4	GE	TC	50:44:51	1
	5	18	GE	TC, T, C	10:00:13	13
	6	6	FLE	HK	82:25:10	3
	7	11	FLE	HK	11:36:09	2
	8	14	FLE	HK	22:37:00	3
	9	5	FLE	HK	51:35:57	1
	10	6	GE	TC, T, C	92:11:08	6
	11	8	FLE	HK	12:50:30	9
	12	12	GE	TC	26:30:13	2
	13	3	GE	TC	13:06:40	3
	14	10	FLE	HK	11:57:37	2
UZ Leuven dataset	15	11	TLE	T	10:00:00	1
	16	4	FLE	HK	10:00:00	1
	17	7	TLE	SU	10:00:00	1
	18	11	FLE	T	10:00:00	1
	19	13	TLE	SU	10:00:00	1
	20	16	TLE	SU	10:00:00	4
	21	7	TLE	SU	10:00:00	8
	22	9	TLE	SU	06:00:00	10
	23	11	TLE	SU	10:00:00	1
	24	6	FLE	T	10:00:00	3
	25	11	TLE	T	10:00:00	1
	26	10	FLE	T, SU	10:00:00	3
	27	8	TLE	SU	10:00:00	1
	28	6	FLE	SU	10:00:00	2
Total	28	[3-18]	-	-	694:32:55	107

† GS: generalized seizures, TLE: partial temporal lobe seizures (partial seizures are currently known as focal seizures in the new ILAE classification system), FLE: partial frontal lobe seizures.

‡ C: clonic, HK: hyperkinetic seizures, SU: subtle, T: tonic, TC: tonic-clonic.

* Duration expressed in hh:mm:ss.

heart rate HR_{peak} and the maximal heart rate gradient ∇HR_{max} already result in a good patient-specific performance for nocturnal heart rate-based seizure detection. In order to keep the complexity of the algorithm sufficiently low for usage with wearable devices, only these two features are used in this chapter.

6.3.2 Adaptive classification

Each HRI should then be classified as either a seizure or non-seizure HRI based on these two features. Although it is possible to update machine learning classifiers in real-time [148], it is computationally too expensive to do it with limited hardware specifications. It also requires the availability of seizure

annotations, which are typically not available or possibly inaccurate in a home environment [18].

Therefore, a heuristic adaptive classifier is proposed in this chapter. Normally, classifiers are characterized by a boundary line, splitting up the data points from the different classes. In our case, this boundary is heuristically constructed by using a very limited set of data points. Normal HRI behavior is characterized by fitting a twodimensional ellipse around the majority of patient-specific data.

Whenever a patient-specific data point (coming from a HRI) is detected, it is stored in a pool of noisefree patient data points PD_{pool}. A HRI is assumed to be noisefree if less than 25% of the absolute differences between consecutive heart rate values during this HRI is higher than 10%. When 5 such HRIs are detected, the adaptive classifier can be initialized. HRIs assumed to be caused by noise do not lead to an update of the classifier.

By assuming that the majority of data is caused by non-epileptic behavior, normal behavior is characterized within an ellipse. Data points inside the ellipse can then be seen as normal heart rate behavior, and data points outside the ellipse can be seen as potential seizure activity.

This ellipse is defined by 3 variables (see figure 6.1):

- The center of the ellipse $c(c_x, c_y)$: Defined as the mean value of the data points collected in PD_{pool}.
- Main directions of the ellipse (u, v): The main directions of the ellipse are found by the principal components from PD_{pool} (with the center $c(c_x, c_y)$ subtracted) by means of principal component analysis [43].
- The widths of the ellipse w_u and w_v along both main axes u and v with origin $c(c_x, c_y)$: These are defined as

$$w_u = std_u * sf \qquad \text{and} \qquad w_v = std_v * sf, \tag{6.1}$$

  with std_u the standard deviation of PD_{pool} along the u axes and sf a fixed scale factor (set to 2.5 according to [43]) and similar for std_v. In order to limit the impact of potential seizure data on the computation of $std_{u,v}$, the widths w_u and w_v are limited so they can not exceed a heuristically defined value of 15. This allows most non-seizure data to be located inside the ellipse, whereas most seizure data will be located outside the ellipse without requiring the corresponding seizure annotations.

The equation of the ellipse classifier along axes (u,v) then becomes

$$f(u, v) : \left(\frac{u}{w_u}\right)^2 + \left(\frac{v}{w_v}\right)^2 = 1 \quad . \tag{6.2}$$

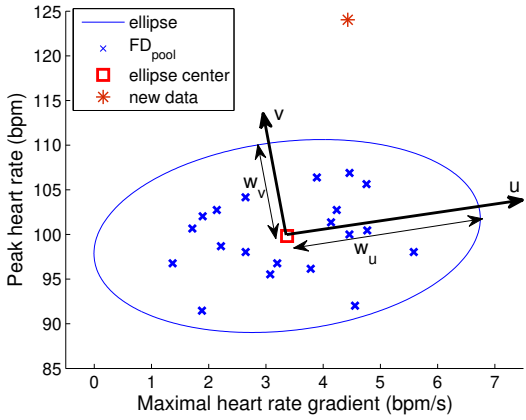


Figure 6.1: Illustration of the proposed adaptive seizure detection algorithm. The decision boundary is formed by the ellipse: data points outside the ellipse cause a warning if they also have a peak heart rate higher than the peak heart rate of the center of the ellipse. In this example, a new seizure data point needs to be classified (marked as a star).

Once the classifier is initialized, the classification becomes straight forward. For each new data point $d(d_x, d_y)$, first subtract the center $c(c_x, c_y)$ of this data point and adjust $d(d_x, d_y)$ to the appropriate axes u and v (called $d(d_u, d_v)$). Next, evaluate $d(d_u, d_v)$ with the equation of the constructed ellipse defined as

$$y(d) = \left(\frac{d_u}{w_u}\right)^2 + \left(\frac{d_v}{w_v}\right)^2 \tag{6.3}$$

if data point d is located inside the ellipse or not. If it is located inside the ellipse ($y(d) <= 1$), it can be seen as normal heart rate behavior. If it is located outside the ellipse ($y(d) > 1$) and the peak heart rate of x_t is higher than the average peak heart rate in PD_{pool} (c_y), it is classified as a seizure HRI. This extra rule is added in order to avoid HRIs with a peak heart rate below the ellipse to cause an alarm, as ictal peak heart rates are assumed to be on average higher than non-ictal peak heart rates [44].

The ellipse is readjusted every time a new noise-free data point is detected in real-time. In order to keep the complexity of the algorithm sufficiently low, only the last 20 data points are used to construct the ellipse.

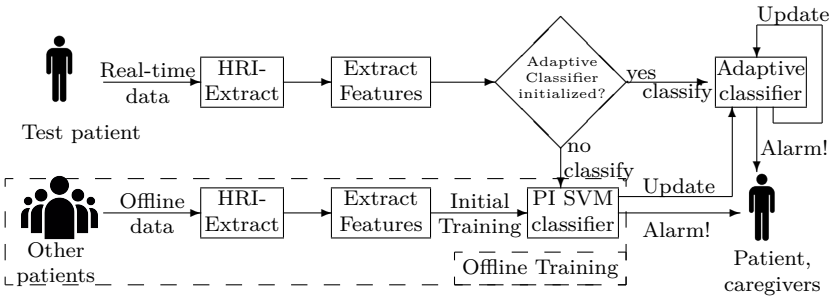


Figure 6.2: Overview of the proposed real-time adaptive seizure detection algorithm. The part in the dashed box shows the initial offline procedure that only needs to be done once. HRI: heart rate increase, PI: patient-independent.

6.3.3 Initial patient-independent classification

The adaptive classifier mentioned above can only be used if 5 patient-specific noise-free data points are collected. Before these 5 data points are collected, the algorithm should also result in a decent performance to lower the impact of this initiation period as much as possible. Therefore, the data points are classified with a patient-independent classifier during this initialization phase . In this case a support vector machine (SVM) classifier is used with the same two features mentioned in Section 6.3.1. The classifier is trained using a leave-one-patient-out approach: the classifier is trained on data from all patients except the one used for testing, similarly as in Section 3.3.5. An overview of the entire procedure for adaptive seizure detection is illustrated in Figure 6.2.

6.3.4 Performance evaluation

The performance of the proposed seizure detection algorithm is evaluated on the data discussed in Section 6.2. The same metrics are used for evaluation as the ones mentioned in Section 3.3.6: sensitivity, FAR (expressed as the number of false positives per night, FP/night, with night defined as 8 hours of recording [180]), detection delay and PPV. The detection of seizures shorter than 20s are counted nor as a true positive, nor as a false positive. Overall seizure-based performance metrics are used in this study, similarly as in [44]. Very similar results are found if a patient-averaged overall performance is used, so only the seizure-based overall performance is mentioned in this study.

Mann-Whitney U tests were used to evaluate whether the results differed significantly, calling the results significantly different when $p < 0.05$ with

Bonferroni correction. The mentioned 95% confidence intervals (CI) for the median estimates are calculated using the rank orders. Values are selected from the ranks closest related to a significance level of 5%. It is also investigate whether the seizure and epilepsy type has an effect on the results found by the proposed adaptive seizure detection algorithm. A first distinction here is made between the seizure onset, either generalized or partial (both temporal and frontal lobe are investigated here). A second distinction is made on the clinical manifestation of the seizure (independent of the seizure onset), making a distinction between tonic, clonic, tonic-clonic, hyperkinetic and non-motor partial (called 'subtle' from now on) seizures.

6.4 Results

An overall sensitivity of 77.6% was found on average on all patients, with on average 2.56 FP/night and 30.7% PPV over the entire recordings including the initialization phase of the adaptive algorithm. The average detection delay is 19.1 s. Variability is however found for these values depending on the clinical nature of the seizures and the seizure onset. Table 6.2 shows the impact of the seizure onset on the performance. Primary generalized and temporal lobe (TL) seizures are detected with around 83% sensitivity, whereas this is only 71.4% for the frontal lobe (FL) seizures. Overall, 17/28 patients (60.7%) show a 100% sensitivity, of which 3/6 patients with generalized seizures, 7/9 with TL onset and 7/13 with FL onset. Patients with generalized (1.92 FP/night) and FL seizures (2.48 FP/night) also have lower false alarm rates (FAR) compared to patients with TL seizures (5.36 FP/night). Despite the overall found differences, none of the sensitivity and FAR results led to significant different results between different seizure onsets for the proposed adaptive algorithm. TL seizures (27 s) are however detected significantly later than FL (17s, $p < 10^{-2}$) and generalized seizures (15 s, $p < 10^{-3}$).

If only the patient-independent algorithm is used on all data (including the initialization phases of the adaptive approach), an overall sensitivity of 81.3% and a FAR of 6 FP/night is found. The adaptive algorithm thus lowers the number of false alarms with 57% over the entire recordings compared to the patient-independent algorithm with a similar sensitivity. The Mann-Whitney U tests show this adaptive algorithm is indeed significantly better compared to the patient-independent algorithm in FAR ($p < 10^{-3}$, median difference: 4.78 FP/night, CI: [2.77,7.20]), but with no significant difference in sensitivity ($p > 0.05$, median difference:0%, CI:[0,0]). The average detection delay of the patient-independent approach is 15.6 s, which is around 3.5 s faster than the adaptive approach. The adaptive algorithm however reduces the FAR drastically

Table 6.2: Overview of the results of the proposed adaptive approach for different presumed seizure onsets.

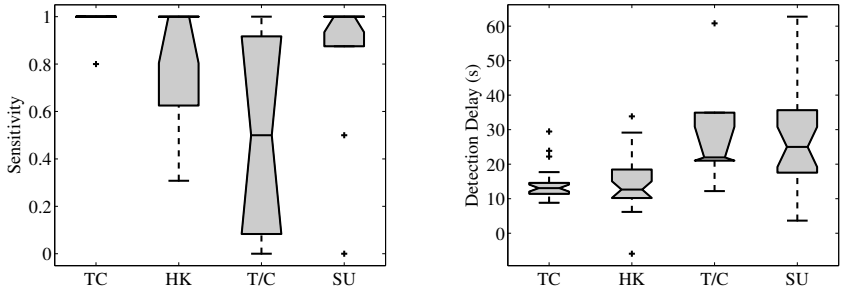
Seizure onset	Sensitivity (%) (detected/total)	Mean delay(s)	False alarm rate (FP/night)
Generalized	83.3 (25/30)	15.0	1.92
Frontal lobe	71.4 (35/49)	16.9	2.48
Temporal lobe	82.1 (23/28)	26.9	5.36
Total	77.6 (83/107)	19.1	2.56

by waiting longer to be more sure on the distinction between epileptic and non-epileptic HRIs. Similar findings on the effect of the onset can be found in the results of the patient-independent algorithm due to the large overlap in sensitivity results.

Figure 6.3(a) shows the boxplots for the sensitivity based on ictal clinical manifestations. 96% (24/25) of tonic-clonic (TC) seizures and 72.5% (29/40) of hyperkinetic (HK) seizures are detected successfully with the proposed seizure detector. Only 46.2% (6/13) of the tonic (T) and clonic (C) seizures are detected, whereas 82.8% (24/29) of the subtle seizures are detected. 5/6 patients had 100% sensitivity for the tonic-clonic seizures, whereas this holds for 6/9 patients with HK seizures and 7/9 patients with subtle seizures. The inter-patient variability for subtle seizures is however lower compared to the HK and T/C seizures. However, again no statistical significant differences between the sensitivity of different seizure types were found for the proposed adaptive algorithm.

Also a difference is found in the detection delays of the clinical more severe seizures (TC and HK) compared to the clinical more subtle seizure (see figure 6.3(b)). The TC and HK seizures ($\pm$14 s) are detected faster than the more subtle seizures and the tonic and clonic seizures ($\pm$28 s). The inter-patient variability for the detection delay is also higher for the subtle seizures compared to the other seizure types. Significant differences were found in detection delays between TC and subtle seizures ($p < 10^{-3}$) and between HK and subtle seizures ($p < 10^{-3}$).

The median duration of the initialization phase is 2.7 hours (CI: [2.27,4.67]). During the initialization phase, the average FAR is 3.52 FP/night for both algorithms. After initialization, the false alarm for the adaptive algorithm is 2.32 FP/night compared to on average 6.40 FP/night for the patient-independent algorithm (63% reduction in false alarms). Figure 6.4 shows the boxplots for the FARs for both algorithms for the groups with different seizure onsets. For all groups both the median and variance in FAR drop strongly with the adaptive algorithm compared to the patient-independent algorithm. The differences



(a) Sensitivity per seizure type.

(b) Detection delay per seizure type.

Figure 6.3: Boxplots for sensitivity (a) and detection delay (b) for the different clinical seizure types. TC: Tonic-clonic, HK: hyperkinetic, T: tonic, C: clonic, SU: subtle seizures.

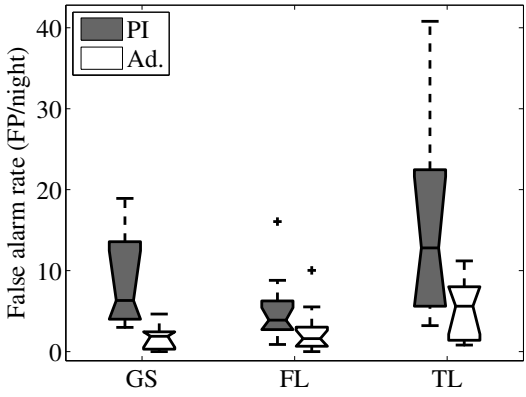


Figure 6.4: False alarm rate per type for patients for the original patient-independent (PI) algorithm and the proposed adaptive (Ad.) algorithm. GS: generalized seizures, FL: frontal lobe, TL: temporal lobe.

are the largest for the TL patients (on average 4.24 FP/night for the adaptive algorithm) and are the smallest for the FL patients (2.32 FP/night). On average 1.60 FP/night are found in patients with mainly generalized seizures.

The majority of false alarms from the adaptive algorithm are found in the early morning after 6 a.m., and around 25% of false alarms are caused when the

patient was already fully awake. Another 25% of false alarms are caused by strong motion artifacts in the ECG, leading to errors in the R peak detection algorithm. Other typical reasons for false alarms were arousals, long periods of nocturnal awake time and non-epileptic spasms.

6.5 Discussion

6.5.1 Sensivity and detection delay

Table 6.2 and figure 6.3 show that the seizure type and onset have a tendency to have an impact on the sensitivity and FAR of the proposed seizure detection algorithm, but none of these results showed to be significantly different. This is assumed to be caused by the limited amount of data points in each group. Almost all generalized seizures are detected, and for these patients the FAR shows to be relatively low compared to the patients with FL or TL seizures. This means this approach works well for these patients, which are the most important patients to monitor at home.

The detection delay results showed significant differences, showing that TC and HK seizures are detected significantly faster than partial subtle seizures. This illustrates that the clinical more important seizures are detected faster than seizures that typically require less attention.

Less than half of the tonic or clonic seizures are detected with the adaptive algorithm. This is due to the fact that most of these seizures occurred in patients with also TC seizures. The TC seizures (and post-ictal activity) result in stronger HRIs, causing the less strong T or C seizures to be seen as more normal behavior. These T and C seizures are indeed detected with the patient-independent algorithm, but not with the adaptive version. Compared to the patient-independent algorithm, only 4 T and C seizures are missed extra (and no seizures of an other type), of which 3 C seizures in patients with TC seizures.

The (most often subtle) TL seizures are detected with a sensitivity of $\pm 80\%$, similar as reported in [44]. The detection of these seizures is however on average slower than for convulsive seizures. This is in line with the findings of [169], which stated that ictal HRIs from patients with TL epilepsy have a longer duration than for example the ictal HRIs from FL patients, thus leading to a slower detection delay. Another reason might be that the late detections of the TL seizures are caused by a limited spread of the temporal lobe compared to extratemporal origins. The detection of the T and C seizures are on average also slower due to the fact that most of them are secondary generalized, leading to a slower detection compared to the primary generalized TC seizures.

6.5.2 False alarm rate

Figure 6.4 shows that the FAR drops strongly after the initialization phase with 63% compared to the patient-independent algorithm. Not only the overall FAR drops, but also the inter-patient FAR variability drops strongly. This can especially be seen for the patients with GS and TL patients. This way, the algorithm is more usable for a wider range of patients. Despite the decreased FAR variability, the FAR variability for the TL patients remains larger compared to that of the GS and FL patients.

The patients were already fully awake during 25% of the false alarms. In practice, a device for nocturnal monitoring can be turned off then to avoid these false alarms. Another quarter of false alarms is caused by strong motion artifacts, which lead to heart beat detection errors. Better noise removal techniques and the usage of more robust sensors could reduce the impact of noise removal on seizure detection performance. Also the addition of other modalities such as ACM and EMG can lead to increased performance for certain types of seizures [126].

6.5.3 General discussion and possible improvements

The proposed adaptive seizure detection algorithm is designed to be low-complex. The adaptive classifier is computationally cheap compared to *HRI-EXTRACT* and the feature extraction procedure, which was already discussed to be sufficiently low-complex in [44]. This way, the algorithm can be implemented directly on a wearable device rather than be connected to a smartphone/server for computation purposes. Photoplethysmography could also be used instead of ECG for extraction of the heart rate in order to increase the wearability of the system [180], but might have a negative impact on the accuracy of the heart rate data [191].

The adaptive algorithm was already initialized after, on average, 2.7 hours. At this point, the adaptive algorithm already did most of the adaptation to the patient characteristics. Manual setting of patient parameters is typically done offline after multiple days or even weeks depending on the amount of recorded seizures. Therefore the proposed algorithm leads much faster to personalization of the detection system compared to the manual alternative. A person's heart rate characteristics vary over time, also when children are growing up [27, 165]. As the algorithm continuously changes over time, it will also adjust to the patient's heart rate characteristics while they grow up, without requiring manual adaptations of the algorithm.

The proposed methodology does not require the availability of annotated patient data. That way, no clinicians need to annotate previously recorded data for each patient. An alternative would be to incorporate user feedback [41], but patients and relatives might not always be aware whether an alarm was correct or not, and missed seizures will remain missed in most cases. Therefore this approach is advised, certainly for nocturnal monitoring of pediatric patients. It works with the assumption that most HRIs are caused by non-epileptic behavior (e.g. arousals). In most cases, this assumption holds well. Only in long series of seizures, the majority of data in PD_{pool} might be epileptic, in which case some seizures might be missed after the correct detection of 5-10 seizures. However, due to the fact that this collection of data always changes in time, this effect goes away again quickly after the series of seizures is stopped.

The proposed algorithm results in a strong decrease in false alarms compared to the patient-independent algorithm. In practice, no extra effort from the patient, clinician or system owner is required with this adaptive algorithm compared to when using a patient-independent algorithm, and should therefore be preferred in real-life usage. However, still too many false alarms are generated in order to be used in practice in the proposed unimodal setting. The proposed method should be combined with other modalities like EMG, accelerometry or wearable EEG using either patient-independent or patient-specific algorithms [7, 31, 78]. Only then will the algorithm lead to an acceptable performance for real usage in practice.

The unimodal heart rate-based seizure detection can be further improved by going into more complex offline or online adaptive seizure detection algorithms. One option would be to also add sleep stage information to the algorithm in order to further finetune it [86]. Another possibility is to apply unsupervised transfer learning in order to allow offline personalization. This will be the topic of Chapter 7.

6.6 Conclusion

The proposed seizure detection algorithm allows to quickly adapt to patient-specific heart rate characteristics, leading to 57% less false alarms compared to a patient-independent algorithm. The adaptation not only leads to an overall decreased false alarm rate, but also to less inter-patient false alarm rate variability. This illustrates that unsupervised personalization could lead to strongly increased performance for automated nocturnal heart rate-based seizure detection.

Chapter 7

Personalizing nocturnal seizure detection using unsupervised transfer learning

An unsupervised real-time personalization approach for heart rate-based seizure detection was proposed in Chapter 6. Although this heuristic approach led to an increase in performance, it can be assumed that a more complex and offline personalization procedure could lead to a further increased performance. In this chapter, a new unsupervised transfer learning approach for support vector machines is proposed, which is able to personalize a reference patient-independent classifier by using unannotated patient-specific data. The proposed method showed to already lead to an increased performance after two nights, leading to performance convergence after around two weeks of patient data.

The goal of the proposed method is shortly summarized in Section 7.1. The datasets used in this study are presented in Section 7.2. The proposed new unsupervised transfer learning approach for personalizing heart rate-based seizure detection is described in Section 7.3. The results of this approach and the reference approaches are described and discussed in Section 7.4 and 7.5. A pilot study of this chapter is described in [45].

7.1 Introduction

In the previous chapter, it was illustrated that the proposed heuristic unsupervised approach leads to an efficient and fast personalization of nocturnal heart rate-based seizure detection. The heuristic approach is of low algorithmic complexity, allowing a real-time personalization to the patient's heart rate characteristics. However, offline personalization would allow more complex methodologies to be used, for example during the day when the device is not used by the patient. In machine learning, unsupervised classification can be performed by for example one-class support vector machine (SVM) [160]. In one-class SVM, the majority of data is characterized automatically based on the data distribution and an a priori defined percentage of accepted outliers ν. The problem with this approach is that the classifier itself decides which data points are the outliers from the dataset, not allowing any addition of prior knowledge in the optimization of the classifier. That way, the selection of outliers could be finetuned in order to improve performance.

In this chapter, an adaptation is proposed to the transfer learning approach discussed in Chapter 4. This approach applies transfer learning to a one-class SVM. Similarly as in Chapter 4, the supervised patient-independent classifier discussed in Chapter 3 is used as reference classifier. The algorithm is then personalized by using a limited amount of unannotated patient-specific data. This approach allows to incorporate prior knowledge from the supervised patient-independent classifier into the personalized unsupervised classifier, which could lead to a more robust unsupervised personalization procedure.

7.2 Data acquisition

The data used in this chapter are obtained from 'Nightwatch' bracelets (provided by Livassured BV) during nocturnal recordings. These bracelets acquire photoplethysmography (PPG) and 3-D accelerometry signals, which are obtained from the upper arm. In this study, only the PPG signal is used for heart rate extraction.

The trial in which the data is acquired, was performed in two phases (Dutch Trial Registry 4115). In the first phase of the trial, 20 patients were monitored for around 3 months to test the initial equipment. This dataset is used for training the reference patient-independent algorithm discussed in Section 7.3.2 and parameter optimization of the personalized approach discussed in Section 7.3.3. During the second phase of the trial, 22 patients were monitored. Some patients used in phase 1 were also present in phase 2, but also new patients

Table 7.1: Detailed overview of the dataset used for testing in this chapter. TC: tonic-clonic seizure, HK: focal hyperkinetic seizure, T>30: generalized tonic seizure with a duration of more than 30 s.

Patient Number	Patient code	Age	Sex	# Nights	# TC	# HK	# T>30	# Major	# All Major
1	102	19	m	40	10	0	0	3	13
2	103	23	m	53	1	14	1	9	25
3	104	17	m	39	10	0	1	1	12
4	153	15	m	85	13	0	11	10	34
5	154	24	m	94	0	68	0	5	73
6	155	21	m	41	2	0	18	6	26
7	301	22	m	45	8	0	0	0	8
8	302	36	f	48	69	10	7	4	90
9	303	23	m	64	32	0	0	0	32
10	304	20	m	96	79	0	56	1	136
11	305	34	f	119	0	72	0	0	72
12	306	29	f	94	15	0	0	0	15
13	309	23	m	111	11	0	0	0	11
14	310	19	f	98	2	29	0	8	39
15	311	23	m	52	2	0	0	0	2
16	312	40	f	99	14	0	0	0	14
17	313	38	m	111	0	0	2	1	3
18	314	22	m	64	1	0	68	0	69
19	315	36	m	48	2	0	6	3	11
20	316	18	f	52	13	0	2	0	15
21	317	38	m	55	0	0	2	0	2
22	318	54	m	43	5	0	6	2	13
Total		[15-54]	16m/6f	1551	289	196	180	53	715

were added to the study. The data obtained from phase 2 of the trial are used for training and testing the proposed transfer learning algorithm of this chapter. A detailed overview of this test dataset is shown in Table 7.1. All patients had an intellectual disability with a history of more than one major nocturnal seizure per month. Patients with a pacemaker, movement disorder or skin pigmentation were excluded from the study. All participants or their legal guardian provided written informed consent to the study, which is approved by the Medical Research Ethics Committee of the University Medical Center in the Netherlands. A detailed description of this dataset can also be found in [5].

Reported events by the Nightwatch bracelet were annotated by trial nurses. Additionally, in a random sample of 10% of all nights, complete video recordings were screened for more detailed annotations. Clinically urgent seizures were annotated as 'major' seizures, including: tonic-clonic (TC) seizures, focal hyperkinetic (HK) seizures, long generalized tonic seizures (of more than 30 s duration) and other major seizures (e.g. TC-like seizures with atypical semiology and clusters of minor seizures). Video was used as the gold standard for seizure

type annotation. All other found seizures not annotated as major seizures are annotated as minor seizures, but are not the main focus of this study.

7.3 Methodology

7.3.1 Preprocessing and feature extraction

PPG is used in this chapter to extract the heart rate, which is done by using the built-in procedure from the Nightwatch bracelet. The constructed tachogram is then processed by the *HRI-EXTRACT* procedure mentioned in Section 3.3.2 to extract the strong sympathetic heart rate increases from the tachogram.

Similarly as in Chapter 6, only the peak heart rate and maximal heart rate gradient are extracted from these heart rate increases. These two features already showed to lead to a maximal performance with a limited amount of features [45, 47]. These features are then used for classification to indicate whether this heart rate increase is caused by an epileptic seizure or not.

7.3.2 Patient-independent reference classifier

The patient-independent classifier is used as reference for the personalized transfer learning approach in Section 7.3.3. Similarly as in Chapter 3, the corresponding supervised support vector machine (SVM) classifier is found by solving the following optimization problem:

$$\begin{aligned} &\min_{w,b,\xi} \frac{1}{2}||w||^2 + C\sum_{i=1}^{N} c_i\xi_i \\ &s.t. \quad \begin{cases} y_i(w^T\varphi(x_i)+b) \geq 1-\xi_i \\ \xi_i \geq 0 \end{cases} \quad ,\forall i \in [1,N] \end{aligned} \tag{7.1}$$

with x_i the data points originating from other patients than the one used for testing, y_i the corresponding ground truth (class -1: seizure, class +1: non-seizure) and ξ_i the error for x_i in this model. An RBF kernel $K(x_i, x_j) = \varphi(x,i)\varphi(x_j)$ with kernel parameter σ is used here similarly as in Chapter 3. Parameters c_i are defined here as well as

$$c_i = \begin{cases} \gamma\frac{(N^+ + N^-)}{2N^-} & : y_i = -1 \\ \frac{N^+ + N^-}{2N^+} & : y_i = +1 \end{cases} \tag{7.2}$$

with N^+ and N^- indicating the number of datapoints belonging to classes +1 and -1 respectively.

In this study, the patient-independent classifier is trained on the data obtained during phase 1 of the trial. If the test patient from phase 2 did not participate in phase 1, the used patient-independent classifier was trained on data from all patients from phase 1. If the patient however also participated in phase 1 of the trial, data from all patients except the test patient from this phase are used for training the patient-independent classifier. Hyperparameter optimization of σ, γ and C for the patient-independent classifier was performed on the data from phase 1 using leave-one-patient-out crossvalidation.

7.3.3 Unsupervised personalization

One-class SVM

In unsupervised classification, no annotations y are available during training. Therefore, outlier detection is done by characterizing the largest portion of data points. In SVM, one option to do this is by solving following minimization problem for patient-specific data points $\widetilde{x}_k$ [160]

$$\begin{aligned} &\min_{\widetilde{w},\rho,\widetilde{\xi}} \frac{1}{2}||\widetilde{w}||^2 - \rho + \frac{1}{\nu M}\sum_{k=1}^{M}\widetilde{\xi}_k \\ &s.t. \quad \begin{cases} \widetilde{w}^T\varphi(\widetilde{x}_k) \geq \rho - \widetilde{\xi}_k \\ \widetilde{\xi}_k \geq 0, \qquad\qquad\qquad\qquad ,\forall k \in [1, M] \end{cases} \end{aligned} \tag{7.3}$$

with $\widetilde{w}$ and ρ the unknown variables which characterize the one-class SVM classifier. Parameter ν represents the fraction of outliers assumed to be present in the dataset, which is assumed to be known a priori. Similarly as in supervised SVM, also the dual problem can be solved

$$\begin{aligned} &\min_{\widetilde{\alpha}} \frac{1}{2}\sum_{k=1}^{M}\sum_{l=1}^{M}\widetilde{\alpha}_k\widetilde{\alpha}_l K(\widetilde{x}_k, \widetilde{x}_l) \\ &s.t. \quad \begin{cases} \sum_{k=1}^{M}\widetilde{\alpha}_k \\ 0 \leq \widetilde{\alpha}_k \leq \frac{1}{\nu M}, \quad ,\forall k \in [1, M] \end{cases} \end{aligned} \tag{7.4}$$

with $K(\widetilde{x}_k, \widetilde{x}_l) = \varphi(\widetilde{x}_k)\varphi(\widetilde{x}_l)$ the applied RBF kernel. From (7.4), it can be seen that ν represents an upper bound on the fraction of outliers present in the dataset. Parameter ν can also be interpreted as a lower bound on the fraction of support vectors. Note that w in one-class SVM is defined as

$$w = \sum_{k=1}^{M} \widetilde{\alpha}_k \varphi(\widetilde{x}_k). \tag{7.5}$$

Unsupervised transfer learning

A new method is proposed which applies unsupervised transfer learning to a supervised reference SVM problem. This way, annotated data can be used in the reference model, allowing a more robust reference model to be used. The new proposed solution for unsupervised transfer learning is defined by the following minimization problem:

$$\begin{aligned} &\min_{\widetilde{w},\rho,\widetilde{\xi}} \frac{1}{2}||\widetilde{w} - w||^2 - \rho + \frac{1}{\nu M}\sum_{k=1}^{M}\widetilde{\xi}_k \\ &s.t. \quad \begin{cases} \widetilde{w}^T\varphi(\widetilde{x}_k) \geq \rho - \widetilde{\xi}_k \\ \widetilde{\xi}_k \geq 0, \end{cases} \quad ,\forall k \in [1, M] \end{aligned} \tag{7.6}$$

with w obtained from the patient-independent classifier. An optimal solution is found for the saddle point in the Langragian

$$\begin{aligned} &\max_{\tilde{\alpha},\tilde{\beta}} \min_{\tilde{w},\rho,\tilde{\xi}} \mathcal{L}(\widetilde{w}, \rho, \widetilde{\xi}; \widetilde{\alpha}, \widetilde{\beta}) \\ &= \max_{\tilde{\alpha},\tilde{\beta}} \min_{\tilde{w},\rho,\tilde{\xi}} \frac{1}{2}||\widetilde{w} - w||^2 - \rho + \frac{1}{\nu M}\sum_{k=1}^{M}\widetilde{\xi}_k \\ &\qquad - \sum_{k=1}^{M}\widetilde{\alpha}_k(\widetilde{w}^T\varphi(\widetilde{x}_k) - \rho + \widetilde{\xi}_k) - \sum_{k=1}^{M}\widetilde{\beta}_k\widetilde{\xi}_k \end{aligned} \tag{7.7}$$

with $\widetilde{\alpha}_k$, $\widetilde{\beta}_k \geq 0$ the multipliers. This then leads to

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \widetilde{w}} = 0 & \rightarrow \tilde{w} = w + \sum_{k=1}^{M} \widetilde{\alpha}_k \varphi\left(\widetilde{x}_k\right) \\ \frac{\partial \mathcal{L}}{\partial \rho} = 0 & \rightarrow \sum_{k=1}^{M} \widetilde{\alpha}_k = 1 \\ \frac{\partial \mathcal{L}}{\partial \widetilde{\xi}_k} = 0 & \rightarrow \frac{1}{\nu M} = \widetilde{\alpha}_k + \widetilde{\beta}_k, \forall k \in [1, M] \end{cases} \tag{7.8}$$

which results in

$$\begin{aligned} \widetilde{w} &= w + \sum_{k=1}^{M} \widetilde{\alpha}_k \varphi(\widetilde{x}_k) \\ &= \sum_{i=1}^{N} \alpha_i y_i \varphi(x_i) + \sum_{k=1}^{M} \widetilde{\alpha}_k \varphi(\widetilde{x}_k) \end{aligned} \tag{7.9}$$

with x_i, y_i and α_i the data points, corresponding ground truths and support vectors from data from other patients, and $\widetilde{\alpha}_k$ the support vectors from the patient-specific data $\widetilde{x}_k$. This shows that the new patient-specific $\widetilde{w}$ can be seen as a combination of the patient-independent w from the reference binary classifier and the patient-specific data points.

Substituting (7.8) in (7.7) leads to the dual quadratic programming problem

$$\begin{aligned} &\min_{\widetilde{\alpha}} \sum_{k=1}^{M} \sum_{l=1}^{M} \widetilde{\alpha}_k \widetilde{\alpha}_l K(\widetilde{x}_k, \widetilde{x}_l) + \sum_{k=1}^{M} \sum_{i=1}^{N} \widetilde{\alpha}_k \alpha_i y_i K(\widetilde{x}_k, x_i) \\ &s.t. \quad \begin{cases} 0 \leq \widetilde{\alpha}_k \leq \frac{1}{\nu M} \\ \sum_{k=1}^{M} \widetilde{\alpha}_k = 1 \end{cases} \quad , \forall k \in [1, M] \end{aligned} \tag{7.10}$$

This indicates that the meaning of parameter ν is similar in this minimization problem as in the original one-class SVM problem (7.3). Solving the above mentioned equation leads to the required values of $\widetilde{\alpha}_k$, but parameter ρ still needs to be defined. This can be done by again using the Karush-Kuhn-Tucker (KKT) conditions, which state that in an optimal solution, the product of the dual variables and the constraints should be zero. The KKT conditions yield

$$\begin{cases} \dfrac{\partial \mathcal{L}}{\partial \widetilde{w}} = 0 \rightarrow \tilde{w} = w + \displaystyle\sum_{k=1}^{M} \widetilde{\alpha}_k \varphi\left(\widetilde{x}_k\right) & \\ \dfrac{\partial \mathcal{L}}{\partial \rho} = 0 \rightarrow \displaystyle\sum_{k=1}^{M} \widetilde{\alpha}_k = 1 & \\ \dfrac{\partial \mathcal{L}}{\partial \widetilde{\xi}_k} = 0 \rightarrow \dfrac{1}{\nu M} = \widetilde{\alpha}_k + \widetilde{\beta}_k & ,\forall k \in [1, M] \\ \widetilde{\alpha}_k \left(\widetilde{w}^T \varphi(\widetilde{x}_k) - \rho + \widetilde{\xi}_k \right) = 0 & ,\forall k \in [1, M] \\ \widetilde{\beta}_k \widetilde{\xi}_k = 0 & ,\forall k \in [1, M] \\ \widetilde{\alpha}_k \geq 0 & ,\forall k \in [1, M] \\ \widetilde{\beta}_k \geq 0 & ,\forall k \in [1, M] \end{cases} \tag{7.11}$$

Similarly as for binary SVM (see Chapter 3), these conditions yield that for a solution $\widetilde{\alpha}_k^* \in \left(0, \frac{1}{\nu M}\right)$, $\widetilde{\xi}_k^* = 0$ so that

$$\widetilde{w}^T \varphi(\widetilde{x}_k) - \rho = 0 \text{ for } \widetilde{\alpha}_k \in \left(0, \frac{1}{\nu M}\right) \tag{7.12}$$

Parameter ρ can then be estimated by taking the average of the values found for the above equation for all $0 < \widetilde{\alpha}_k < \frac{1}{\nu M}$.

The classification of a new data point $\widetilde{x}_n$ is then defined by

$$\widetilde{y}(\widetilde{x}_n) = \sum_{k=1}^{M} \widetilde{\alpha}_k K(\widetilde{x}_k, \widetilde{x}_n) + \sum_{i=1}^{N} \alpha_i K(x_i, \widetilde{x}_n) - \rho \quad . \tag{7.13}$$

The same kernel K is used here as the one used in the patient-independent classifier. Similarly as in all SVM applications, the summation in (7.13) only needs to be evaluated for the real support vectors with $\widetilde{\alpha}_k > 0$.

The personalized transfer learning approach is tested using 5-fold crosstesting on the patient-specific data from phase 2 of the trial. This means that 4 folds of patient-specific data are used for training, whereas the fifth fold is used for testing. This is repeated so that each fold is used once as test set.

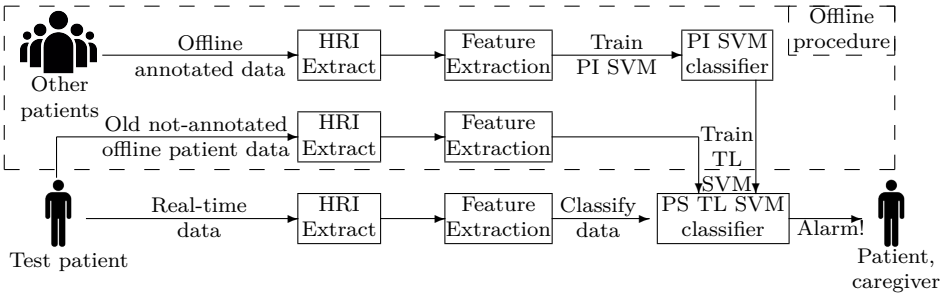


Figure 7.1: Overview of the proposed unsupervised transfer learning approach for personalized heart rate-based seizure detection. PS TL SVM: patient-specific transfer learning support vector machine; PI SVM: patient-independent support vector machine.

ν estimation

Parameter ν is a hyperparameter that needs to be tuned before the minimization problem (7.10) can be solved. It indicates an upper bound on the fraction of outliers in the dataset, but this percentage is unknown in this case. In the application of seizure detection, it can also greatly vary between different subjects, or within different recordings from one patient. Therefore, it is important to have a good automated estimation of it in order to obtain a good result for the proposed unsupervised transfer learning approach.

Standard procedures for the estimation of ν can not take into account the prior knowledge of the reference classifier [71]. These approaches typically search for a percentage of outliers based on the available (patient-specific) data and their distribution. They however often require a large amount of data in order to get a robust estimation of the fraction of outliers. This makes them hard to use for this application, as the transfer learning method should work for a limited amount of data. Therefore, parameter ν is estimated here by analyzing the distribution of the latent variable outputs of the patient-independent classifier on the patient-specific training data. This way, outliers are defined as outliers in function of the patient-independent classifier output, only counting for outliers which would be seen as seizures in the reference classifier. This allows for a more accurate ν estimation in the proposed methodology.

The outputs $y(\widetilde{x}_k)$ from the patient-independent classifier of the patient-specific data points $\widetilde{x}_k$ typically follow a distribution similar to a Gaussian distribution (see Figure 7.2). As can be assumed that the majority of data is non-epileptic, this peak (located at a patient-independent SVM output of $peak_{index}$) in the histogram is caused by non-epileptic behavior. The cut-off point to differentiate

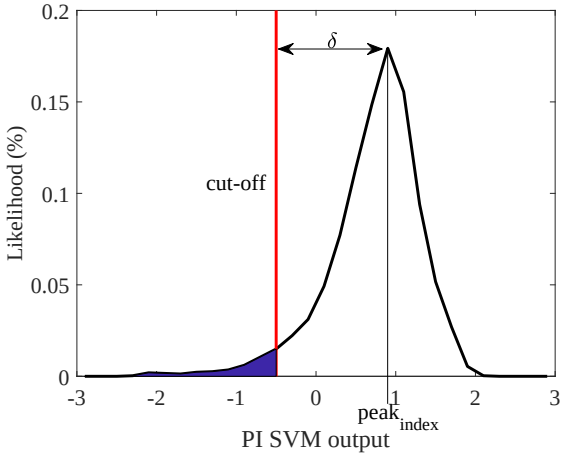


Figure 7.2: Illustration of the used procedure for automated estimation of parameter ν in (7.3). $peak_{index}$ indicates the patient-independent (PI) SVM output for which the highest likelihood is found in the available patient-specific data. The cutoff value (the vertical red line) is found by subtracting $peak_{index}$ with a fixed value δ. The blue area indicates the data points which are seen as outliers for this patient. Parameter ν is then estimated as the blue area over the entire area under the likelihood curve.

outliers from normal heart rate behavior is therefore based on when this Gaussian-like distribution is showing a strong decrease. The optimal value of this cutoff value ($peak_{index} - \delta$) was trained during the validation of this approach on the data from phase 1. Parameter ν is then estimated as the fraction of data points with a classifier output lower than this cutoff value (the area marked in blue in Figure 7.2 over the complete area under the curve). An overview of the entire followed procedure is summarized in Algorithm 7.1.

7.3.4 Simulation for evaluating the impact of training size

One of the strong aspects of transfer learning is the fact that it allows to make a patient-specific classifier with a limited amount of patient-specific data. It is however of course important to have an idea on how much data is required for the proposed unsupervised personalization approach. In order to evaluate this, it is evaluated what the impact on performance is if only 1, 2, 3, 4, 7, 10, 14, 21 and 28 nights of a specific patient are available for training the personalized

Algorithm 7.1: Procedure for unsupervised personalization using transfer learning.

Input: Online HR data HR_{new} and offline HR data HR_{old} of patient P_p from phase 2 of the study, patient-independent classifier SVM_{PI} defined by w and b trained on phase 1 data (except data from P_p) with kernel projection φ and hyperparameters C and γ trained using LOPO CV
Output: Alarm in case of an assumed seizure for patient P_p
Offline personalized construction with offline patient-specific data HR_{old}
Train personalized model by using offline patient-specific data (4/5 folds of patient-specific data: HR_{old}, 1/5 folds: HR_{new})
Extract latent variables $LV_{PI} = w\varphi(\tilde{x}_{old}) + b$ for $\forall \tilde{x}_{old} \in HR_{old}$
Compute histogram of LV_{PI}
Estimate parameter ν based on histogram of LV_{PI} (see Figure 7.2)
Solve minimization problem (7.6) in order to find $\widetilde{w}$
Get parameter ρ through (7.12)
The personalized SVM solution SVM_{pers} is then defined by $\widetilde{w}$ and ρ
Real-time analysis of HR data HR_{new} ***from patient*** P_p
for each new HR measurement HR_k within HR_{new} at time t_{HR_k} **do**
  Extract significant HRIs (see Algorithm 3.1)
  if the end of a significant HRI is detected
    Extract selected features from the significant HRI
    Classify using SVM_{pers}
    if class output of classifier = 1
      Alarm!
    end if
  end if
end for

approach. These nights are chosen at random from the available nights in the training set, and simulations are repeated 100 times.

7.3.5 Performance evaluation

Similarly as in the previous chapters, the performance of the proposed methodology is done by using the sensitivity, false alarm rate (FAR) and positive predictive value (PPV). The FAR is expressed as false positives per night (=8 hours). Seizures are assumed to be detected if an alarm is generated within 3 minutes of the seizure onset [181]. Data from phase 1 is used for training the reference patient-independent classifier. The algorithms are evaluated on patients involved in phase 2 of the trial. For each patient in phase 2, 80% of the data is used for training the personalized classifier, and evaluated on the remaining 20% of patient data. This is then repeated five times so that every

Table 7.2: Overview of the overall results of the different applied methodologies in this chapter. PI: patient-independent, OC SVM: one-class SVM, OC SVM+TL: proposed one-class SVM with transfer learning.

Method	Sensitivity (%)	FAR (FP/night)	PPV (%)	#minor seizures detected
PI	**78.8**	1.11	21.1	**258 (69.4%)**
OC SVM	66.1	3.01	6.0	222 (59.7%)
OC SVM+TL	**78.8**	**0.75**	**28.1**	243 (65.5%)

data point is used once in the test set. Hyperparameter optimization is done by also performing these tests on the data from phase 1, and parameter values leading to optimal results in phase 1 are used for the evaluation of the patients in phase 2. The above-mentioned patient-independent classifier, personalized one-class SVM classifier and personalized one-class SVM classifier with transfer learning are evaluated in this chapter. Estimation of parameter ν in normal one-class SVM is done by using the procedure described in [71]. The detection of minor seizures was nor counted as a correct detection, nor as a false alarm.

7.4 Results

The reference patient-independent classifier resulted in on average 78.8% sensitivity with an average FAR of 1.11 FP/night on the test dataset (see Table 7.2). By applying the proposed transfer learning approach, the FAR decreased to 0.75 FP/night while maintaining the same sensitivity as the patient-independent approach. This is a reduction of 32.4% in FAR compared to the patient-independent classification results. The overall PPV improved from 21.1% to 28.1%. If the normal one-class SVM classifier is used for personalization, an overall sensitivity of 66.1% is obtained with on average 3.01 FP/night.

Table 7.3 shows the sensitivity results for the different major seizure types for the proposed transfer learning approach. Long generalized tonic seizures and generalized tonic-clonic led to the highest sensitivities (88% and 84.1%), whereas the results were lower for focal hyperkinetic and other major seizure types (72.5% and 71.7%).

Figure 7.3 shows the impact of the seizure type on the ROC curve for variable δ values in the proposed transfer learning approach. For a FAR higher than 0.5 FP/night, the sensitivity is the highest for the long generalized tonic seizures. For FAR lower than 0.5 FP/night, the sensitivity is highest for generalized tonic-clonic seizures. A generalized tonic-clonic sensitivity of 80% can be achieved

Table 7.3: Performance of the proposed transfer learning approach per seizure type. TC: tonic-clonic seizure, HK: focal hyperkinetic seizure, T>30: generalized tonic seizure with a duration of at least 30 s.

Seizure type	# seizures detected	# seizures in total	Sensitivity (%)
TC	243	289	84.1
HK	140	193	72.5
T>30	159	180	88.3
Other major	38	53	71.7
All major	580	615	81.1

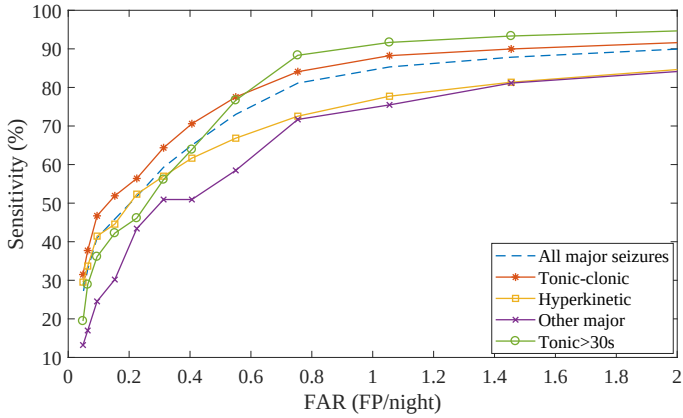


Figure 7.3: Impact of seizure type on performance of the proposed unsupervised transfer learning approach. The FAR and sensitivity results are obtained for different values for δ (range:[-2.9,2.9], in steps of 0.2, only showing results with a FAR smaller than 2 FP/night). Sensitivity results are here obtained as seizure-based overall results.

with around 1 false alarm every two nights with the proposed approach. The sensitivity is the lowest for all FAR for the other major seizure type. On average, generalized tonic-clonic seizures lead to higher peak heart rates and higher maximal heart rate gradients compared to the other investigated major seizure types as illustrated in Figure 7.4.

As discussed in Section 7.3.4, an additional simulation was performed to evaluate the impact of the amount of data on the performance of the personalized approach. The results for this simulation are illustrated in Figure 7.5, which shows the impact of the overall PPV for the different amounts of nights used for training. If only 1 night is used during training, the personalized approach leads

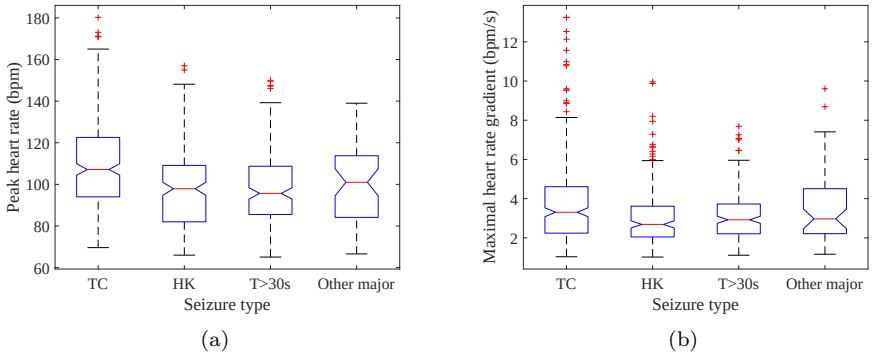


Figure 7.4: Boxplots of the ictal peak heart rate (a) and ictal maximal heart rate gradient (b) for different seizure types as found by the used preprocessing procedure. TC: tonic-clonic, HK: hyperkinetic, T: tonic.

to a decrease in performance compared to the patient-independent approach. Also the variation on results is very high. When two nights of data are used during training, the average PPV is already improved compared to the patient-independent performance. This increase in performance continues until around 2 weeks of nocturnal data, after which the proposed transfer learning approach converges towards its maximal PPV value. The variation on PPV results however continues to decrease further when 3 and 4 weeks of nocturnal recordings are used during training.

7.5 Discussion

7.5.1 Performance comparison of the different approaches

Table 7.2 shows that the proposed transfer learning approach indeed leads to a strong decrease in FAR compared to the patient-independent approach while maintaining the same overall sensitivity. The unsupervised procedure of standard one-class SVM however led to a decrease in performance compared to the patient-independent approach. In automated outlier detection, the classifier decides itself which data points are assumed to be the outliers based on the data distribution due to a lack of prior knowledge. This classifier thus does not know which outliers are of interest for this application. This is illustrated in Figure 7.6, which illustrates the classifier boundaries of the three mentioned classifiers. In one-class SVM, a boundary can be found around the center of the

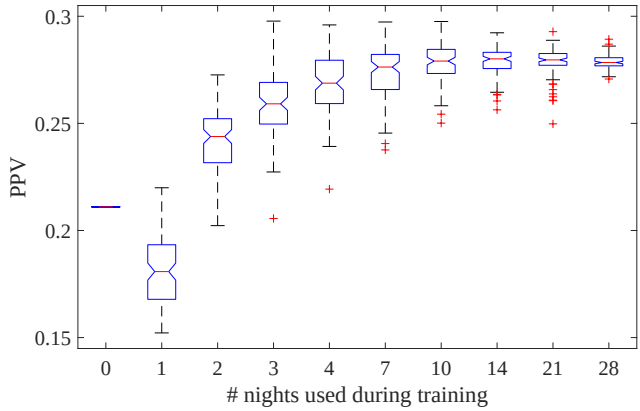


Figure 7.5: Impact of number of nights in training on performance of the proposed personalized transfer learning approach.

data cloud. This leads to outliers in every direction of this cloud, which leads to bad performance in this application. When transfer learning is used, the knowledge contained in the binary patient-independent classifier (summarized in w) allows to indicate where the outliers of interest are located. The proposed methodology for the estimation of ν allows to estimate the number of outliers in this area of interest, leading to an improved performance.

Another reason for the low performance of the normal one-class SVM classifier is the fact that epileptic heart rate behavior can not always be seen as clear outliers compared to non-epileptic behavior, even during nocturnal recordings. This makes it difficult to robustly estimate the ν parameter based only on the data distribution.

Another way in which ν can be defined is as function of an alarm rate which is acceptable for the user (e.g. 1 alarm/night) based on the alarm rate in the training set. That way an upper boundary can be put on the FAR of the personalized approach, but it is unknown then what effect this has on the sensitivity performance for that patient. This is due to the fact that some patients (e.g. younger patients) require the acceptance of a higher FAR in order to get a sufficient high sensitivity value.

The proposed method resulted in a decrease in FAR of 32.4% compared to the reference patient-independent classifier while maintaining a similar sensitivity. This is a similar FAR decrease as found in Chapter 4 by applying supervised transfer learning on day and night recordings from focal temporal lobe epilepsy

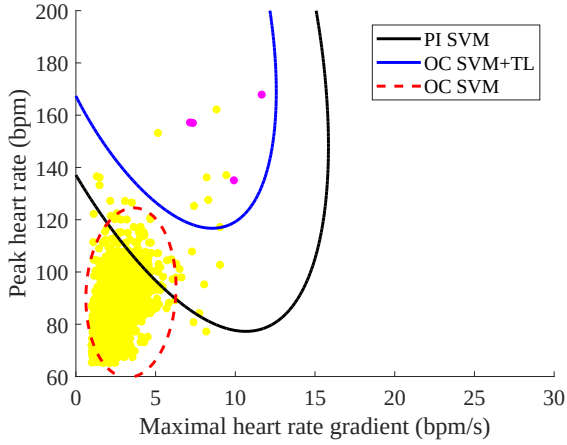


Figure 7.6: Example of the boundaries of the different classifiers. PI SVM: patient-independent SVM classifier, OC SVM+TL: one-class SVM classifier with transfer learning, OC SVM: one-class SVM without transfer learning using only patient-specific data. Data points in yellow indicate non-epileptic data points, pink data points indicate seizure data points.

patients. The FAR reduction is however smaller than the reduction found in Chapter 6. This is due to the fact that in this chapter, mainly adults are investigated, whereas in Chapter 6 only pediatric patients were evaluated. The heart rate increases (both ictal and interictal) in pediatric patients are however much stronger, also including even more inter-patient variability than in adults only. This is illustrated in Figure 7.7, which shows the differences in the ictal features for both this dataset (mainly adults) and the dataset used in Chapter 6 (all children). A larger variation can be seen in the features, certainly in the heart rate gradient. Therefore, a personalized approach has more added value for pediatric patients, as a patient-independent approach will lead to worse overall results than when applied to only adults (or a combination of children and adults). This can be seen by the fact that both the patient-independent and personalized approach led to a much smaller FAR compared to the results found in Chapter 6. Simple adjustments of thresholds in the preprocessing procedure *HRI-EXTRACT* can solve a part of the problem, but is insufficient to take into account the large variation, which depends on age, fitness, disease, drugs and many more factors. In future work, the proposed work will also be evaluated on pediatric patients to evaluate the performance for this population.

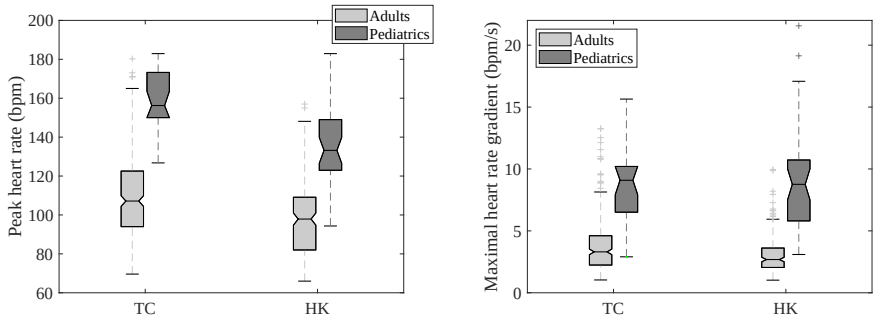


Figure 7.7: Boxplots of the ictal peak heart rate (a) and ictal maximal heart rate gradient (b) for generalized tonic-clonic (TC) and hyperkinetic (HK) seizures for both the mainly adults patients of the dataset mentioned in this chapter and the pediatric patients described in Chapter 6.

7.5.2 Impact of seizure type on performance

Table 7.3 shows that the obtained sensitivity values for generalized tonic-clonic and long generalized tonic seizures were on average around 10% higher than those for focal hyperkinetic and other major seizure types. In Chapter 6, it was also shown that the personalized approach there also led to a smaller sensitivity value for hyperkinetic seizures. There was however a smaller sensitivity for tonic and clonic seizures compared to the results found in this study. This is mainly due to the fact that a restriction is set on the minimal duration of the generalized tonic seizures in this chapter. Longer seizures tend to more often lead to a (stronger) heart rate increase, which increases the chance of successful detection with heart rate-based seizure detection.

Figure 7.3 shows that the sensitivity of the long generalized tonic seizures is higher than the sensitivity of generalized tonic-clonic seizures for a FAR higher than $\pm$0.5FP/night, although the difference is not found to be statistically different at the proposed δ value. This seems a bit contradictory to the results from Figure 7.4, which show that the extracted features are stronger for generalized tonic-clonic seizures. Figure 7.4 however only shows the results for seizures for which actually a heart rate increase was found by the *HRI-EXTRACT* preprocessing procedure. During some generalized tonic-clonic seizures, strong noise in the PPG (and thus also in the tachogram) occurred due to the strong ictal movements. This resulted either in no visible heart rate increase, or sometimes smaller heart rate increases than actually experienced

by the patient due to errors during the heart rate extraction. Therefore, no accurate detection was possible for these generalized tonic-clonic seizures, making it difficult to get a tonic-clonic sensitivity higher than 90%. In generalized tonic seizures, this problem occurred less due to less extreme ictal movements, allowing to obtain higher sensitivity values with a decent FAR performance. The used preprocessing procedure *HRI-EXTRACT* was initially implemented in Section 3.3.2 using ECG. It might however be that the method needs further finetuning in order to be able to deal with a tachogram built up from a PPG rather than ECG. This was however not the main target of this study, but might lead to an increased performance of the preprocessing procedure.

The other major seizure type showed a wide variation of seizure types. This means that it can contain seizure types that show strong and less strong heart rate increases. The other major seizure type has an overall lower percentage of ictal heart rate changes (around $\pm 85\%$ as indicated in Figure 7.3, similar as HK seizures but much smaller than generalized TC and long T seizures). Figure 7.4 shows that a larger variation in the extracted features (more clearly in the maximal heart rate gradient) is found in this other major seizure type compared to the focal hyperkinetic and long generalized tonic seizures. This is due to the fact that some of the other major seizures include seizures with a strong heart rate increase (e.g. some of the tonic-clonic-like seizures with unclear semiology) and seizures with rather small or no heart rate increases (e.g. clusters of minor seizures). This seizure type also showed to be the hardest to detect accurately with varying accepted FAR as illustrated in Figure 7.3.

Although the main focus of this chapter was the detection of major seizures, still 65.5% of the minor seizures were detected with the proposed personalized approach. Alarming in case of such a minor event is still useful to comfort the patient, but is less crucial compared to the detection of major types. It is however interesting to also have a notion of the minor seizure frequency as well, for example for seizure diary applications.

7.5.3 Impact amount of data on performance

One of the added values of transfer learning is that it allows to personalize the heart rate-based seizure detection with a limited amount of data. Figure 7.5 illustrates the effect of the amount of patient-specific data on the performance of the proposed unsupervised transfer learning approach. If only 1 night of data is used, a decrease in performance is found. This is mainly due to the fact that during one night, an insufficient amount of (both ictal and inter-ictal) heart rate increases is typically found, which leads to an insufficiently robust

personalized classifier. The estimation of parameter ν can only be trustworthy if a sufficient amount of such patient-specific heart rate increases is collected.

By using only two nights of data, the average performance is already increased compared to the reference patient-independent performance. Parameter ν can be more robustly estimated, and thus leads to an improved performance, without requiring any recorded seizures in the training set. By further including more nocturnal recordings in the training set, the average performance continues to increase, whereas the variability on the PPV results decrease, leading to a more robust result. After 2 weeks of training, the average results however stabilizes, but the variability of the results continues to decrease. These results show that the method is able to lead to a strong increase in performance after only a limited amount of nocturnal recordings.

In the proposed approach using two features, performance convergence was found after around 2 weeks of patient-specific data. If more features are used in this approach, potentially a higher performance can be obtained, but it will also require more time for a robust personalized classifier. For example taking into account the duration of sympathetic activation or a more detailed measurement of parasympathetic activity (i.e. high-frequency monitoring) could lead to an increased performance [38, 99, 142].

7.5.4 General discussion

Despite the fact that the proposed unsupervised personalization approach led to an increased performance, also this approach still generates too many false alarms on average for usage in practice. If however this personalized approach is used as part of a multimodal setup (for example for the detection of generalized tonic-clonic seizures), the multimodal performance can be improved with a similar gain as the performance increase found in the unimodal case. The method can also be applied to other modalities such as EMG and accelerometry.

The only drawback of this approach compared to the unsupervised approach discussed in Chapter 6 is the fact that it needs to be trained offline. Therefore, some performance is lost during the night itself between the offline updates. However, from Figure 7.5 it can be assumed that this impact would be minimal after 2 week of personalization, which is an acceptable waiting term for refractory epilepsy patients. As the proposed approach only works for nocturnal monitoring, the algorithm can easily be offline personalized further when the device is not worn by the patient. This retraining can for example be done when recharging the device during the day.

The unsupervised approach is useful as it does not allow annotations to be made

by either expert clinician or patients themselves. This removes the problems of either having a small amount of accurately annotated data (by clinicians) or a large amount of inaccurately annotated data (by the patient). Certainly for patients with intellectual disabilities, this unsupervised approach is of added value, in particular for nocturnal recordings.

7.6 Conclusion

A new unsupervised transfer learning solution is proposed which is able to adapt a binary support vector machine classifier using a limited amount of unannotated data. It was tested to personalize a nocturnal heart rate-based seizure detector with a limited amount of unannotated patient-specific data. The proposed unsupervised personalization approach led to a false alarm reduction of around 32.4% compared to the reference patient-independent classifier. Personalization already led to an increased performance after only 2 nights of data, which illustrates that the proposed method can be used for a fast and accurate personalization of heart rate-based seizure detection.

Chapter 8

Multimodal Generalized Tonic-Clonic Seizure Detection

The previous chapters discussed heart rate-based seizure detection, both patient-independent and personalized approaches. The results showed that in order to reach a performance good enough for usage in practice, multimodal approaches are required. In this chapter, first unimodal seizure detection algorithms using accelerometry, electromyography and electrocardiography are compared for the detection of nocturnal generalized tonic-clonic seizures in pediatric patients. Different sensor locations are evaluated for these modalities. Also the combination of these sensors is evaluated in a newly proposed multimodal seizure detection algorithm which takes into account the differences in time between ictal activations in the different modalities. All sensor combinations are compared in terms of performance, detection delay and wearability. Also the added value of a personalized unimodal approach in a multimodal setting is illustrated.

An introduction to seizure detection using electromyography and accelerometry is given in Section 8.1. The dataset used in this chapter is described in Section 8.2. The different unimodal and multimodal seizure detection algorithms are described in Section 8.3. The obtained results are shown and discussed in Section 8.4 and 8.5. This work is partly described in [46].

8.1 Introduction

ECG-based seizure detection has been discussed thoroughly during the previous chapters. However, seizure detection can be done with a large variety of modalities. Two of the most commonly used modalities for seizure detection in a home environment are electromyography (EMG) and accelerometry (ACM) [184]. In this chapter, the detection of nocturnal tonic-clonic seizures is evaluated for these two modalities, together with the ECG-based seizure detection described in Chapter 3. For the first time in literature, unimodal algorithms using these three modalities with sensors placed on different locations on the body are compared for the detection of generalized tonic-clonic seizures. These seizures are analyzed as they are the most interesting seizure type to detect in practice, and also have the highest possibility for accurate detection using these three modalities [113]. They will also be combined with each other in order to reach a real-time multimodal seizure detection algorithm.

8.1.1 Tonic-clonic seizure detection using EMG

EMG measures muscle activity, and thus allows to detect muscle contractions by placing electrodes on the skin (so-called skin EMG, or sEMG). Therefore, they are ideal for the detection of seizures which involve stiffening of a particular muscle. Often, either the deltoid or biceps muscle is analyzed in automated seizure detection due to the frequent stiffening of these muscles during generalized tonic(-clonic) seizures. In generalized tonic-clonic seizures, epileptic EMG activity is mainly visible during the tonic phase of the seizure. Figure 8.1 illustrates the ictal EMG activity during a generalized tonic-clonic seizure, which is mainly visible in the high frequency activation in the EMG signals during the tonic phase of the seizure.

The amount of research in the literature in this field is relatively limited . In [31], tonic-clonic seizure detection is performed by putting a fixed threshold on the zero-crossing rate in a specific time window. A wide variety of EMG features are extracted in [126], and feature selection selection was performed to select an optimal subset of features for EMG-based tonic-clonic seizure detection. Different frequency features were used in [112] for automated detection of tonic seizures. In [174], seizure detection is performed by comparing the activity with a maximum voluntary contraction (MVC) for each individual patient, alarming when this new activity was found to be sufficiently different compared to this MVC.

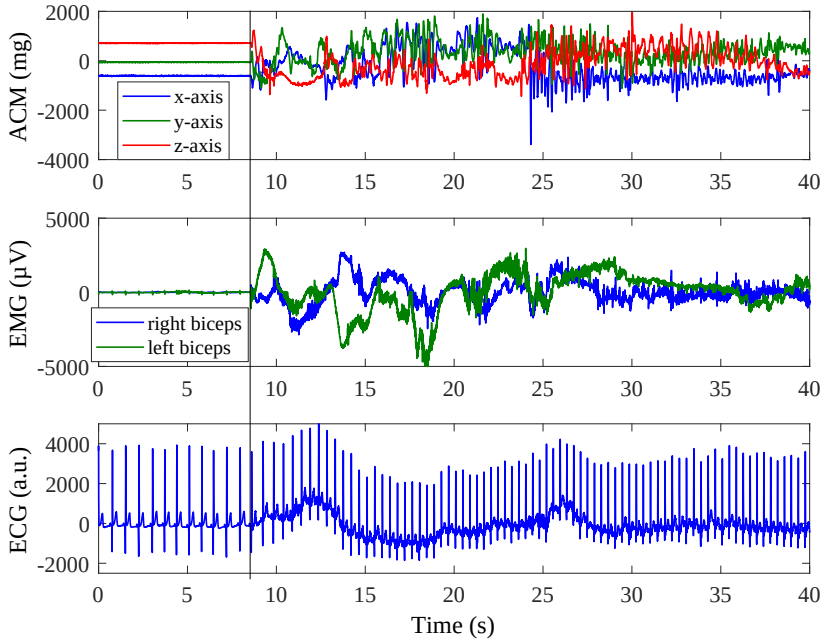


Figure 8.1: Illustration of ictal activity in accelerometry (ACM), electromyogram (EMG) and electrocardiography (ECG) during a tonic-clonic seizure. The vertical red line indicates the clinical seizure onset. The tonic phase of this seizure ends at 24 s, after which the clonic phase starts. During the tonic phase, the high frequency activity of the EMG is increased. During the clonic phase, strong periodic patterns can be seen in the three ACM channels. During both tonic and clonic phase, the heart rate gradually increases (RR intervals decrease).

8.1.2 Tonic-clonic seizure detection using accelerometry

Accelerometers (ACM) measure the acceleration of the sensors in the different directions. When attached to a patient, it allows the detection of seizures with strong motor involvement. Examples of seizures that can be detected using ACM analysis are generalized (tonic-)clonic seizures and focal hyperkinetic seizures. They both typically cause strong amplitude movements following a periodic pattern.

Generalized tonic-clonic seizure detection using ACM in the literature is mainly done in a window-based approach. Different features that give information on the strength, frequency and direction of the patients' movements can be

extracted and classified using different types of machine learning techniques [9, 11, 35, 36, 117, 126, 135, 137, 185]. Similarly as in the EMG-based seizure detection, mostly patient-independent approaches are tested on adult patients in the literature. Sensor placement is typically done around the wrist, but also has been tested on the upper arm, ankles and in smartphones [7, 11, 126, 150, 181, 184]. Similar as for EMG-based seizure detection, testing in the literature was mostly performed on adult patients rather than pediatric patients.

8.1.3 Multimodal seizure detection

As already discussed in Section 2.4.2, two types of multimodal combination can be used for combining different modalities. In late integration, the unimodal features are extracted and classified per individual modality using an unimodal decision step (e.g. a classifier or a rule-based decision). The different unimodal outcomes are then combined in an extra combinatory phase, after which a multimodal decision can be found (see Figure 2.13(a)). In early integration, features are extracted from the different modalities, and are directly given as input to a multimodal decision step, which directly results in a multimodal decision (see Figure 2.13(b)).

In literature, mostly late integration approaches are used to combine ACM, EMG and/or ECG sensors. One of the reasons for this is that for the early integration approach, it is important to have synchronous activity in the investigated modalities. This is however rarely the case, as for example EMG and ACM are rarely fully activated at the same time during a generalized tonic-clonic seizure (see Figure 8.1). As most features from different modalities do not have correlation with each other, combining features from different modalities does not result in much added value compared to combining unimodal classifier outcomes. A late integration approach is however typically easier to train with a relatively limited amount of available data, and also easier to interpret than early integration approaches.

The late integration combination of different unimodal outcomes is typically done using a rule-based approach in the literature. Applied procedures in the literature are for example the AND-operator (requiring all modalities to be active during a seizure) , majority voting (requiring the majority of sensors to be ictally active) or the OR-operator (requiring only 1 sensor to be active) [66, 76, 126, 181]. In [28], a specific time pattern between oxygen saturation, heart rate and EDA features are required in order to cause a multimodal alarm. This pattern was however only found in a subset of the patients, and is thus not widely applicable for all patients having generalized tonic-clonic seizures.

The combination of accelerometers, EMG and ECG was never realized using an AND-operator approach in the literature. In this chapter, a variation of the AND-operator combination is used in this chapter. The AND-operator assumes all used sensors to be active, which is ideal for the detection of generalized tonic-clonic seizures in which it is assumed that all three evaluated modalities are activated by epileptic seizures. The standard AND-operator however assumes all modalities to be active at the same time. This is only usable when sensors from the same modality are combined (e.g. 2 EMG sensors). Therefore, a multimodal combination is proposed here which allows to take into account these time differences in epileptic activity in the different modalities.

8.2 Data acquisition

The different unimodal and multimodal algorithms are evaluated on data obtained at the Pulderbos revalidation center for youth and adolescents. Seven pediatric patients with 22 generalized tonic-clonic seizures are monitored during 224 hours of nocturnal data. Each patient had sensors from three different modalities attached. Four accelerometer sensors were attached to each wrist and ankle (right wrist: RW, left wrist: LW, right ankle: RA, left ankle: LA), two EMG sensors to the biceps muscle on each arm (right biceps: RB, left biceps: LB), and single-lead ECG obtained from the chest (see Figure 8.2). Sensor locations were chosen because they are most often ictally activated. All signals were acquired with 250 Hz sampling frequency.

Seizures were annotated by expert clinicians using the gold video-EEG monitoring standard. Only generalized tonic-clonic seizures of at least 10 s duration are evaluated in this study. All participants or their legal guardian gave written consent to this study.

8.3 Methodology

8.3.1 Seizure detection using EMG

The preprocessing procedure for tonic-clonic seizure detection is based on [126] according to the following steps. First, a notch filter is used to remove the 50 Hz power-line interference from the EMG signal. Then, a Butterworth high-pass filter with a cut-off frequency of 20 Hz is used to remove movement artifacts and baseline wander.

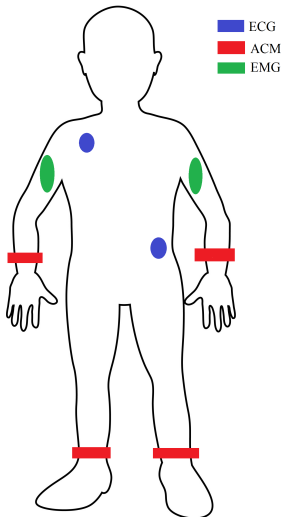


Figure 8.2: Sensor placement for the different modalities in the dataset used in this chapter. Accelerometers (ACM): placed at both wrists and ankles. Surface electromyogram (sEMG): located at both biceps, consists of two sensors, one placed at the belly of the muscle and one placed on the tendon. This results in one EMG signal per biceps. Electrocardiogram (ECG): two sensors placed at the chest, one near the right shoulder and one near the left hip, which lead to the lead II ECG signal.

A window-based approach is used for the EMG-based seizure detection algorithm. Windows of 2 s with 75% overlap are used. An extra preprocessing step is performed in order to lower the class unbalance for later classifier training. Feature extraction is only performed if the standard deviation in the window is larger than 30 μV. During feature extraction, the zero-crossing rate of the second derivative signal and the Wilson amplitude are extracted (see Appendix A for a detailed description of the used features). These features are selected based on the feature selection procedure discussed in [126].

The extracted features are then classified using a weighted support vector machine (SVM) classifier, defined by the following minimization problem (see Section 3.3.5 for a detailed description of all variables)

$$\min_{w,b,\xi} \frac{1}{2}||w||^2 + C\sum_{i=1}^{N} c_i\xi_i$$
$$s.t. \quad \begin{cases} y_i(w^T\varphi(x_i)+b) \geq 1-\xi_i \\ \xi_i \geq 0 \end{cases} \quad ,\forall i \in [1,N] \tag{8.1}$$

with

$$c_i = \begin{cases} \gamma\frac{(N^+ + N^-)}{2N^-} & : y_i = -1 \\ \frac{N^+ + N^-}{2N^+} & : y_i = +1 \end{cases} \tag{8.2}$$

with N^+ and N^- indicating the number of datapoints belonging to classes +1 (non-seizure) and -1 (seizure) respectively. During training, windows which contain at least 50% seizure activity are annotated as epileptic. Optimization of hyperparameters C, γ and RBF kernel parameter σ are automatically defined by using leave-one-out crossvalidation on the training set. An alarm is generated if 8 subsequent EMG windows (6 s duration in total) are classified as epileptic activity as described in [126].

8.3.2 Seizure detection using accelerometry

The preprocessing procedure for accelerometry-based seizure detection is also based on the implementation of [126] according to the following steps. First, a Chebyshev low-pass filter with a cutoff frequency of 45 Hz is applied to remove the 50 Hz powerline and aliasing artifacts from the ACM signal. After downsampling the ACM signal with a factor 2, a Chebyshev high-pass filter with a cutoff frequency of 0.2 Hz is applied to remove baseline wandering.

Similarly as in the EMG-based detection, 2 s windows are used for further feature extraction. Also here an extra preprocessing step is performed in order to limit the later class unbalance for classification training: Only windows in which the standard deviation of all channels was higher than 10 mg are used for further processing. This value was experimentally set based on the findings in [126].

Feature extraction is performed for all channels individually. Four different features are extracted from each channel, resulting into 12 features per window per sensor. The extracted features are based on the findings from [126], which are the number of local maxima in the ACM channel, Shannon entropy of diagonal line lengths in a recurrence plot [52, 150], frequency of the first local peak in the spectrum and power within pseudofrequency range 9-22.5 Hz from

a continuous wavelet transform (see Appendix A for a detailed description of these features).

These features are then classified using a patient-independent SVM classifier as defined in (8.1). The accelerometry-based seizure detection algorithm then generates an alarm if 8 subsequent windows are classified as epileptic activity by the classifier as described in [126].

8.3.3 Seizure detection using ECG

The ECG-based seizure detection is based on the procedure mentioned in Chapter 3. First, the heart rate is extracted in real-time from the ECG by an R peak detection algorithm which is based on adaptive filtering on the derivative ECG signal. Then, strong heart rate increases (HRI) are extracted using the *HRI-EXTRACT* preprocessing procedure described in Section 3.3.2.

If such a strong HRI is found, the peak heart rate during the HRI and the average heart rate and standard deviation of the pre-HRI period are extracted from the tachogram. These features are also classified using a weighted SVM classifier as defined in (8.1). If a HRI is classified as an epileptic HRI, an alarm is generated by the ECG-based seizure detection algorithm.

8.3.4 Personalized seizure detection using ECG

The personalized heart rate-based approach mentioned in Chapter 6 is also applied here in order to illustrate its added value in a multimodal algorithm. The heuristic unsupervised approach is preferred here over the transfer learning approach due to the limited amount of data per patient in the used dataset, which is too small for most patients. The method quickly adjusts to patient-specific characteristics after only a couple of hours on average, making it ideal for such a small amount of data per patient (only a couple of nights per patient in this dataset).

8.3.5 Multimodal methodology

In this study, only combinations of two sensors are investigated thoroughly. The combination of different modalities is performed here with a late integration approach. The outcome of the unimodal approaches (binary epileptic (-1) or non-epileptic (1) outcome) is combined by an approach similar to an AND-operator, requiring both modalities to be active in order to result in a multimodal alarm.

Table 8.1: Overview of the used δ values for multimodal seizure detection (see (8.3)).

Modality combination	δ (s)
EMG+EMG	5
ACM+ACM	15
ACM+ECG	20
ACM+EMG	20
ECG+EMG	20

A simple AND-operator however requires the two modalities to be active at the same time, which is not always the case as mentioned in Section 8.1.3. In order to solve the issue of asynchronous epileptic activity in different modalities, a variation to the classical AND-operator is applied. Let $z_s(t)$ be the unimodal output ($z_i(t) \in \{-1, 1\}, \forall t \geq 0$) for sensor s at time t. This output is then adjusted so that its minimal value is maintained for a certain time δ, creating a new signal $\hat{z}_s(t)$, defined as

$$\hat{z}_s(t) = min(z_s(\tau) | \tau \in [t - \delta, t]), \forall t \geq 0. \tag{8.3}$$

This way, epileptic activity in a sensor is maintained in unimodal memory for a period δ. Parameter δ depends on the applied sensor combination, not only the used signal s. Combining an EMG sensor with an ECG sensor would require a different optimal δ value than when two EMG sensors are combined due to the difference in activation pattern. The used δ values for the different combinations are experimentally chosen based on the typical differences in detection delays found for the different unimodal algorithms (see Table 8.1 for the used values in this chapter). In general, the combination of sensors from the same modality require lower δ values due to increased epileptic activity synchronization in those signals.

A multimodal alarm is generated at time t if $\hat{z}_s(t) = -1$ for both sensors $s = 1, 2$. An overview of this multimodal approach is depicted in Figure 8.3. Three sensors can be combined in a similar way. Here, three sensors are combined in two steps. First, two sensors are combined as described above. Afterwards, the result of this combination is then combined with the third sensor as if the first combination was an individual sensor as in Figure 8.3.

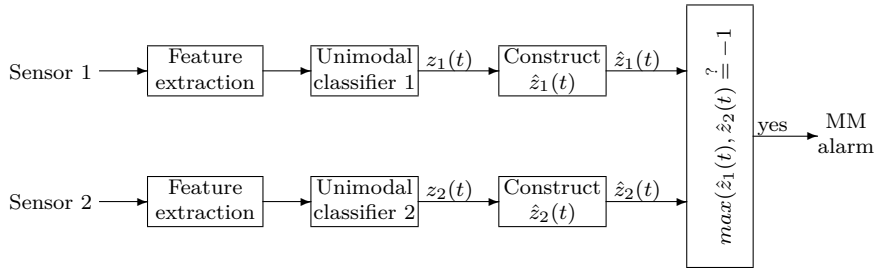


Figure 8.3: Overview of the proposed multimodal (MM) seizure detection algorithm.

8.3.6 Algorithm evaluation

The different unimodal patient-independent algorithms are evaluated using leave-one-patient-out crosstesting. Unimodal algorithms are trained per sensor, using only data acquired from the sensor also used for testing. This is done in order to illustrate the training and testing procedure if only this sensor would be available for training and testing. Similarly as in the previous chapters, the sensitivity, false alarm rate (FAR, expressed in (FP/night, with a night defined as 8 hours [181]), positive predictive value (PPV), F_3 score and detection delay are used for evaluating the different unimodal and multimodal seizure detection algorithms.

8.4 Results

8.4.1 Unimodal results

Table 8.2 gives an overview of the different unimodal results for tonic-clonic seizure detection using either ECG, EMG or accelerometry sensors. All algorithms led to sensitivities of at least 80%. Because the sensitivity and FAR vary between the different algorithms, the ROC-like curves of all unimodal algorithms are shown in Figure 8.4, which allows to give a better algorithm comparison. From this figure, it can be seen that the ACM-RW sensor (sensitivity: 95.5%, FAR: 5.6 FP/night) results in a better overall performance compared to the other evaluated sensors. Based on the F_3 scores in Table 8.2, also the ECG (sensitivity: 100%, FAR: 6.96 FP/night) and EMG-RB (sensitivity:86.4%, FAR: 4.93%) can be seen as one of the better unimodal algorithms.

Table 8.2: Overview of the unimodal and multimodal results. Bold values indicate the best value for a specific evaluation metric for the unimodal and multimodal algorithms. Se: sensitivity, FAR: false alarm rate, PI: patient-independent, PS: patient-specific.

		Used sensors	Se (%)	FAR (FP/night)	Delay (s)	F_3 score (%)
Unimodal	PI	ECG	**100.0**	6.82	14.0	53.5
		ACM-RW	95.5	5.43	15.8	**56.6**
		ACM-LW	86.4	10.36	14.9	37.5
		ACM-RA	90.9	9.86	19.2	40.5
		ACM-LA	90.9	11.54	15.2	37.0
		EMG-RB	86.4	**4.93**	8.7	53.5
		EMG-LB	95.5	6.97	**8.3**	50.7
	PS	ECG-Personalized	95.5	**2.75**	15.2	**70.7**
Multimodal	Same modalities	ACM-RW+ACM-LW	86.4	2.18	18.6	68.3
		ACM-RW+ACM-LA	86.4	1.61	16.3	72.5
		ACM-RW+ACM-RA	86.4	1.57	20.9	72.8
		ACM-LW+ACM-RA	77.3	2.64	20.8	58.8
		ACM-LW+ACM-LA	86.4	3.00	18.1	63.1
		ACM-LA+ACM-RA	86.4	6.90	22.2	46.3
		EMG-RB+EMG-LB	81.8	1.21	**11.7**	72.0
	Different modalities	ECG+ACM-LW	81.8	0.82	20.4	75.3
		ECG+ACM-RW	86.4	0.50	19.8	82.2
		ECG+ACM-LA	86.4	0.82	21.5	79.1
		ECG+ACM-RA	90.9	0.61	22.0	**85.1**
		ECG+EMG-RB	86.4	1.04	14.6	77.2
		ECG+EMG-LB	**95.5**	1.46	15.4	80.8
		ACM-RW+EMG-RB	81.8	0.25	17.6	80.7
		ACM-RW+EMG-LB	90.9	0.64	19.1	84.7
		ACM-LW+EMG-RB	77.3	0.64	18.7	73.0
		ACM-LW+EMG-LB	81.8	1.43	17.8	70.3
		ACM-RA+EMG-RB	86.4	**0.29**	20.1	84.4
		ACM-RA+EMG-LB	90.9	0.68	19.5	84.4
		ACM-LA+EMG-RB	77.3	0.43	18.1	74.9
		ACM-LA+EMG-LB	86.4	0.79	18.1	79.5

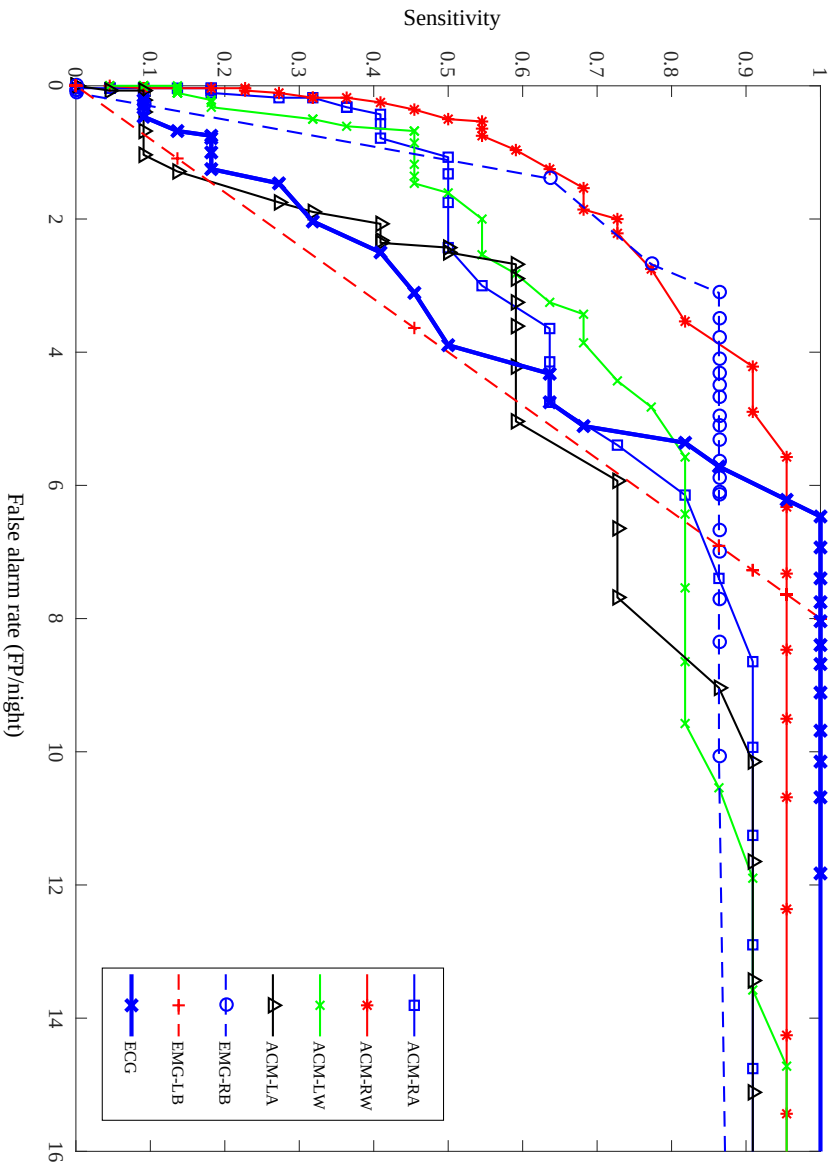


Figure 8.4: Sensitivity over false alarm rate ROC-like curve for the evaluated unimodal seizure detection algorithms.

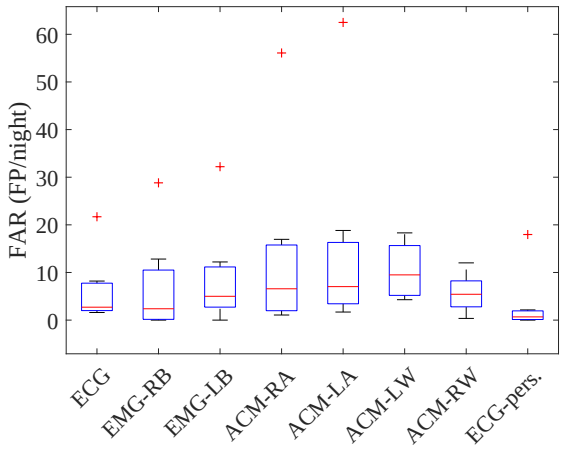


Figure 8.5: Boxplots indicating the inter-patient FAR variability between the different unimodal seizure detection algorithms. ECG-pers.: personalized ECG-based seizure detection method described in Section 8.3.4.

If for every patient the best unimodal algorithm is used, an overall sensitivity of 100% would be found with on average 2 FP/night. In this case, 5 different sensors would be taken as best sensor, with only the ECG and ACM-RW sensor to be chosen twice as the best sensor. Overall, the inter-patient FAR variability showed to be very high for most unimodal algorithms, with the smallest variations found for ECG and ACM-RW (see Figure 8.5).

The EMG-based algorithms led to the lowest detection delay with an average delay of ±8 s compared to the EEG seizure onset (see Figure 8.6). Also the delay variability showed to be the smallest for the EMG-based seizure detection. The ACM-based algorithms showed to be the slowest on average (15.0-19.3 s), also leading to the largest variability.

8.4.2 Multimodal results

Table 8.2 also shows the results for the different combinations of two sensors in a patient-independent multimodal algorithm. The best F_3 score was found to be ECG+ACM-RA, which led to 90.9% sensitivity and 0.60 FP/night. Also the combination of ACM-RA with EMG-RB led to one of the best multimodal solutions (sensitivity: 86.4%, FAR: 0.32 FP/night). If these three most successful

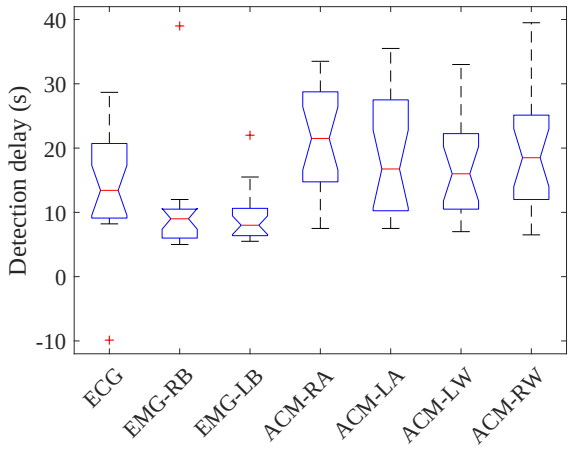


Figure 8.6: Boxplots for detection delay for the different unimodal seizure detection algorithms.

sensors were combined, a sensitivity of 86.4% was found with no false alarm in the entire dataset.

From Table 8.2, it can also be seen that combining sensors from the same modality (ECG, ACM or EMG) led on average to a lower performance than when sensors from different modalities were combined (average F_3 score: 64.5% vs. 79.4%). All combinations of sensors from multiple modalities led to at least 75% reduction in FAR, whereas 59% of the combinations led to at least 90% reduction in FAR. PPV values often increase from around 10% to up to 50%. In 18/21 of the multimodal combinations, the sensitivity remained higher than 80%.

The resulting multimodal detection delays depend on the slowest modality used in the combination of sensors. As can be expected from the unimodal results in Figure 8.6, the combinations including accelerometry lead to the slowest detections with an average delay of 19.4 s (see Table 8.2). The fastest detection is obtained by the combination of the two EMG sensors, which led to an average detection delay of 11.7 s.

8.4.3 Unimodal personalization in a multimodal combination

In order to demonstrate the added value of personalizing an unimodal algorithm in a multimodal setting, also the personalized ECG approach mentioned in

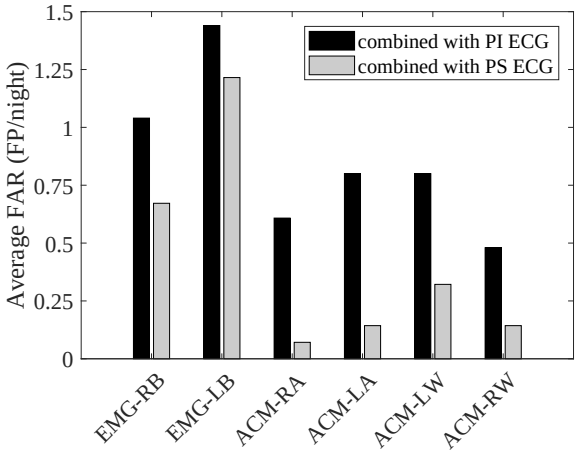


Figure 8.7: Average false alarm rate (FAR) for the combinations with patient-independent (PI) and patient-specific (PS) ECG-based seizure detection.

Chapter 6 is applied here, and combined with the other modalities. Unimodally, the personalized approach leads to on average 95.5% sensitivity with on average 2.96 FP/night (see Table 8.2). This is a FAR reduction of around 57% compared to the patient-independent (PI) ECG approach, but with a small decrease in sensitivity (-4.5%, caused by one missed seizure). Unimodally, this personalized approach clearly outperforms all other unimodal patient-independent performances. Figure 8.5 also shows that this personalized algorithm leads to a strong decrease in inter-patient FAR variability with the exception of one strong outlier. The unimodal personalized approach even led to a better performance than some multimodal combinations.

If this personalized approach is combined with the other unimodal algorithms, also the FAR decreases in this multimodal setting. This is illustrated in Figure 8.7, which shows the FAR without and with personalized ECG-based seizure detection (notice that for all combinations, the combinations including the personalized ECG approach lead to 4.5% less sensitivity than the original fully patient-independent combinations). For combinations with ACM-RA and ACM-LA, there is for example a strong decrease in FAR (0.61 to 0.07 FP/night and 0.80 to 0.14 FP/night). For the combinations with EMG, the FAR decrease is however less strong (EMG-RB: 1.04 to 0.67 FP/night, EMG-LB: 1.44 to 1.22 FP/night). The average FAR decrease by applying the personalized ECG approach instead of the PI ECG approach in a multimodal system is 59%.

8.5 Discussion

8.5.1 Comparison of different unimodal approaches

Figure 8.4 shows the sensitivity in function of the FAR for the different unimodal seizure detection algorithms. Based on the overall performance, ACM-RW seems to lead to the most robust algorithm, leading to good sensitivity performances for all FAR. Also the EMG-RB sensor leads to one of the best unimodal performances. However, in both cases, the algorithms seem to have a certain sensitivity maximum which is hard to improve with increasing FAR (95.5% and 86.4% respectively). Similar effects can be seen for all non-ECG sensors, for which none are able to reach 100% sensitivity with a FAR lower than 16 FP/night. This means that for these sensors, some seizures are nearly impossible to detect accurately. There are three main reasons for this. First, ACM sensors might not be activated due to obstruction of the monitored limb. If for example the person is lying on his left arm during a seizure, this sensor might not be able to detect this seizure due to a lack of movement. Second, it might be that the monitored limb is not activated at all during a seizure. Sometimes, only a subset of limbs is activated, which might cause no specific pattern in those sensors. Third, in some cases, either the tonic or clonic phase of the generalized tonic-clonic seizures are very short (<5 s). Therefore, it is difficult to accurately detect these seizures by using only ACM or EMG sensors, as they are typically only strongly activated during one of these phases.

ECG-based seizure detection on the other hand led here to 100% sensitivity with a relatively limited FAR. The downside of ECG-based seizure detection is however its sensitivity to noise. This can cause a lot of false alarms, but also could lower the sensitivity (see Chapter 7). This is illustrated in Figure 8.4, which shows that for a small FAR, the ECG leads almost to the worst sensitivity results compared to the other unimodal algorithms.

ECG-based seizure detection however led together with ACM-RW to the smallest inter-patient variability in FAR (see Figure 8.5). ECG led to a smaller median FAR than ACM-RW, but had a large outlier of 22 FP/night. This outlier was caused by a patient for which the ECG recordings were very noisy, which caused a lot of false alarms due to R peak detection errors. The inter-patient variability was larger in the other unimodal modalities. For some of these algorithms, they however also led to perfect performance for some patients, but caused too many false alarms in other patients. This illustrates that also in EMG and ACM-based seizure detection, personalization could be beneficial for improved unimodal performance. The methods proposed in Chapter 4 and Chapter 7 could be used for a fast and robust personalization in these modalities as well.

ACM-RW thus showed to lead to the best overall performance, better than for example the ACM-LW sensor. This is in contradiction with [126], which found that the non-dominant hand (normally left) should lead to a better performance due to less non-epileptic activity. In this study however, sensor classifiers are trained by using only data from that particular sensor. This led to a larger amount of training data of a wider variety of movements for the RW sensor compared to the LW sensor, leading to a better trained classifier for the RW sensor. If however the RW classifier is applied to the LW sensor data, a nearly similar performance can be found for both RW and LW. A similar effect can also be seen for EMG-RB and EMG-LB sensors, in which the large difference in performance can be strongly reduced in a similar way. No such effect is visible between the two ankle sensors, which are identified as the two worst performing sensors for unimodal seizure detection in this dataset.

From Figure 8.6 it can be seen that overall EMG-based seizure detection led to the fastest detection of around 8 s on average. However, the detection of seizures for accelerometers was on average at most 19.3 s. For monitoring applications in which the patient and caregiver are not in the same building, this difference in detection delay can however be neglected compared to the required traveling time of the caregiver. For direct intervention, these 10 s difference in delay can however be important if it is really required to arrive to the patient during the seizure rather than postically.

Unfortunately, the unimodal algorithms all led to a too large FAR for sufficient high sensitivity values in order to be used in practice as part of an unimodal seizure warning system. Some algorithms from literature were able to reach sufficiently good results for usage in practice for adult patients [10, 11, 81]. These results could however not be replicated here. The typical generalized tonic-clonic EMG pattern discussed in [10] for example was only found in around 50% of the seizures in this dataset. This is due to the fact that in children, the tonic-clonic patterns can be different compared to those in adults. The seizures in this dataset are also relatively short, with around half of the seizures having duration of at most 30s. Therefore, the time required for ACM and EMG activation is limited in duration in order to detect a sufficient amount of seizures. In the literature, the evaluated seizures are typically longer, and require longer ACM and EMG activation before an alarm is set. This could also be applied by the applied algorithms in this study, which would lead to a decreased FAR, but also a smaller overall sensitivity for shorter seizures [126]. It also leads to an increased detection delay.

Different unimodal algorithms were able to reach perfect performance in some patients. This ideal sensor varied however per patient, leading to 5 different ideal sensors on 7 patients in this dataset. Ideally, this sensor is known beforehand and can be chosen based on visual inspection. This might be the case for the

selection of ACM sensors, but it is more challenging for the selection of the EMG and ECG sensors. Therefore, it is advised to let the patient wear a variety of sensors for a number of nights until sufficient seizure data is gathered. The optimal sensor for that patient can then be selected in an manually or in automated way based on this collected dataset.

There is however an important caveat on this study. It has only been evaluated on a limited amount of patients (7) with a limited amount of tonic-clonic seizures (22). It is therefore difficult to draw hard conclusions from this study in terms of which sensor leads to the best performance. The found results however confirm the hypothesis that the ideal sensor for tonic-clonic seizure detection is patient-specific, and that accelerometers and EMG seem overall more suited than ECG for the detection of tonic-clonic seizures.

8.5.2 Multimodal tonic-clonic seizure detection

The results for the different multimodal combinations are shown in Table 8.2. In general, it can be seen that combinations of sensors of a different modality led to a better performance than the combination of sensors from the same modality. This is illustrated in Figure 8.8(a), which shows the average F_3 improvement by adding a sensor to an unimodal algorithm. When a sensor is added to an unimodal ACM algorithm, the increase in F_3 is much larger for EMG and ECG sensors ($\pm$ 37% increase) compared to another ACM sensor (22% improvement). A similar, though less strong in amplitude, effect can be seen if a sensor is added to an EMG electrode. This shows that sensors from different modalities should be used in a multimodal seizure detection system in order to obtain optimal performance. Figure 8.9 also illustrates this by showing the average F_3 scores for each sensor when combined with a different sensor. Also here it is clear that sensors from a different modality clearly lead to an increased performance compared to when two sensors from the same modality are combined.

Figure 8.8(b) shows the average F_3 scores over all combinations in which a specific sensor is used. A distinction is made between combinations of sensors from the same modality and different modalities. This again demonstrates that the combinations of different modalities led on average to better performances. Specially large differences were found for the two ACM ankle sensors, where a difference of around 30% was found between both types of combinations. On average, combinations including the ACM-RA sensor and an EMG/ECG sensor led to the highest F_3 score. This is remarkable due to the fact that unimodally, this sensor was one of the worst performing sensors. It has however much added value compared to the other sensor combinations. 5/7 sensors acquire signals from the upper part of the body, and causes that led to false alarms in one

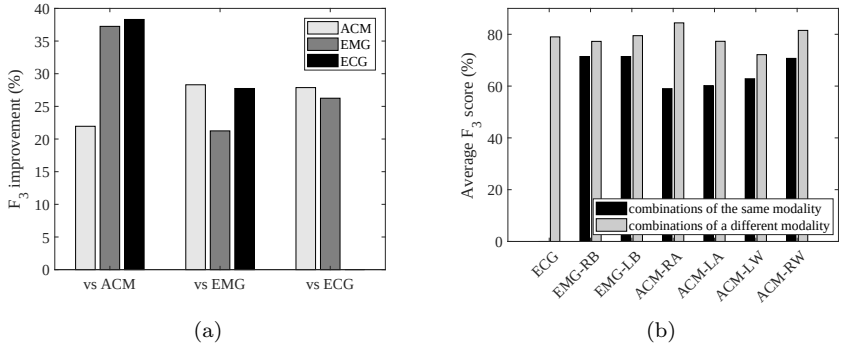


Figure 8.8: Impact of different sensors on the multimodal F_3 performance. (a) F_3 score improvement by including a sensor to an unimodal algorithm grouped per modality. (b) Average multimodal F_3 score for each sensor involved in the multimodal setting. Differentiation is made for combinations with a different modality and combinations of the same modality.

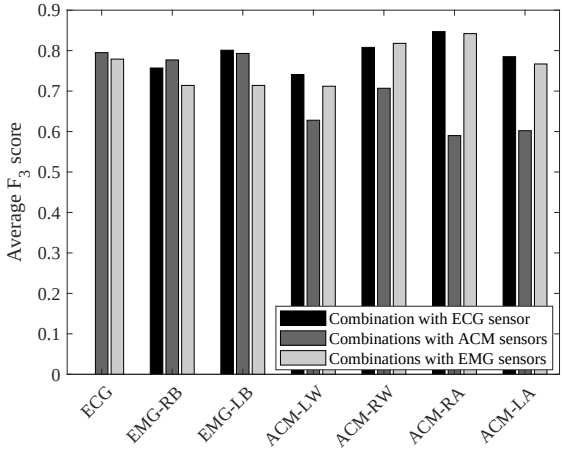


Figure 8.9: Average F_3 scores for each sensor when combined with another sensor, averaged out for each type of modality.

sensor often also caused false alarms in some of the other sensors. The ankle sensors were less likely to be involved in such motions, therefore less likely to also be affected by events that also caused false alarms in the upper body sensors. This shows that preferably sensors on different parts of the body are

used, which can also been seen in the results of Table 8.2. Sensors placed on the other side (left/right and upper body/ankle) often lead to better performances than sensors placed near to each other due to this overlap in false alarm causes.

8.5.3 Added value of unimodal personalization in a multimodal algorithm

The used personalized approach for ECG-based seizure detection resulted in a FAR decrease of around 57% unimodally. It clearly outperforms the other unimodal patient-independent algorithms. Nevertheless, the used procedure for personalization does not require any intervention from either patient or clinician, leading to personalization in a fully automated way. This means that in practice, both the patient-independent and personalized approach would work with the same amount of effort for the user. Therefore, it has only positive effects to use this personalized approach in a real-life setting, as it clearly leads to an improved performance, also in a multimodal setting. Similar unsupervised personalization approaches could and should be used for the other modalities as well, as also these modalities showed to have patient-dependency in their performance.

The effect of personalizing one modality in the multimodal setting is illustrated in Figure 8.7. Although there was overall a 59% reduction in FAR (similar as in the unimodal case), this reduction is not equally large for all combinations. For the combination with ACM-RA, this reduction was 88.4%, whereas the reduction for EMG-LB was only 15.6%. There are mainly 2 reasons for this variation. First of all, the gain by personalizing the ECG-based seizure detection is patient-dependent, not resulting in an equally large FAR decrease in all patients. For some patients, there was a strong FAR reduction (sometimes already coming near perfect detection), whereas for one patient, the FAR only dropped slightly towards still almost 20 FP/night. Similarly as for ECG-based seizure detection, also the other modalities show strong inter-patient variability in FAR (see Figure 8.5). If the reduction of the personalized ECG approach is mainly in patients for which the other modality already resulted in near perfect performance, then the personalization only has a limited added value (but also no negative effect) compared to the patient-independent alternative. This is for example the case for the combination with EMG-LB, in which most false alarms are generated by 1 patient for which both sensors cause a lot of false alarms. If however the personalization leads to a strong reduction in false alarms in patients for which the other modality gave bad results, then a much stronger FAR decrease can be found (e.g. the combination with ACM-RA).

The other reason is the fact that false alarms caused by noise in the ECG

(and thus also in the heart rate signal) can not be removed by personalizing the ECG-based seizure detection. This is due to the fact that these heart rate increases are often seen as very strong heart rate increases due to R peak detection errors, which makes it difficult to ignore them with the used heart rate features. Often noise in the ECG was triggered by strong motion in the upper part of the body, often also causing false alarms in one of the wrist or biceps sensors. This effect was less strong in the sensors on the ankles, which caused them to have more gain from personalization of the ECG-based seizure detector.

The best multimodal combination including the personalized ECG-based seizure detector (personalized ECG+ACM-RA) led to on average 86.4% sensitivity with 0.07 FP/night. This is only 2 false alarms in the entire dataset, and these false alarms are caused by other seizure activity (1 generalized tonic-clonic seizure shorter than 10 s and 1 versive seizure). This is a similar performance as found by combining the three best patient-independent unimodal sensor algorithms. This means that by personalizing one of the two unimodal algorithms, requiring a third sensor can be avoided, which is of course of great importance in terms of wearability. In practice, the amount of sensors should be limited to a minimum. This result shows that personalizing only one sensor could lead to requiring one sensor less on the patient's body without much loss in performance.

8.5.4 General discussion

The ACM-RW sensor was found to be optimal for unimodal seizure detection, although also the EMG-RB and ECG sensor showed to have good results. They however resulted in lower unimodal performances than often found in the literature [10, 11, 81]. This is due to the used dataset, which targets pediatric patients with often very short seizures. Pediatric patients showed to have not so clear motoric patterns as discussed in the literature for adults, which results in a worse performance due to this larger inter-seizure variation. In the literature, often a minimal duration of seizures is required for detection, which is typically longer than 20/30 s, whereas half of the investigated generalized tonic-clonic seizures in this dataset was shorter than 30 s. Also the ECG-based seizure detection leads to a larger FAR in pediatric patients than in adults due to the increased inter-patient heart rate variability (see for example the results shown in Chapter 7). It can be assumed that in adults, both unimodal and multimodal nocturnal seizure detection algorithms could lead to better results. This study however shows a first comparison between ACM, EMG and ECG for unimodal seizure detection with sensors placed on different parts of the body for pediatric patients.

Only generalized tonic-clonic seizure are investigated in this study. One of the advantages of ECG-based seizure detection is that it is usable on a wider range of seizure types than EMG and ACM, which can only be used for the detection of motor seizures. In some applications, the detection of a wider range of seizures can be interesting as well. Therefore, it is important to know if the patient is also interested in the detection of for example clonic seizures or focal impaired awareness seizures, as it can also influence the modality choice. In advanced multimodal alarm system, it is therefore advised to incorporate ECG. It allows to detect the widest range of seizure types, but also allows to monitoring other aspects of the patient, such as sleep analysis, stress monitoring and potential dangerous heart conditions. These are impossible with for example EMG and ACM sensors.

A comparison is done here only in terms of performance. It is of course also important to discuss the wearability of the different sensors. ECG is used for the mentioned heart rate-based seizure detection. ECG however requires two sensors to be placed on the chest (or one big patch in some ECG hardware), which is not comfortable for long-term monitoring. ECG could be replaced by PPG as for example done in Chapter 7, where PPG was obtained from the upper arm. This would make the data acquisition much less obtrusive for the patient. Available EMG-based seizure detection devices are currently still relatively bulky [10, 81]. Accelerometer devices can be made very small and subtle. They are typically also already available in most wearable devices (e.g. smartwatch), which allow easy reusage of other wearable devices.

The results show that sensors from different modalities placed in different parts of the body lead to the best multimodal results. In practice, this is however cumbersome, as it would require multiple devices on the patient's body. This would strongly affect the wearability of the alarm system for a patient. If however the sensors are found to be sufficiently subtle and easy to wear, carrying two such sensors can be done in a subtle and non-stigmatizing way. Further research should then go into an efficient communication and transmission of data in such a wireless body area network (WBAN). In practice, however, most multimodal systems contain sensors grouped in one device [55, 181]. This is a good implementation in terms of wearability, but artifacts often contemplate both modalities then. This creates a larger overlap in false alarm causes between both signals, leading to an increased FAR.

Different modalities could also be used to improve the quality of signals and features used for automated seizure detection. R peak detection in the ECG can for example be improved under noisy circumstances by using information from different modalities. In Appendix B, an algorithm is proposed that allows to more accurately detect heart beats in multimodal data. This algorithm can lead to an increased performance for ECG-based seizure detection due to the

more accurate heart rate extraction. It can however have a negative effect on the computation time, which is something that should be kept in mind in this real-time setting using wearable devices.

In general, the combination of different modalities can be further finetuned. Instead of using a binary unimodal outcome, continuous data such as the SVM latent variables could be used to allow more flexibility in the combination. The combination for example can then be a weighted average over the different sensors, giving more importance to more reliable sensors. The combination can also be done by a machine learning classifier, which automatically trains which unimodal inputs should have how much importance in defining the multimodal outcome. Another example of a potential improvement is a multimodal combination which takes into account missing data (e.g. disconnected sensor) or noisy data (e.g. turn of sensor input from a sensor under heavy noise). Similarly as in unimodal seizure detection, also multimodal seizure detection can strongly benefit from personalization. The ideal combination of sensors should be selected automatically per patient based on prior acquired data from previously recorded seizures for that patient. Not only the selection of sensors can be personalized, also the combination can be further optimized per patient.

8.6 Conclusion

A comparison of unimodal algorithms using sensors on different parts of the body was performed. Multimodal algorithms however typically lead to a decrease of at least 75% in false alarm rate, and sometimes decreases up to 90% are found. This illustrates the power of adding one extra sensor to an unimodal system. Multimodal systems are advised for automated warning systems for pediatric patients, in which the different sensors should be ideally from a different modality and placed at different locations of the body. Personalization of one unimodal algorithm can lead to a strong performance increase in a multimodal approach, which can lead to a similar performance increase as adding an extra sensor to the warning system.

Chapter 9

Conclusions and future directions

9.1 Conclusions

9.1.1 ECG-based patient-independent seizure detection

A new algorithm for automated patient-independent seizure detection using single-lead ECG was proposed in Chapter 3. A new preprocessing procedure was applied, and the extracted features were classified using a support vector machine classifier. These were both new aspects for the real-time detection of seizures using ECG, which resulted in a strong performance increase for the detection of focal impaired awareness seizures in temporal lobe seizures compared to the literature. The proposed patient-independent algorithm resulted in a 50% decrease in false alarm rate compared to the literature.

Despite the strong improvement of the results, the proposed method still leads to a too large false alarm rate (2 FP/h) to be used in practice. Therefore, unimodal algorithms using ECG are not likely to be currently used in an unimodal warning system. It could however be used as part of a closed-loop VNS system [19]. In these systems, the automated seizure detection causes an activation of the vagus nerve, which could potentially limit the effect of the seizure or even shorten the seizure duration [84]. In this application, having many false alarms is less critical than in the warning system application, as a false alarm has a limited effect on the patient. The proposed method can

enhance existing algorithms in currently available closed-loop VNS devices as it limits the FAR, while maintaining a nearly equally fast and accurate detection of the seizures.

9.1.2 Personalization of ECG-based seizure detection

Different personalization approaches were proposed in the thesis to further improve the seizure detection performance. These approaches are meant to adapt to the patient's heart rate characteristics with a limited amount of patient-specific (annotated or not-annotated) data. Four different personalization approaches were proposed, each with their specific strengths and target usage in practice.

Chapter 4 proposed a transfer learning technique which allowed to personalize ECG-based seizure detection using a limited amount of annotated patient data. It can be efficiently trained offline after collecting some days of patient-specific data. This approach can for example be used if accurate annotations are available, for example after a prechirurgical evaluation in the hospital. The annotations can then be made by the clinician. The results showed that supervised transfer learning is indeed an ideal solution for a fast and robust offline personalization of seizure detection algorithms.

It might however be that not enough seizure data is gathered during such a prechirurgical evaluation for a robust personalization. In that case, a different supervised approach was proposed in Chapter 5, which allows to use data collected in the patient's home environment. It adapts in real-time to the patient's heart rate characteristics by using patient feedback on previously generated alarms. An extra procedure was implemented in order to compensate for potential incorrect patient annotations, which was based on the transfer learning principles discussed in Chapter 4. This method showed to be robust against incorrect patient feedback, which is an unavoidable issue that needs to be solved when this approach would be applied in practice.

Unsupervised personalization approaches are mainly of added value for nocturnal monitoring as during nocturnal recordings, epileptic activity can be seen more clearly as outliers in the data distribution. A novel heuristic real-time adaptation procedure was proposed in Chapter 6, which only requires on average a couple of hours to adapt to each patient. A newly proposed unsupervised transfer learning approach was tested for the offline personalization of ECG-based seizure detection in Chapter 7. It also led to an increased overall performance, but also requires more patient-specific data than the method proposed in Chapter 6 to adapt to a patient. Therefore, the heuristic approach might be ideal for first nights/weeks of monitoring a specific patient. Afterwards, the transfer learning

approach might lead to a better personalization due to the increased complexity of the algorithm. Ideally, both algorithms are combined, preferring the first during the first weeks of monitoring, after which the latter can be applied for further personalized optimization.

The unsupervised methods are mainly important in practice, as they require no added input/effort from any user (developer, patient or clinician) since the personalization is accomplished completely automatically within the algorithm. They thus lead to an improved personalized performance without extra cost (except for a limited amount of extra computational requirements) compared to the patient-independent algorithms currently used in real-life devices. The unsupervised approaches are however currently limited to usage during nocturnal recordings. Supervised approaches still lead to stronger performance increases overall, but require the occurrence of at least a couple of seizures in the specific patient. This can mainly be an issue in patients with a low seizure frequency, which will still have to wait a large time before their algorithm can be sufficiently personalized. This requirement for seizures is however not a problem for the proposed unsupervised approaches.

In all proposed methods, only a limited amount of ECG-based features were used. This is done in order to be able to get a robust SVM classifier with a limited amount of (seizure) data. Typically, the more features that are used for training a classifier, the more data is required in order to train a robust classifier. This is even more important for the proposed transfer learning approaches, for which the strength is that it is able to adapt to a patient with a limited amount of data. However, if more features would be added, the performance (both patient-independent and personalized) might increase, but also the time for personalization would strongly increase. It needs to be further analyzed in the future whether adding more features indeed has added value in terms of personalization efficiency, and what the effect is on the time to personalization.

These personalized approaches showed to lead to decreases in FAR between 30% and 57% depending on methodology, seizure type and patient group. Nevertheless, they on average still lead to too many false alarms for unimodal usage in a warning system. The proposed approaches try to find heart rate increases in order to detect seizures. In some seizures, however, the ANS is not altered, and ECG-based seizure detection will not be able to detect these seizures under any circumstances. Also very short seizures are very difficult to detect using ECG (similar as in most modalities), as only sufficiently long ANS activations result in heart rate increases that are differentiable from non-epileptic heart rate increases. Another important factor remains the occurrence of motion artifacts, which cause a large fraction of the found false alarms in ECG-based seizure detection. Improved artifact removal techniques and hardware are required in order to further limit the amount of false alarms in

unimodal ECG-based seizure detection in the future.

ECG-based seizure detection thus remains very challenging for unimodal seizure detection, even with personalization approaches. Often, it is very difficult to differentiate epileptic and non-epileptic heart rate behavior efficiently, which causes a mediocre performance in unimodal systems. It could be that if more patient-specific data is obtained, the performance of the algorithms could be further improved in the future. The algorithms can then take into account more accurately the extra different factors influencing the heart rate dynamics, such as stress, sleep and the circadian rhythm.

Although the ECG-based seizure detection algorithms are less differentiable than other algorithms such as ACM and EMG-based algorithms, they are more widely applicable than these algorithms. Heart rate increases are found in around 70% of seizures, which allow them to be detected by ECG-based algorithms. Because the accuracy of unimodal ECG-based algorithms is too low, they are ideal to be used in a multimodal combination as they are active in most seizure types. Another reason why ECG would be preferred in a monitoring system is the great variability in possible extra uses for which it can be used. By including ECG in a monitoring system, also other health information from the patient can be gathered, such as stress, sleep and potential dangerous heart conditions. These can greatly improve a seizure warning system into a more general monitoring application for epileptic patients.

The different results demonstrated that specially children could benefit from the used personalization approaches. More inter-patient variability can be found within children compared to adults. This causes patient-independent algorithms to also perform worse in pediatric patients compared to adults, which increases the strength of personalization approaches in those patients.

The different proposed algorithms for personalization can easily be reused for personalizing/adapting other types of algorithms from different research topics. The newly proposed algorithms for real-time adaptable classifiers can for example also be used for real-time monitoring of patients in the ICU. For these patients, every moment counts, so an algorithm which quickly adapts to a patient's heart rate characteristics can be of great added value in generating alarms with increased accuracy. It can for example also be used for automated recognition of physical activities in patients with chronic movement disorders [17].

Also the offline personalization using the newly proposed unsupervised and supervised transfer learning approaches can be easily applied to other research topics in which typically only a limited amount of data is available. Automated sleep analysis and apnea detection are typical examples of such applications

which could benefit from efficient personalization using the proposed approaches [196]. The applied techniques can also be used for adapting to other situations rather than different individuals. The methods can for example can also be applied to adapt noise detection algorithms in order to be reused for applying the algorithm on new sensors.

9.1.3 Multimodal seizure detection

Different unimodal algorithms using ACM, EMG and ECG were compared for the automated detection of nocturnal generalized tonic-clonic seizures. From this comparison, placing an ACM sensor on of the wrist is assumed to be the best choice in terms of performance and wearability. It can be easily concealed in a bracelet (or already available in a smartwatch), which is the most subtle and non-stigmatizing solution compared to the investigated alternatives.

However, for multimodal combinations, sensors from different modalities should preferably be placed at different locations of the body. This is to avoid that the origins of false alarms in one sensor also do not cause false alarms in other sensors (e.g. rhythmic clapping might cause false alarms in both wrist sensors). In the proposed multimodal setup, both sensors are required to be active, but the added value of this approach is diminished if too much correlation between both sensors is found. Therefore, ideally the sensors are placed at different locations of the body. This however creates a less comfortable situation for the patient as it requires multiple sensors on the body. If the different sensors are however very subtle, it can be acceptable (e.g. two subtle wristband-like sensors). More work however needs to be done in order to improve the hardware to obtain better biomedical signals with smaller devices. Also the wireless connection of sensors should be further optimized (see Section 9.2).

The newly proposed combination method which takes into account the time difference in activations in different modalities led to a decrease of at least 75% in FAR in all combinations compared to the unimodal algorithms, and to at least 90% FAR reduction in 57% of the multimodal combinations. The proposed AND-operator method which takes into account the potential time differences in epileptic activity showed to work well on this dataset, but further evaluation needs to be done on larger datasets. For focal to bilateral generalized tonic-clonic seizures, the used time delay values δ might be insufficient for some multimodal combinations. If the time before the generalization of the seizure is too long, the differences in time between epileptic activity in the ECG and ACM/EMG might be larger than the currently used values. Broader ranges of δ will then be required, potentially also increasing the FAR of the multimodal system.

Personalization of one sensor could also lead to a great performance increase in FAR in a multimodal algorithm. It was found that the combination of 3 patient-independent algorithms led to a similar performance as the combination of 1 patient-independent and 1 personalized algorithm. This illustrates that personalization should be used, also in multimodal systems. In practice, three of the four proposed personalization approaches does not require any added effort from either patient or clinician compared to currently used monitoring devices, meaning that these personalization approaches can be easily implemented in current solutions for improved performance. The proposed personalization approaches are also applicable to other modalities such as ACM and EMG, which have shown to also possibly benefit from personalized solutions.

9.2 Future directions

The proposed ECG-based algorithms showed to be insufficiently accurate for unimodal warning systems in most cases. Potentially, seizure detection could be improved by including more data from the state of the patient. Algorithms that monitor for example sleep or stress could allow more accurate detection per patient [37, 86, 196]. Also more accurate monitoring of the circadian rhythm could lead to improved performance [105, 123].

Inclusion of meta-data or context data could also lead to increased seizure detection performance, but is currently not well investigated. Some patients for example have most of their seizures in the early morning, so taking into account the time could also lead to improved personalized performance in some patients. The location of the patient might be interesting to either make the algorithm more flexible (i.e. allowing more false alarms to increase sensitivity, for example in dangerous locations such as kitchen and staircase) or to make it stricter (i.e. allowing less false alarms and potentially decreased sensitivity, for example during physical exercise).

In the literature, deep learning also showed to give increased performance compared to classical machine learning solutions. However, a very large amount of annotated data is required in order to accurately train such a deep learning neural network. Typically, not enough data is available to train such a network for patient-independent automated seizure detection, let alone for personalized approaches. In general, large public datasets for seizure detection algorithms are missing for this type of research. Evaluation in the literature is typically done by using private datasets of limited sizes, which leads to a large amount of algorithms that are not reproducible on new datasets. A large available

public dataset would take away this concern, and could potentially trigger more researchers to explore this topic in the future.

All studies were currently performed in a controlled hospital or revalidation center environment. The algorithms should be tested in a real-life home setting to further investigate their accuracy and practical usability. These approaches should also be implemented for efficiently personalizing seizure detection using other modalities such as accelerometry and electromyography.

In the literature, most multimodal combinations are performed with an approach of a relative limited complexity. Typically, it is restricted using a rule-based approach such as an AND-operator or OR-operator on the binary unimodal outputs. More complex multisensor combinations could be used for improved performance. The combination can for example be done in an automated (and potentially personalized) way, which might give more importance to certain sensors over others in a sort of combinatory machine learning classifier. The combination can also be further improved by not simply combining the unimodal outputs in a fixed way, but rather dynamically, constantly adapting to the current status of the patient/system (e.g. temporary turn off the impact of too noisy sensors).

It was found that multimodal algorithms with sensors placed on different parts of the body lead to an increased performance. However, currently no research went into how the combinations of such wireless body area network (WBAN) of sensors could be used efficiently in practice for epileptic seizure detection. For such a network, it is for example important to keep the battery lifetime intact, certainly if the different devices have to be of limited size. Different options are possible, for example by choosing one master sensor which is constantly active. If potential epileptic activity is detected in the master sensor, it can activate the other sensors ("slaves") to improve the accuracy of the seizure detection. If the potential epileptic activity is over, the other sensors can be turned off again (see Figure 9.1). This leads to a strong reduction in battery lifetime of the slaves, but requires a robust master sensor which is able to quickly detect potential epileptic activity. The battery life-time can also be strongly increased by using context-aware sensors, which automatically decide when a signal needs to be sampled at maximal frequency and accuracy. Also distributed signal processing techniques can be applied, which reduces communication to a minimum by performing as many computations on the sensor device itself [14].

Only the multimodal detection of generalized tonic-clonic seizures was investigated in this thesis. Also the detection of other seizures can be of interest for automated warning systems. Specially the detection of subtle focal seizures is currently not well addressed in the literature, for which the different proposed ECG-based seizure detection algorithms can be of great added value.

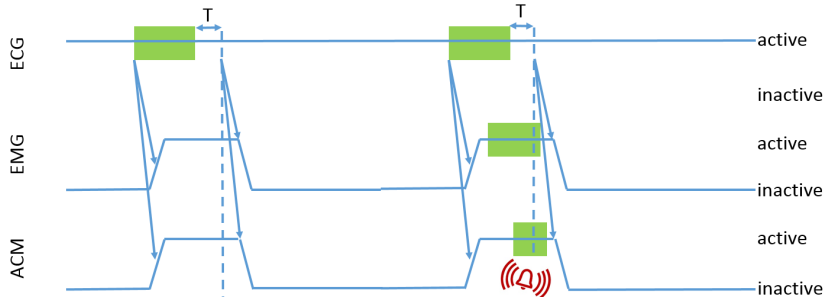


Figure 9.1: Illustration of a multimodal master-slave setup. In this example, the ECG sensor is the master sensor, and the ACM and EMG sensors are the so-called slaves. The ACM and EMG sensors are only active when they are awoken by the master sensor after epileptic activity (marked in green) is detected. After a period T after the last epileptic activity in the master, the slaves are turned off again. A multimodal alarm is then generated if all three sensors contained epileptic activity (see second master activation).

Which sensors can be used for detection, will depend on the patients and seizure type. The optimal selection of sensors could be chosen automatically. The patient could initially wear a larger number of sensors from different modalities all over his/her body for a limited amount of time. Afterwards, an algorithm can decide which combination of sensors would lead to an optimal performance for a minimal amount of sensors (see Figure 9.2). This way, one automated procedure could be used for all patients and seizure types, automatically choosing the optimal setting for each individual patient.

Initial steps towards automated monitoring applications had already been taken in the literature, which mainly worked out straight-forward generic detection algorithms on a limited set of seizure types. The work in this thesis illustrates that more advanced procedures such as improved multimodal combinations and efficient personalization can strongly improve the accuracy of detection algorithms for a wider range of seizure types. These algorithms can pave the way towards practically usable monitoring systems to improve the quality of life of a large variety of epilepsy patients and their relatives.

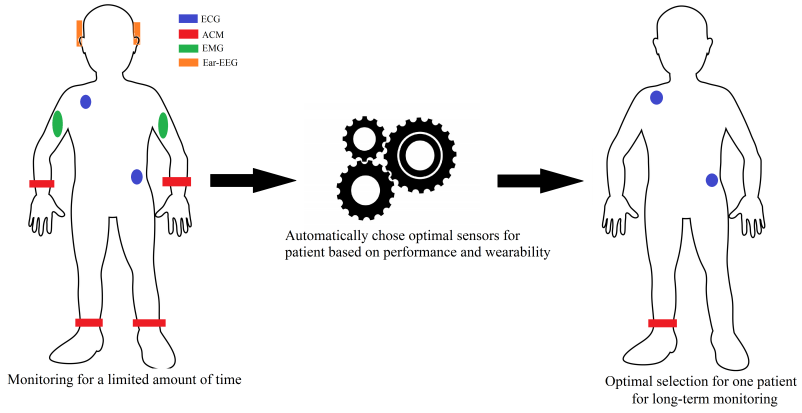


Figure 9.2: Illustration of automated selection of sensors for a multimodal seizure detection system.

Appendix A

Features extracted for automated seizure detection

This appendix gives an overview of the different features extracted for automated seizure detection throughout the thesis. Section A.1 gives an overview of all extracted features for the ECG-based seizure detection approach discussed in Chapter 3. Sections A.2 and A.3 discuss the features used for EMG and ACM-based generalized tonic-clonic seizure detection used in Chapter 8.

A.1 ECG-based seizure detection features extracted during *HRI-EXTRACT*

This section gives an overview of the different features that are extracted from a detected heart rate increase (HRI) according to Algorithm 3.1 mentioned in Section 3.3.2. The time of the start of the HRI is defined as t_{start} and the end of the HRI as t_{end}.

A.1.1 Features extracted during the HRI

Features related to the shape of the HRI

The following features indicate specific features which give information on the shape, duration, etc. of the detected HRI. The mentioned features are extracted

using the filtered heart rate signal.

1. Duration of the HRI in seconds ($\Delta t_{HRI} = t_{end} - t_{start}$).
2. Duration of the HRI in number of detected heart beats.
3. Maximal heart rate gradient (slope) ∇HR_{max} found during the HRI. The heart rate slope is defined as the slope of a linear fit on 10 consecutive heart rate values from the filtered heart rate signal.
4. The R-square value of the above-mentioned fit for the found maximal slope, which gives an indication of the goodness of fit of this linear fit, computed as
$$R^2 = 1 - \frac{\sum_i^{HRI} w_i(\hat{y}_i - \bar{y})^2}{\sum_{i \in HRI} w_i(y_i - \bar{y})^2} \tag{A.1}$$
with w_i the found weights from the fit, $\hat{y}_i$ the found estimated values with the linear fit, y the original values and $\bar{y}$ the mean value of the original values.
5. The maximal heart rate found during the HRI (called peak heart rate HR_{peak}, found at time t_{end})
6. The gain of heart rate in beats per minute (bpm) (ΔHR) during the HRI compared to the start of the HRI (HR_{start} at time t_{start}), and is defined as
$$\Delta HR = HR_{peak} - HR_{start} \tag{A.2}$$
7. The percentual gain of heart rate ($\Delta HR_{\%}$) during the HRI compared to the start of the HRI, defined as
$$\Delta HR_{\%} = \frac{HR_{peak} - HR_{start}}{HR_{start}} \tag{A.3}$$
8. The gain of heart rate in bpm during the HRI compared to the average heart rate in the pre-HRI period ($\overline{HR}_{pre}$)
$$\Delta HR_{pre} = HR_{peak} - \overline{HR}_{pre} \tag{A.4}$$
9. Average heart rate in rest (HR_{rest}). HR_{rest} stands for the estimated average HR at rest for the patient under testing and is computed on-the-run. HR_{rest} is set to $\overline{HR}_{pre}$ if
   - $\overline{HR}_{pre} < HR_{rest}$,
   - the linear fit on the data of HR_{pre} results in an absolute gradient $< 0.2\ bpm/s$ and

- the linear fit is found with $R^2 > 0.7$.

10. Absolute gain in heart rate during the HRI, but with reference to HR_{rest} instead of HR_{start}

$$\Delta HR_{rest} = HR_{peak} - HR_{rest} \tag{A.5}$$

11. Percentual gain in heart rate during the HRI, but with reference to HR_{rest} instead of HR_{start}

$$\Delta HR_{\%rest} = \frac{HR_{peak} - HR_{rest}}{HR_{rest}} \tag{A.6}$$

12. Percentage of differences in RR-intervals that deflect more than 20% from its successive RR-interval. This feature gives an indication of the amount of potential R-peak detection errors found during the HRI.

Heart rate variability during the HRI

The following features include heart rate variability (HRV) features, which can be divided in time domain features and frequency domain features. All HRV features are computed based on the unfiltered heart rate signal during the HRI as obtained by the R peak detection algorithm.

Time domain features:

13. Mean heart rate during the HRI.

14. Standard deviation of the NN interval (SDNN), with NN being the normal-normal beats (so similar to RR-intervals but with removal of abnormal beats), defined as

$$SDNN_{HRI} = \sqrt{\frac{1}{N-1} \sum_{i=1|i\in hri}^{N} \left(RR_i - \overline{RR}\right)^2} \tag{A.7}$$

15. Root mean square of the successive differences (RMSSD) of the RR-intervals in the HRI

$$RMSSD_{HRI} = \sqrt{\frac{1}{N-1} \left(\sum_{i=1|i\in HRI}^{N-1} \left(RR_{i+1} - RR_i\right)^2 \right)} \tag{A.8}$$

16. Standard deviation of successive differences (SDSD) of the RR-intervals in the HRI

$$SDSD_{HRI} = \sqrt{\frac{\sum_{i=1|i\in HRI}^{N-1} \left(\Delta RR_i - \overline{\Delta RR}\right)^2}{N}} \tag{A.9}$$

with $\Delta RR_i = RR_{i+1} - RR_i$.

Frequency domain features:

17. The absolute power in the very low frequency (VLF) domain (0-0.04 Hz), low frequency (LF) domain (0.04-0.15 Hz) and high frequency (HF) domain (0.15-0.4 Hz) during the HRI [23].
18. Total power over all frequency bands during the HRI.
19. Relative power of the individual power bands compared to the total power during the HRI.
20. The LF/HF ratio during the HRI [99].
21. Spectral entropy during the HRI, defined as

$$\text{spectral entropy} = \sum_{f=1}^{M} PSD_n(f) log\left(PSD_n(f)\right) \tag{A.10}$$

with $PSD_n(f)$ the normalised power spectral density (PSD) at a certain frequency f.

A.1.2 Features extracted from pre-HRI period

22. Mean heart rate during the pre-HRI period ($\overline{HR}_{pre}$).
23. SDNN during the pre-HRI period

$$SDNN_{pre-HRI} = \sqrt{\frac{1}{N-1} \sum_{i=1|i\in pre-hri}^{N} \left(RR_i - \overline{RR}\right)^2} \tag{A.11}$$

24. RMSSD of the RR-intervals in the pre-HRI period

$$RMSSD_{pre-HRI} = \sqrt{\frac{1}{N-1} \left(\sum_{i=1|i\in pre-HRI}^{N-1} \left(RR_{i+1} - RR_i\right)^2 \right)} \tag{A.12}$$

25. SDSD of the RR-intervals in the pre-HRI period

$$SDSD_{pre-HRI} = \sqrt{\frac{\sum_{i=1|i\in pre-HRI}^{N-1} \left(\Delta RR_i - \overline{\Delta RR}\right)^2}{N}} \tag{A.13}$$

with $\Delta RR_i = RR_{i+1} - RR_i$.

26. The absolute power in the VLF domain, LF domain and HF domain during the pre-HRI period.
27. Total power over all frequency bands during the pre-HRI period.
28. Relative power of the individual power bands compared to the total power during the pre-HRI period.
29. The LF/HF ratio during the pre-HRI period.
30. Spectral entropy during the pre-HRI period.

Also some of the Hjorth parameters [90] are applied here on the RR-sequence of the pre-HRI period:

31. Hjorth's activity (similar as normal variance):

$$\text{activity}(RR) = \frac{1}{N} \sum_{i=1|i\in pre-hri}^{N} \left(RR_i - \overline{RR}\right)^2 \tag{A.14}$$

with $\bar{RR}$ the mean RR during the pre-HRI.

32. Hjorth's mobility

$$\text{mobility}(RR) = \sqrt{\frac{\text{activity}\left(\Delta RR\right)}{\text{activity}(RR)}} \tag{A.15}$$

33. Hjorth's complexity

$$\text{complexity}(RR) = \frac{\text{mobility}\left(\Delta RR\right)}{\text{mobility}(RR)} \tag{A.16}$$

A.2 EMG-based seizure detection features

This section gives a detailed description of the features used for patient-independent EMG-based generalized tonic-clonic seizure detection as described in Section 8.3.1.

1. Zero-crossing rate of the second derivative signal: The zero-crossing rate indicates the number of times the (in this case second derivative) signal changes signs within the window.

2. Wilson amplitude for an EMG signal $emg(n)$ is defined here as

$$\begin{aligned} &\text{Wilson amplitude} = \sum_{n=2}^{N} \phi_n \\ &\text{with} \quad \phi_n = \begin{cases} 1 & ,|emg(n) - emg(n-1)| > 50\ \mu V \\ 0 & ,otherwise \end{cases} \end{aligned} \tag{A.17}$$

A.3 ACM-based seizure detection features

This section gives a detailed description of the features used for patient-independent accelerometry-based generalized tonic-clonic seizure detection as described in Section 8.3.2. Feature extraction is done per channel within a sensor.

1. The number of local maxima in the ACM channel: estimated by the number of zero crossings of the first derivative of the ACM signal.

2. Frequency of the first local peak in the frequency spectrum: The power spectral density is acquired by using the Welch method.

3. Shannon entropy of diagonal line lengths in a cross-recurrence plot [52, 150]: A recurrence plot shows the time instances for when the phase space trajectory of a dynamical system visits nearly the same area in the phase space. A recurrence is a time the trajectory returns to a location it has visited before. It can be mathematically formulated as

$$\mathbf{R}_{i,j} = \begin{cases} 1 & \text{if} ||x_i - x_j|| \leq \epsilon \\ 0 & \text{otherwise} \end{cases} \tag{A.18}$$

with $1 \leq i, j \leq N$ with N the number of considered states and ϵ a threshold value. The computation of the recurrence plot is done by using the CRP toolbox [121]. Here, the Shannon entropy of diagonal line lengths in the recurrence plot are used, which is defined as

$$E_{rp} = -\sum_{l=l_{min}}^{N} p(l) \ln p(l) \tag{A.19}$$

with $p(l)$ the probability that a diagonal line in the recurrence plot $\mathbf{R}$ has exactly length l.

4. Power within pseudofrequency range 9-22.5 Hz from a continuous wavelet transform: The continuous wavelet transform (CWT) of a continuous signal x(t) is defined as

$$CWT_{\psi}(a,b) = \frac{1}{\sqrt{a}} \int_{-\infty}^{+\infty} x(t)\psi^{*}\left(\frac{t-b}{a}\right) dt \tag{A.20}$$

with ψ^{*} the complex conjugate of the mother wavelet. A Morlet wavelet is used in this thesis as mother wavelet. The power in the frequency band of interest δ_a can then be found by

$$W_{\delta_a} = \frac{\sum_{a\in\delta_a}\sum_b CWT(a,b)}{\sum_a\sum_b CWT(a,b)}. \tag{A.21}$$

Appendix B

Heart beat detection in multimodal data

This appendix proposes two algorithms for robust heart beat detection in multimodal data, which was submitted as entry for the 2014 Computing in Cardiology challenge [163]. The goal of this study is to robustly extract the heart rate from multimodal data if the reference ECG signal is contaminated with noise. The first proposed algorithm combines heart beat annotations from all signals of interest (e.g. ECG signals, blood pressure signals, PPG signals, etc. as labeled in the dataset) using majority voting, heart beat location estimation and Hjorth's mobility of the resulting RR intervals. These signals of interest can be detected automatically by analyzing their power spectral density or by using the available signal type identifiers, which is applied in the second method. This way, signals with incorrect signal labels or signals with strong ECG noise can still be used to improve the performance of the multimodal heart beat detection.

These algorithms could also be used in a multimodal setting for automated seizure detection, as noisy ECG signals showed to be an important reason for low seizure detection performance in the ECG [44]. This study is also described in detail in [40].

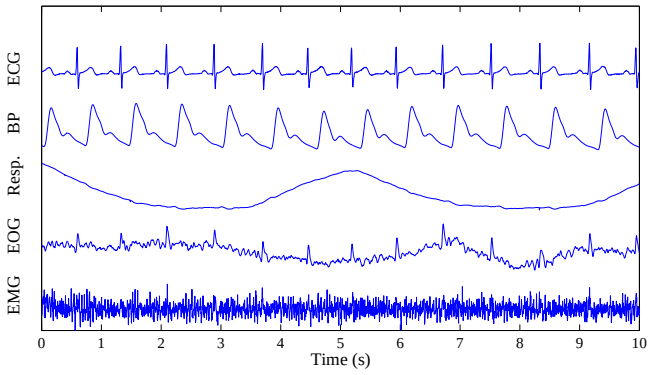


Figure B.1: Example of a multimodal data segment. In this example the BP signal and the EOG signal with R peak interference can improve the performance of the heart beat detection in case the ECG signal would have been noisy.

B.1 Introduction

Continuous extraction of the heart rate and heart rate variability is crucial in multiple monitoring applications. Typically only the available ECG signal is used for its computation. It can however occur that this signal becomes too noisy or the electrodes can even become disconnected, possibly leading to false or missing heart rate information, which could have crucial effects on the monitoring efficiency.

In several monitoring applications other biomedical signals are also acquired simultaneously. Pulsatile signals such as blood pressure (BP) or signals with strong R peak interference could be used for improved heart beat detection (see Figure B.1). Motion artifacts in the ECG can for example also be removed if the corresponding noise is also acquired in another signal in de dataset (i.e. accelerometer) by applying an adaptive filter [153], which can also increase the heart beat detection performance. The goal is to combine the heart beat information from all these signals in order to make a single decision about the R peak locations. This was also the subject of the 2014 Physionet/Computing in Cardiology challenge (from now on referred to as the 'challenge') to which the proposed methods were submitted [163].

Signal type identifiers (*ECG*, *BP*, etc.) were available in the used datasets (see Section B.2), but this study proposes a method that is able to automatically detect these signals that can improve the heart beat detection. In this way the algorithm is able to detect signals with R peak interference (e.g. EEG, EOG)

and deal with missing or erroneous signal type identifiers (IDs). This procedure and the methods for multimodal heart beat detection are described in Section B.3. Finally, the results are given and discussed in Section B.4 and Section B.5. A precursor of this study is published in [39].

B.2 Datasets and challenge information

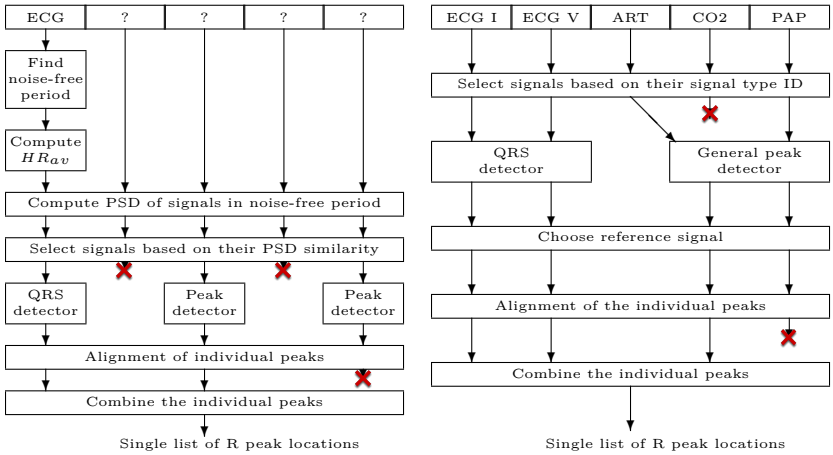
A detailed overview of the detailed challenge information, restrictions and data can be found in [163]. In short, the dataset contains signals of at most 10 minutes. The signals are multi-parameter recordings of adults, including both healthy volunteers and patients with a wide range of health problems. Each recording contains between four and eight signals, of which the first is always an ECG signal. The other signals could be any variety of simultaneously recorded physiologic signals. Sampling frequencies range between 120 and 1000 Hz. Example of acquired signals include amongst others extra ECG channels, EEG, blood pressure, accelerometry and respiration. Signal labels are available, but could contain incorrect or inconsistent information, which should be kept in mind when constructing the algorithm.

In this study, the mentioned training set contains the 200 recordings that were available after the challenge, and the test set contains the 200 other remaining hidden recordings. It should however be noted that the discussed methods were initially developed with the use of only the initial challenge training set (containing only 100 recordings).

Because the initial available training set did not contain sufficient challenging examples (typical performance $> 99.85\%$), an extra dataset of 100 recordings was extracted from the MGH/MF Waveform database [202]. These were extracted as the recordings between the 20th and 30th minute of recordings 001 to 107. Recordings 041, 043, 045, 047-049 and 083 were left out due to errors during reading of the original data/annotations. The time intervals were chosen to guarantee that correct reference annotations were available. These recordings typically contained additional ECG, pulmonary arterial pressure (PAP), arterial BP (ART) and central venous pressure (CVP) signals and can thus optimally profit from the methods discussed in this study.

B.3 Methodology

Two algorithms for multimodal heart beat detection are elaborated in this study (see Figure B.2). At first, a multimodal algorithm MM_{NI} that does not



(a) MM_{NI} algorithm (b) MM_{WI} algorithm

Figure B.2: Overview of the algorithms proposed in this study. In the MM_{NI} algorithm (a), only the signal type of the first signal is certainly known (ECG). In the MM_{WI} algorithm (b), all available signal IDs are used. Crosses indicate after which steps signals can be removed from further processing.

require the availability of signal type IDs (except for the known reference ECG signal) is discussed in Section B.3.1 to Section B.3.3. Section B.3.4 describes an additional algorithm MM_{WI} that does use these signal IDs.

B.3.1 Automatic relevant signal detection using power spectral density similarity

At first, the potential signals of interest are selected automatically. These signals typically have the same periodic behavior as the ECG, which will be exploited in this procedure. With this procedure, there is no need for using the given (possibly false or missing) signal IDs, making the algorithm more robust when applying it on other databases. Pulsatile signals are examples of signals that need to be detected, but also other signals which contain R peak interference can improve the heart beat detection.

First, a noise-free segment of the ECG is selected based on a method described in [199] as follows. The ECG is divided into k segments $C_i, 1 \leq i \leq k$, of 5 seconds and each segment is filtered using a bandpass Butterworth filter with cut-off frequencies at 1 Hz and 40 Hz. Next, the autocorrelation function (ACF)

of each segment is computed. A similarity matrix $\mathbf{S}$ is then constructed with entries defined as

$$\mathbf{S}_{i,j} = \frac{A_i^T A_j}{\|A_i\|\|A_j\|}, \qquad i,j = 1,2,...,k \tag{B.1}$$

with A_i the ACF of segment C_i. Now let

$$M_i = \sum_{j=1}^{k} \mathbf{S}_{i,j}, \qquad i = 1,2,...,k \tag{B.2}$$

so that M_i gives an indication about the similarity between segment C_i compared to all other segments in the signal. We assume that the signal contains more noise-free than noisy segments, so a high similarity in a segment would indicate a noise-free segment. A period of l subsequent segments is then selected with maximal M_i values:

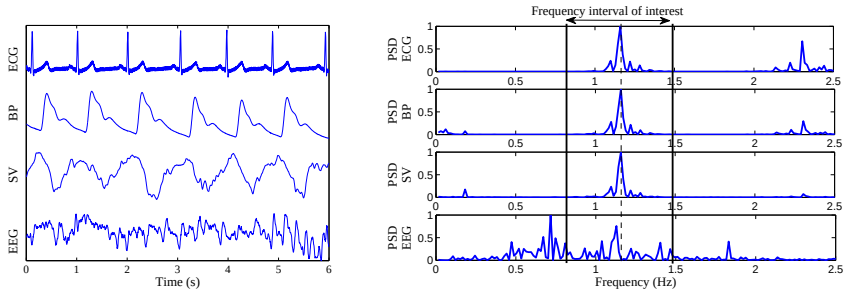
$$a = \arg\max_i \sum_{n=0}^{l-1} M_{i-n}, \tag{B.3}$$

with l chosen as the number of segments corresponding to 10% of the signal length (typically 1 minute). The resulting block of segments $C_{a-l+1,a-l+2,...,a}$ then forms the noise-free ECG period used in further processing.

The R peak detection algorithm from [208] is then applied to this ECG segment and the average heart rate HR_{av} of this segment is obtained. Next, the normalized power spectral density (PSD) is computed for all signals in this time segment. Signals with the same periodic behavior as the ECG should contain a peak in their PSD around the average heart rate frequency HR_{av} (see Figure B.3). In order to get a more distinctive peak around HR_{av} in the ECG PSD, flattened ECG signals are used as input for the PSD computation. These flattened signals were found by subtracting the upper and lower secant envelopes from each other [196]. The similarity (computed similarly as in (B.1)) between the PSD values in the interval [HR_{av}-0.3 Hz, HR_{av}+0.3 Hz] of the ECG signal and the other signals must exceed the threshold of 0.75 in order to be detected as a signal of interest. The value of this threshold is optimized heuristically using the available training set. Only these signals will be processed from now on. The procedure for detection of the noise-free ECG period was necessary in order to get a reliable estimate of the HR_{av} value.

B.3.2 Individual peak detection

In the next phase, heart beat locations are found for each signal of interest for the entire record length. For the first signal, a different QRS detection algorithm



(a) Original signals.

(b) PSD of signals.

Figure B.3: Illustration of the automatic signal detection procedure. The left figure shows four signals during a 6 seconds epoch of the selected noise-free ECG period. The PSD values of all signals are shown in the right figure. The dotted line in the right figure indicates the average heart rate frequency HR_{av} in the selected noise-free period in the ECG signal. The frequency segment that is used for computing the correlation coefficient is marked by the solid vertical lines. In this example only the BP and stroke volume (SV) signal are selected as signals of interest.

(*gqrs* from the WFDB toolbox) is used here compared to the one discussed in Section B.3.1 [164]. Although the latter resulted in a better performance, it was computationally too expensive to run on the entire signal length. Therefore it was only executed on the noise-free ECG period to get a more reliable estimate of HR_{av}.

For the other signals, a basic general peak detection algorithm is used. First, an exponentially decreasing envelope e_{sig} is computed from each signal sig

$$e_{sig}(n) = \max_i sig(n - w + i)exp(-if_s^{-1}), \qquad i = 1, 2, ..., w \tag{B.4}$$

with f_s the sampling frequency of the signal and w the used window size, set here to $\lceil f_s/2 \rceil$. Heart beat locations are then detected if their corresponding envelope value is the highest value in a neighborhood of 300 ms around its location. Despite its simplicity, the algorithm performed well on noise-free BP, stroke volume (SV) and ECG signals with positive peaks at a low computational load. By detecting the peaks of the signal envelopes, this method is more robust when finding peaks in additional ECG signals as the T peaks are less likely to be detected this way.

Because the peaks in the different types of signals do not occur at the same time as the ECG R peaks, the acquired peak positions need to be aligned with respect to the R peak locations before combining them in the next phase. The average

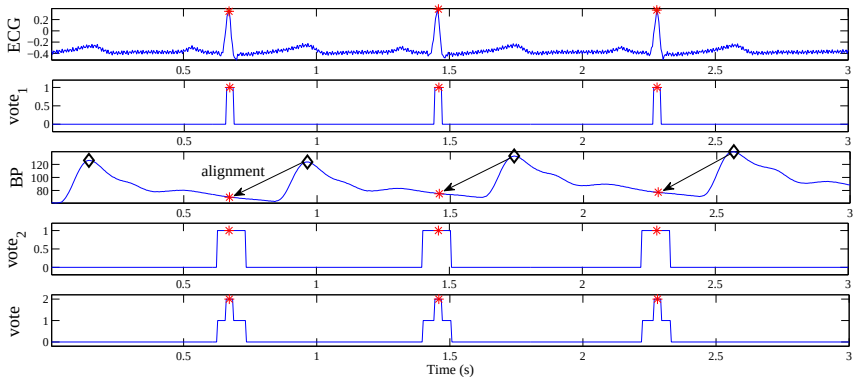


Figure B.4: Illustration of the majority voting procedure for combining heart beats from different signals. Signals $vote_1$ and $vote_2$ indicate temporary versions of *vote* for each signal separately. The peaks from the BP signal are indicated by diamond symbols, the arrows indicate the peaks for this signal after alignment. Signal *vote* is defined as the sum of $vote_1$ and $vote_2$. The real R peak annotations are marked by an asterisk.

delay between the peaks of the ECG and the other signals in the noise-free ECG period are computed and used for aligning each record separately. Derivation of the average delay is done by taking the most occurring delay between each signal and the reference ECG signal by computing the histogram of these delays. In this way, the algorithm can adapt to different kinds of signals of interest in a general way. If however the standard deviation of these delay values is too high (> 150 ms), the corresponding signal is removed from further analysis. This may occur if a signal was falsely selected or if the corresponding peak locations are too inaccurate to increase the heart beat detection performance.

B.3.3 Combination of individual peak locations

Finally, the different individual peak locations from all signals need to be combined into a single list of R peak locations. This procedure runs in two phases. In a first phase, an initial list of R peak locations is found by using majority voting on the individual peak locations. This procedure is implemented here by defining a signal *vote*, which has the same length as the signals and all values are initialized at zero. For each peak location *loc* in each signal of interest, the values $vote(loc - window : loc + window)$ are increased with 1. For R peak locations coming from the known reference ECG signal, *window*

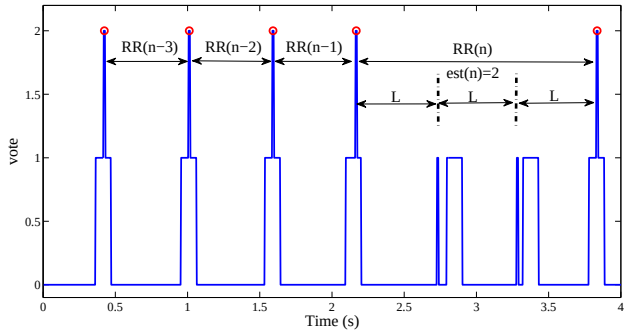


Figure B.5: Illustration of the combination of peak locations of different signals. The red circles indicate the initial set of R peak locations *peaks*. In this case, peak locations from ECG (narrow peaks) and BP (wide peaks) are combined. The chained lines show the estimated R peak locations $estloc_2^n$, which lay near the ECG peaks. The ECG peak locations are thus added to *peaks* in this example.

equals 10 ms, for other signals it is set to 50 ms. This is due to the fact that the pulsatile signals tend to contain less accurate peak locations and could thus vary more compared to the corresponding R peaks in the ECG signal. This procedure is illustrated in Figure B.4. An initial set of peak locations *peaks* is then found by selecting peak values in *vote* which are greater than or equal to $\lfloor \#signals/2 \rfloor + 1$, with $\#signals$ the number of signals of interest.

In the second phase, this initial set of heart beat locations *peaks* is reinspected in a chronological fashion in order to see if heart beats have been missed. Define $RR(n)$ as the interval between $peaks(n)$ and $peaks(n+1)$. An estimation of the number of missing R peaks $est(n)$ in $RR(n)$ is defined as

$$est(n) = \begin{cases} 0 & : val(n) \leq 1.5 \\ 1 & : 1.5 < val(n) < 2.6 \\ 2 & : 2.6 \leq val(n) < 3.5 \\ round(val(n)) - 1 & : else \end{cases}$$

with

$$val(n) = \frac{RR(n)}{median(RR(n-3:n-1))}$$

If $est(n) > 0$, further analysis of $RR(n)$ is required and peak locations in this interval need to be added. In this case, only peak locations originating from one signal can be added to *peaks*. The question remains from which signal

Algorithm B.1: Combining peak locations from different signals of interest

Input: loc_i with $1 \leq i \leq \#signals$ and $\#signals$ the number of signals of interest
Output: $peaks$ containing R peak locations
Construct $vote$ based on individual peak locations loc_i (see Figure B.4)
Find peak locations $peaks$ so that $vote(peaks) \geq \lfloor \#signals/2 \rfloor + 1$
for n=1:length($peaks$) **do**
    $RR(n) = [peaks(n), peaks(n+1)]$
    Estimate number of missing peaks $est(n)$ in $RR(n)$
    Define loc_i^n as the individual peak locations from signal i in $RR(n)$
    if $est(n) < 6$ **and** $\exists c \in [1, \#signals] : length(loc_c^n) = est(n)$
        $b = \arg\min_c \|estloc_{est(n)}^n - loc_c^n\|^2$
    else if $est(n) < 6$ **and** $\exists c \in [1, \#signals] : length(loc_c^n) = est(n) \pm 1$
        $b = \arg\min_c \|estloc_{est(n)\pm 1}^n - loc_c^n\|^2$
    else
        $b = \arg\min_i \mathrm{var}(\mathrm{diff}(loc_i^n))$, $i \in [1, \#signals]$
    end if
    add loc_b^n to $peaks$
end for

to choose these peak locations in $RR(n)$. Let loc_i^n be the individual peak locations in $RR(n)$ from signal i. If at least one list of peak locations loc_i^n contains exactly $est(n)$ heart beats, one searches for one loc_i^n that follows this constraint and has a minimal Euclidean distance to the estimated missing beat locations $estloc_{est(n)}^n$. $estloc_j^n$ is defined here as a list of peak locations in $RR(n)$ containing j peaks, with peaks located so that they result in equally large RR-intervals in $RR(n)$ (see Figure B.5). When none of the peak location lists loc_i^n contains exactly $est(n)$ peaks, this procedure is repeated for $est(n) \pm 1$ and $estloc_{est(n)\pm 1}^n$. If again no such list of beat locations is available or if $est(n) > 5$, the peak locations loc_i^n with the lowest variance in RR-intervals is chosen. Estimating the number of missing heart beats becomes too inaccurate if $est(n) > 5$ and is therefore avoided in this procedure. In these cases, loc_i^n are extended with $peaks(n-3:n-1)$ and $peaks(n)$ in order to check if these RR intervals fit in the surrounding RR values. When $peaks$ is fully analyzed, it contains the final list of heart beat locations. An overview of the combining procedure is shown in Algorithm B.1.

B.3.4 Heart beat detection using signal type identifiers

A second version of the algorithm (called MM_{WI}) was also submitted for the challenge. In this version, the available signal type identifiers are used

in order to define the potential signals of interest to be used for heart beat detection optimization. It is thus not able to detect signals which contain R peak interference, but is less likely to miss detection of signals with pulsatile signal IDs. The flow chart of this algorithm is shown in Figure B.2(b) and follows these steps:

1. First the signals of interest are detected by searching for signal IDs containing the string 'ECG', 'BP', 'PLETH', 'PAP', 'ART', 'PPG' or 'pressure'. These signals were chosen as they typically yield robust information about the heart beat locations and the general peak detector performed well on these signals. SV and CVP signals were left out as they too often had too much variance in their peak locations, most often leading to a decrease in performance.

2. R peaks in signals with ECG IDs were computed by means of the *gqrs* algorithm because of its limited computational load. Heart beats in the other signals were detected by the general peak detector described in Section B.3.2.

3. It can occur that the first signal is not trustworthy as a reference signal for peak alignment. This is typically the case if the first signal is too noisy over a big part of the length. Therefore it is possible to change the reference signal in this method. It is chosen as the signal that has the lowest Hjorth's mobility in its heart rate values [90]. This parameter gives a good indication about the smoothness of the resulting tachogram. Non-ECG reference signals are aligned towards the R peak locations by using standard delay values for it's signal type (250 ms for PAP, 200 ms for other pulsatile signals).

4. All peak locations are aligned to these reference peak locations by computing the average delays over the entire signal length (similar to in MM_{NI}). Signals with a too large standard deviation (> 150 ms) on these delays are removed from the list of signals of interest.

5. Different peaks are combined as discussed in Section B.3.3.

B.4 Results

B.4.1 Automatic signal detection

The correctness of the automatic detection of signals of interest was evaluated on the challenge training set and the mentioned subset of the MGH/MF dataset

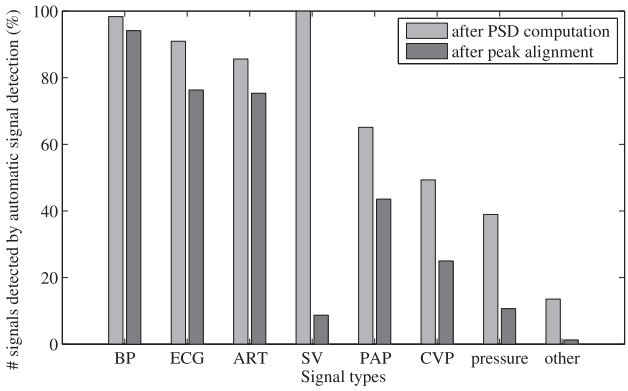


Figure B.6: Percentages of signals with given identifier that are detected by the proposed method for automatic signal detection. The first bars indicate the sensitivity percentages of each signal type after inspecting the PSD similarity (see Section B.3.1), the seconds bars indicate the percentages of used signals after removing the signals with too high variability in delay values (see Section B.3.2).

(see Figure B.6). It should be noted that some of these assumed signals of interest might actually be too noisy to use, so the actual sensitivity values of useful signals are a bit higher.

As indicated, the signals that were used for the final beat locations were chosen in two steps. In the first step the similarity of the PSD values was used in order to get a first list of signals of interest. The most important signals of interest (BP, ECG and ART) are correctly detected in at least 91% of the cases after this first step.

Additional signals of interest were removed if the found alignment delays between the signals had a too high variance. The loss in detection sensitivity is limited here for the most important types of signals. The sensitivity drop is however larger for the CVP and PAP signals, and almost none of the SV signals are remained for further processing after this step. After this procedure, 10 non-pulsatile signals remained for further processing. All of them were however signals with R peak interference, so no signals were falsely selected for the final heart beat computation.

Table B.1: Results (in %) for the different proposed methods for multimodal heart beat detection. Also the results for the other final official top 5 challenge entries and for the unimodal *gqrs* algorithm are shown. $MM_{NI} - limited$ and $MM_{WI} - limited$ indicate the limited versions of both algorithms which can only use two signals. The last column indicates the average performance over both discussed datasets.

	Challenge training set			MGH/MF subset			Overall
Method	Se_{gr}	$+P_{gr}$	Av	Se_{gr}	$+P_{gr}$	Av	Av
gqrs	94.11	91.00	92.55	83.25	94.32	88.78	91.30
MM_{NI}	94.60	92.36	93.48	90.76	96.71	93.73	93.56
MM_{WI}	**95.40**	93.32	**94.36**	97.11	97.75	97.43	**95.38**
$MM_{NI} - limited$	95.04	92.32	93.68	89.31	96.66	92.98	93.44
$MM_{WI} - limited$	93.99	93.09	93.54	93.14	96.83	94.99	94.02
Method [102]	93.91	92.48	93.20	92.37	92.15	92.26	92.88
Method [200]	93.54	**94.61**	94.07	**97.22**	**97.96**	**97.59**	95.25
Method [74]	94.22	93.29	93.75	78.77	95.80	87.29	91.60
Method [175]	93.09	92.10	92.60	82.89	94.67	88.78	91.32

B.4.2 Multimodal heart beat detection

The performance of the different heart beat detection algorithms is evaluated by using the sensitivity (Se) and positive predictive value (+P) of the detected R peaks. R peak positions were labeled as a true positive if they were within a distance of 150 ms of the correct reference annotations. Two variants of both measurements are used here. Se_{gr} indicates the sensitivity over the entire dataset, whereas Se_{av} indicates the average sensitivity over the recordings (similar for $+P_{gr}$ and $+P_{av}$). The overall performance (Ov) is defined as the average of these 4 measurements [163]. The average performance (Av) is the average of Se_{gr} and $+P_{gr}$.

Table B.1 shows the results from both mentioned algorithms on the training set and the subset of the MGH/MF database. Because the Physionet server had a restricted processing time for each entry, which typically only allowed peak detection in two signals, also limited versions of the algorithms were elaborated that use maximally two signals (called $MM_{NI} - limited$ and $MM_{WI} - limited$). These two selected signals are the reference ECG signal and the first detected signal of interest (either by using the PSD procedure in $MM_{NI} - limited$ or using the signal ID's in $MM_{WI} - limited$).

The results of the discussed algorithms and the original versions of these algorithms as discussed in [39] on the hidden test set, are shown in Table B.2. The initial version of the algorithm ended third in the last stage of the official

Table B.2: Results for the different proposed methods for multimodal heart beat detection on the hidden challenge test set. The old version of the MM_{WI} algorithm is the algorithm as described in [39].

Method	Se_{gr} (%)	$+P_{gr}$(%)	Se_{av} (%)	$+P_{av}$(%)	Ov(%)
$MM_{WI} - limited$ (old)	89.49	86.06	87.09	87.78	87.60
$MM_{WI} - limited$	**90.74**	**90.15**	**89.59**	**89.62**	**90.02**
$MM_{NI} - limited$	88.96	87.52	86.44	85.30	87.05

challenge [163]. A maximum score of 90.02% was found for the $MM_{WI} - limited$ algorithm on the hidden test set of the challenge.

The results of the other official top 5 entries from the 2014 Physionet/Computing in Cardiology challenge [75] and the unimodal *gqrs* algorithm (applied on the available ECG signal) are also shown in Table B.1.

All proposed multimodal algorithms result in a clear performance increase compared to when only the ECG signal is used. Especially for the subset of the MGH/MF database this increase is significant: a performance increase of 8.65% and 5.06% for the MM_{WI} and MM_{NI} algorithms. The performance gain in the challenge training set is a bit smaller as half of this dataset (the original training set) already had a performance of around 99.6% for the *gqrs* algorithm, which did not leave much room for further improvement. A similar trend is noticeable for the limited algorithm versions, but most often with a lower performance increase.

B.5 Discussion

B.5.1 Automatic signal detection

During the first phase of the automatic signal detection procedure, typical reasons for missing the most important signals (ECG, BP and ART) were noisy (reference) signals or too many false R peak detections in the reference ECG signal. Missed detections of PAP and CVP signals were typically due to too much interbeat variation in these signals, which would otherwise lead to bad peak detection in further processing. Some EEG, EOG and EMG signals contaminated with R peaks were also detected with this method, which would not be possible with the MM_{WI} algorithm. A few low-frequency signals (e.g. respiration) were however falsely detected if they accidentally had a harmonic frequency peak at HR_{av}.

After the second stage of the algorithm, most signals with inaccurate heart beat locations were removed from the signals of interest. This can be seen in the big drop in detection sensitivity for the SV signals, and to a lesser extent for the CVP and PAP signals. Some important signals were however also dropped, either by the absence of periodic positive peaks or due to bad R peak detection, but the sensitivity drop for these signals was significantly lower.

The introduced method for automatic signal detection does not only detect most of the pulsatile signals, but also discovered new signals that can improve the heart beat detection. Apart from that advantage, other advantages of this approach are that the algorithm can also be used reliably if

- there are no signal type identifiers available in the database;
- signal type errors can often occur; or
- the way of signal identifying is unknown or uncertain (full names versus abbreviations for IDs, different types of abbreviations, different language, etc.).

The method however still requires the availability of one known ECG signal in the available multimodal dataset.

One of the main drawbacks of this method is that it assumes that all signals are noise-free during the selected noise-free ECG period. This is however not always the case. A first problem is that the approach for detecting the noise-free ECG period assumes that the majority of the ECG is noise-free, but this is however not the case in some recordings. This could possibly lead to missing all potential signals of interest in that recording. A second problem is that the approach also assumes that when the reference ECG signal is noise-free, all other signals are also noise-free, which is not always true. A signal that is temporarily detached during the noise-free ECG period might be useful outside this period, but will not be selected as a signal of interest. A potential solution to the latter problem would be to apply the same method for detection of a noise-free period on each signal instead of only the known ECG. To see if a signal $s, 2 \leq s \leq N$, with N the number of signals in the recording, is a signal of interest, one can combine both similarity matrices in order to find a period in time in which both signals have a maximal similarity

$$a_{(s)} = \arg\max_{j} \sum_{n=0}^{l-1} M_{j-n}^{(1)} + M_{j-n}^{(s)} \tag{B.5}$$

The PSD values and their similarity to the ECG PSD values can then be computed in this segment. This procedure is left for future work. It however

still requires the reference ECG signal to be noise-free for at least 50% of the signal.

B.5.2 Multimodal heart beat detection

In both datasets and both algorithm versions, MM_{WI} resulted in a better performance than MM_{NI}, but this difference was minimal in case of the limited versions. Several reasons for the difference in performance can be found. In case the majority of the reference ECG signal is noisy, a noisy ECG segment is more likely to be selected. This can produce errors in HR_{av}, possibly leading to missing all potential signals of interest or aligning towards bad R peak locations. A second reason is the fact that MM_{WI} is able to switch the reference signal. By doing this, a noisy or disconnected ECG can be exchanged for another signal with a more logical RR sequence. The alignment step is therefore more reliable, typically leading to less shifted heart beat locations. By adding this procedure to the MM_{WI} algorithm, a performance increase of around 1.5% was found on the hidden challenge test set. MM_{NI} was however not able to find any peak location in 5 recordings in the hidden test set because of its dependency on the ECG signal. Finally, MM_{WI} is able to use the specific QRS algorithm *gqrs* on possible additional ECG signals in the MGH/MF database. MM_{NI} however had to use the general peak detector, which is less accurate and is not able to detect negative R peaks. If however the majority of the reference ECG is indeed noise-free, both algorithms have a very similar performance. For the initial challenge training set, MM_{NI} even worked slightly better as it was able to detect some extra signals with R peak interference, which was not possible with MM_{WI}.

The original version of the proposed algorithms was described in [39]. The main differences of the current versions compared to the original version are the addition of the automatic selection of the reference signal (MM_{WI}), the lowering of the PSD similarity threshold (MM_{NI}) and fixing of some programming bugs (MM_{NI} & MM_{WI}). By making these adjustments on the MM_{WI} algorithm, a performance increase of 2.42% was accomplished on the new hidden test set for the MM_{WI} algorithm (see Table B.2).

Compared to the other official top 5 challenge entries, both proposed limited methods would get a third place on the challenge training set. The difference with the better performing algorithms here is the fact that these use another R peak detector, which are probably better suited for ECG signals with pacemaker artifacts. The used *gqrs* algorithm is not able to deal well with these artifacts, typically leading to either shifted R peak locations or additional peaks for these artifacts. Beats from other signals could be aligned to these erroneous R peak

locations, which can lead to large series of shifted R peak locations. However for the MGH/MF subset, only the algorithm proposed in [200] gives a better performance. The algorithm does not put a limit on the number of signals to be used, but as it is typically around 5 signals in this dataset, computing it might be too long for using it on the Physionet server. The non-limited version of the MM_{WI} algorithm however gives the best performance over both datasets.

It should be noted that a fairly high amount of false and missed R peak detections occurred due to the mentioned shifted R peak locations. These sometimes occur in large series of shifted peaks, but the resulting immediate heart rate measurements only differ in the beginning and end of these series. The intermediate heart rate values are most often correct in this case. Because typically only the heart rate values themselves are used in practice, the real error rate is even lower than the R peak detection error.

The proposed algorithms worked well during different observed arrhythmias, leading to a performance increase compared to unimodal heart beat detection. Because the algorithms try to find the most "logical" RR sequence in the available R peak locations, it can however occur that a wrong R peak location is chosen in case of a special beat (e.g. a premature beat). These special beats lead to an abnormal looking heart rate value compared to previous values, and might therefore not be preferred if a more smooth RR sequence can be obtained with another available R peak location from another signal. This however only occurs if the corresponding R peak locations from both signals do not resemble well and if at least one of the signals contains artifacts during this period. The number of these errors were however limited in the observed datasets. The presence of arrhythmia's does not interfere with the detection of signals of interest.

The general peak detection algorithm was rather basic in order to be able to detect peaks in multiple different signals of interest. Nevertheless it was able to give similar results as specific BP and R peak detectors in case of positive peaks and limited noise. If the specific BP detector from the WFDB toolbox would be used in the MM_{WI} algorithm, performance slightly decreased on the discussed training sets. More enhanced QRS and BP peak detectors could however be used in order to further optimize the performance, but were not possible during the 2014 Computing in Cardiology challenge due to the limited available computation time [163].

B.6 Conclusion

The proposed multimodal algorithms were able to detect the heart beats more robustly compared to unimodal algorithms. The most important signals that can help to improve the heart beat detection were reliably detected in at least 90% of the cases, and also other signals with strong R peak interference were detected with the proposed method. The individual peaks of the different signals of interest were successfully combined by a procedure containing majority voting, heart beat location estimation and Hjorth's mobility. Its performance can be further increased by optimizing the individual peak detection algorithms or by incorporating reliability scores to the different signals during the voting procedure. Also noise could be removed from the ECG by using noise information from other available signals, which was however not the primary goal of the 2014 Computing in Cardiology challenge. The methods could be used to improve heart beat detection in several multimodal heart beat detection applications like in the ICU or for epileptic heart rate-based seizure detection in a home environment [44, 46].

Bibliography

[1] Adeli, H., and Ghosh-Dastidar, S. *Automated EEG-based diagnosis of neurological disorders: Inventing the future of neurology*. CRC press, 2010.

[2] 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.

[3] Akbani, R., Kwek, S., and Japkowicz, N. Applying support vector machines to imbalanced datasets. In *Machine Learning: ECML 2004*. Springer, 2004, pp. 39–50.

[4] 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.

[5] Arends, J., Thijs, R. D., Gutter, T., Ungureanu, C., Cluitmans, P., Van Dijk, J., van Andel, J., Tan, F., de Weerd, A., Vledder, B., et al. Multimodal nocturnal seizure detection in a residential care setting: A long-term prospective trial. *Neurology* (2018), 10–1212.

[6] Arends, J. B., van Dorp, J., van Hoek, D., Kramer, N., van Mierlo, P., van der Vorst, D., and Tan, F. I. Diagnostic accuracy of audio-based seizure detection in patients with severe epilepsy and an intellectual disability. *Epilepsy & Behavior 62* (2016), 180–185.

[7] Arends, J. B. A. M. Movement-based seizure detection. *Epilepsia 59*, S1 (2018), 30–35.

[8] Baumgartner, C., Lurger, S., and Leutmezer, F. Autonomic symptoms during epileptic seizures. *Epileptic Disorders 3*, 3 (2001), 103–16.

[9] Becq, G., Kahane, P., Minotti, L., Bonnet, S., and Guillemaud, R. Classification of epileptic motor manifestations and detection of tonic–clonic seizures with acceleration norm entropy. *IEEE Transactions on Biomedical Engineering 60*, 8 (2013), 2080–2088.

[10] Beniczky, S., Conradsen, I., Henning, O., Fabricius, M., and Wolf, P. Automated real-time detection of tonic-clonic seizures using a wearable emg device. *Neurology 90*, 5 (2018), e428–e434.

[11] Beniczky, S., Polster, T., Kjaer, T. W., and Hjalgrim, H. Detection of generalized tonic–clonic seizures by a wireless wrist accelerometer: a prospective, multicenter study. *Epilepsia 54*, 4 (2013), e48–e51.

[12] Bergmann, J. H., Chandaria, V., and McGregor, A. Wearable and implantable sensors: the patient's perspective. *Sensors 12*, 12 (2012), 16695–16709.

[13] Berntson, G. G., Bigger, J. T., Eckberg, D. L., Grossman, P., Kaufmann, P. G., Malik, M., Nagaraja, H. N., Porges, S. W., Saul, J. P., Stone, P. H., et al. Heart rate variability: origins, methods, and interpretive caveats. *Psychophysiology 34*, 6 (1997), 623–648.

[14] Bertrand, A. Distributed signal processing for wireless eeg sensor networks. *IEEE Transactions on Neural Systems and Rehabilitation Engineering 23*, 6 (2015), 923–935.

[15] Bialer, M., Johannessen, S. I., Levy, R. H., Perucca, E., Tomson, T., White, H. S., and Koepp, M. J. Seizure detection and neuromodulation: A summary of data presented at the xiii conference on new antiepileptic drug and devices (eilat xiii). *Epilepsy research 130* (2017), 27–36.

[16] Bidwell, J., Khuwatsamrit, T., Askew, B., Ehrenberg, J. A., and Helmers, S. Seizure reporting technologies for epilepsy treatment: A review of clinical information needs and supporting technologies. *Seizure 32* (2015), 109–117.

[17] Billiet, L., Swinnen, T., De Vlam, K., Westhovens, R., and Van Huffel, S. Recognition of physical activities from a single arm-worn accelerometer: A multiway approach. In *Informatics* (2018), vol. 5, Multidisciplinary Digital Publishing Institute, p. 20.

[18] Blachut, B., Hoppe, C., Surges, R., Elger, C., and Helmstaedter, C. Subjective seizure counts by epilepsy clinical drug trial participants are not reliable. *Epilepsy & Behavior 67* (2017), 122–127.

[19] Boon, P., Vonck, K., van Rijckevorsel, K., El Tahry, R., Elger, C. E., Mullatti, N., Schulze-Bonhage, A., Wagner, L., Diehl, B., Hamer, H., et al. A prospective, multicenter study of cardiac-based seizure detection to activate vagus nerve stimulation. *Seizure-European Journal of Epilepsy 32* (2015), 52–61.

[20] Brotherstone, R., Blackhall, B., and McLellan, A. Lengthening of corrected qt during epileptic seizures. *Epilepsia 51*, 2 (2010), 221–232.

[21] Bruno, E., Biondi, A., and Richardson, M. P. Pre-ictal heart rate changes: A systematic review and meta-analysis. *Seizure 55* (2018), 48–56.

[22] Bruno, E., Simblett, S., Lang, A., Biondi, A., Odoi, C., Schulze-Bonhage, A., Wykes, T., and Richardson, M. P. Wearable technology in epilepsy: The views of patients, caregivers, and healthcare professionals. *Epilepsy & Behavior 85* (2018), 141–149.

[23] Camm, A., Malik, M., Bigger, J., Breithardt, G., Cerutti, S., Cohen, R., Coumel, P., Fallen, E., Kennedy, H., Kleiger, R., et al. Heart rate variability: standards of measurement, physiological interpretation and clinical use. task force of the european society of cardiology and the north american society of pacing and electrophysiology. *Circulation 93*, 5 (1996), 1043–1065.

[24] Carlson, C., Arnedo, V., Cahill, M., and Devinsky, O. Detecting nocturnal convulsions: efficacy of the mp5 monitor. *Seizure-European Journal of Epilepsy 18*, 3 (2009), 225–227.

[25] Carrette, E., Vonck, K., and Boon, P. The management of pharmacologically refractory epilepsy. *Int J Clin Rev 1*, 02 (2011).

[26] Chung, S. S., Gerber, P., and Kirlin, K. A. Ictal eye closure is a reliable indicator for psychogenic nonepileptic seizures. *Neurology 66*, 11 (2006), 1730–1731.

[27] Clifford, G. D., Azuaje, F., McSharry, P., et al. *Advanced methods and tools for ECG data analysis*. Artech house Norwood, MA, 2006.

[28] Cogan, D., Birjandtalab, J., Nourani, M., Harvey, J., and Nagaraddi, V. Multi-biosignal analysis for epileptic seizure monitoring. *International journal of neural systems 27*, 01 (2017), 1650031.

[29] Cogan, D., Heydarzadeh, M., and Nourani, M. Personalization of noneeg-based seizure detection systems. In *Engineering in Medicine*

and Biology Society (EMBC), 2016 IEEE 38th Annual International Conference of the (2016), IEEE, pp. 6349–6352.

[30] Commission on Classification and Terminology of the International League Against Epilepsy. Proposal for revised classification of epilepsies and epileptic syndromes. *Epilepsia 30* (1989), 389–399.

[31] Conradsen, I., Beniczky, S., Hoppe, K., Wolf, P., and Sorensen, H. B. Automated algorithm for generalized tonic–clonic epileptic seizure onset detection based on semg zero-crossing rate. *Biomedical Engineering, IEEE Transactions on 59*, 2 (2012), 579–585.

[32] Conradsen, I., Beniczky, S., Wolf, P., Terney, D., Sams, T., and Sorensen, H. B. Multi-modal intelligent seizure acquisition (misa) system—a new approach towards seizure detection based on full body motion measures. In *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE* (2009), IEEE, pp. 2591–2595.

[33] Conradsen, I., Wolf, P., Sams, T., Sorensen, H. B., and Beniczky, S. Patterns of muscle activation during generalized tonic and tonic–clonic epileptic seizures. *Epilepsia 52*, 11 (2011), 2125–2132.

[34] Cuppens, K., Bonroy, B., Van de Vel, A., Ceulemans, B., Lagae, L., Tuytelaars, T., Van Huffel, S., and Vanrumste, B. Automatic video detection of nocturnal epileptic movement based on motion tracks. In *Proceedings of the international conference on bio-inspired systems and signal processing (Biosignals)* (2012), INSTICC, pp. 342–345.

[35] Cuppens, K., Karsmakers, P., Van de Vel, A., Bonroy, B., Milosevic, M., Luca, S., Croonenborghs, T., Ceulemans, B., Lagae, L., Van Huffel, S., et al. Accelerometry-based home monitoring for detection of nocturnal hypermotor seizures based on novelty detection. *Biomedical and Health Informatics, IEEE Journal of 18*, 3 (2014), 1026–1033.

[36] Dalton, A., Patel, S., Chowdhury, A. R., Welsh, M., Pang, T., Schachter, S., OLaighin, G., and Bonato, P. Development of a body sensor network to detect motor patterns of epileptic seizures. *IEEE Transactions on Biomedical Engineering 59*, 11 (2012), 3204–3211.

[37] 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.

[38] De Cooman, T., Carrette, E., Boon, P., Meurs, A., and Van Huffel, S. Online seizure detection in adults with temporal lobe epilepsy using single-lead ECG. In *Proceedings of the 22nd European Signal Processing Conference (EUSIPCO), 2014* (2014), IEEE, pp. 1532–1536.

[39] De Cooman, T., Goovaerts, G., Varon, C., Widjaja, D., and Van Huffel, S. Heart Beat Detection in Multimodal Data Using Signal Recognition and Beat Location Estimation. In *Proc. of the 41st Annual Computing in Cardiology (CinC2014)* (2014), pp. 549–552.

[40] De Cooman, T., Goovaerts, G., Varon, C., Widjaja, D., Willemen, T., and Van Huffel, S. Heart beat detection in multimodal data using automatic relevant signal detection. *Physiological measurement 36*, 8 (2015), 1691.

[41] De Cooman, T., Kjær, T. W., Van Huffel, S., and Sørensen, H. B. D. Adaptive heart rate-based epileptic seizure detection using real-time user feedback. *Physiological measurement 39*, 1 (2017), 014005.

[42] De Cooman, T., Van de Vel, A., Ceulemans, B., Lagae, L., Van Paesschen, W., Vanrumste, B., and Van Huffel, S. Estimation of the Maximal Heart Rate to Improve Online Tonic-Clonic Seizure Detection using ECG. In *Computing in Cardiology Conference (CinC), 2015* (2015), IEEE, pp. 977–980.

[43] De Cooman, T., Van de Vel, A., Ceulemans, B., Lagae, L., Vanrumste, B., and Van Huffel, S. Online Detection of Tonic-Clonic Seizures in Pediatric Patients using ECG and Low-Complexity Incremental Novelty Detection. In *Proc. of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC2015)* (2015), IEEE, pp. 5597–5600.

[44] De Cooman, T., Varon, C., Hunyadi, B., Van Paesschen, W., Lagae, L., and Van Huffel, S. Online automated seizure detection in temporal lobe epilepsy patients using single-lead ecg. *International Journal of Neural Systems* (2017), 1750022.

[45] De Cooman, T., Varon, C., Van de Vel, A., Ceulemans, B., Lagae, L., and Van Huffel, S. Semi-supervised one-class transfer learning for heart rate based epileptic seizure detection. In *Proceedings of the 44th Computing in Cardiology Conference (CinC 2017)* (2017), pp. 1–4.

[46] De Cooman, T., Varon, C., Van de Vel, A., Ceulemans, B., Lagae, L., and Van Huffel, S. Comparison and combination of

electrocardiogram, electromyogram and accelerometry for tonic-clonic seizure detection in children. In *2018 IEEE EMBS International Conference on Biomedical Health Informatics (BHI)* (March 2018), pp. 438–441.

[47] De Cooman, T., Varon, C., Van de Vel, A., Jansen, K., Ceulemans, B., Lagae, L., and Van Huffel, S. Adaptive nocturnal seizure detection using heart rate and low-complexity novelty detection. *Seizure 59* (2018), 48–53.

[48] Delamont, R. S., and Walker, M. C. Pre-ictal autonomic changes. *Epilepsy research 97*, 3 (2011), 267–272.

[49] Devinsky, O. Effects of seizures on autonomic and cardiovascular function. *Epilepsy currents 4*, 2 (2004), 43–46.

[50] Di Gennaro, G., Quarato, P. P., Sebastiano, F., Esposito, V., Onorati, P., Grammaldo, L. G., Meldolesi, G. N., Mascia, A., Falco, C., Scoppetta, C., et al. Ictal heart rate increase precedes eeg discharge in drug-resistant mesial temporal lobe seizures. *Clinical Neurophysiology 115*, 5 (2004), 1169–1177.

[51] Doyle, O., Temko, A., Marnane, W., Lightbody, G., and Boylan, G. Heart rate based automatic seizure detection in the newborn. *Medical engineering & physics 32*, 8 (2010), 829–839.

[52] Eckmann, J.-P., Kamphorst, S. O., and Ruelle, D. Recurrence plots of dynamical systems. *EPL (Europhysics Letters) 4*, 9 (1987), 973–977.

[53] Elger, C. E., and Hoppe, C. Diagnostic challenges in epilepsy: seizure under-reporting and seizure detection. *The Lancet Neurology 17*, 3 (2018), 279–288.

[54] Elliott, R. E., Morsi, A., Kalhorn, S. P., Marcus, J., Sellin, J., Kang, M., Silverberg, A., Rivera, E., Geller, E., Carlson, C., et al. Vagus nerve stimulation in 436 consecutive patients with treatment-resistant epilepsy: long-term outcomes and predictors of response. *Epilepsy & Behavior 20*, 1 (2011), 57–63.

[55] Empatica. Embrace. `https://www.empatica.com/embrace/`. Accessed: 2018-06-22.

[56] Englot, D. J., Chang, E. F., and Auguste, K. I. Vagus nerve stimulation for epilepsy: a meta-analysis of efficacy and predictors of response: a review. *Journal of neurosurgery 115*, 6 (2011), 1248–1255.

[57] Epstein, M. A., Sperling, M. R., and O'connor, M. J. Cardiac rhythm during temporal lobe seizures. *Neurology 42*, 1 (1992), 50–50.

[58] 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.

[59] Fisher, R. S., Blum, D. E., DiVentura, B., Vannest, J., Hixson, J. D., Moss, R., Herman, S. T., Fureman, B. E., and French, J. A. Seizure diaries for clinical research and practice: limitations and future prospects. *Epilepsy Behav. 24*, 3 (2012), 304–310.

[60] Fisher, R. S., Cross, J. H., D'souza, C., French, J. A., Haut, S. R., Higurashi, N., Hirsch, E., Jansen, F. E., Lagae, L., Moshé, S. L., et al. Instruction manual for the ilae 2017 operational classification of seizure types. *Epilepsia 58*, 4 (2017), 531–542.

[61] Fisher, R. S., Cross, J. H., French, J. A., Higurashi, N., Hirsch, E., Jansen, F. E., Lagae, L., Moshé, S. L., Peltola, J., Roulet Perez, E., et al. Operational classification of seizure types by the international league against epilepsy: Position paper of the ilae commission for classification and terminology. *Epilepsia 58*, 4 (2017), 522–530.

[62] Fountas, K. N., and Smith, J. A novel closed-loop stimulation system in the control of focal, medically refractory epilepsy. In *Operative Neuromodulation*. Springer, 2007, pp. 357–362.

[63] Freestone, D. R., Karoly, P. J., and Cook, M. J. A forward-looking review of seizure prediction. *Current opinion in neurology 30*, 2 (2017), 167–173.

[64] Fujiwara, K., Miyajima, M., Yamakawa, T., Abe, E., Suzuki, Y., Sawada, Y., Kano, M., Maehara, T., Ohta, K., Sasai-Sakuma, T., et al. Epileptic seizure prediction based on multivariate statistical process control of heart rate variability features. *IEEE Transactions on Biomedical Engineering 63*, 6 (2016), 1321–1332.

[65] Fulton, S., Poppel, K. V., McGregor, A., Ellis, M., Patters, A., and Wheless, J. Prospective study of 2 bed alarms for detection of nocturnal seizures. *Journal of child neurology 28*, 11 (2013), 1430–1433.

[66] Fürbass, F., Kampusch, S., Kaniusas, E., Koren, J., Pirker, S., Hopfengärtner, R., Stefan, H., Kluge, T., and Baumgartner, C. Automatic multimodal detection for long-term seizure documentation in epilepsy. *Clinical Neurophysiology 128*, 8 (2017), 1466–1472.

[67] Gadhoumi, K., Lina, J.-M., Mormann, F., and Gotman, J. Seizure prediction for therapeutic devices: A review. *Journal of neuroscience methods 260* (2016), 270–282.

[68] Gaitatzis, A., and Sander, J. W. The mortality of epilepsy revisited. *Epileptic Disorders 6*, 1 (2004), 3–13.

[69] Garbey, M., Sun, N., Merla, A., and Pavlidis, I. Contact-free measurement of cardiac pulse based on the analysis of thermal imagery. *IEEE transactions on Biomedical Engineering 54*, 8 (2007), 1418–1426.

[70] Geertsema, E. E., Thijs, R. D., Gutter, T., Vledder, B., Arends, J. B., Leijten, F. S., Visser, G. H., and Kalitzin, S. N. Automated video-based detection of nocturnal convulsive seizures in a residential care setting. *Epilepsia 59*, S1 (2018), 53–60.

[71] Ghafoori, Z., Rajasegarar, S., Erfani, S. M., Karunasekera, S., and Leckie, C. A. Unsupervised parameter estimation for one-class support vector machines. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining* (2016), Springer, pp. 183–195.

[72] Ghosh-Dastidar, S., Adeli, H., and Dadmehr, N. Principal component analysis-enhanced cosine radial basis function neural network for robust epilepsy and seizure detection. *IEEE Transactions on Biomedical Engineering 55*, 2 (2008), 512–518.

[73] Ghougassian, D. F., d'Souza, W., Cook, M. J., and O'Brien, T. J. Evaluating the utility of inpatient video-eeg monitoring. *Epilepsia 45*, 8 (2004), 928–932.

[74] Gieraltowski, J., Ciuchcinski, K., Grzegorczyk, I., Kozna, K., Solinski, M., and Podziemski, P. Algorithm for Detection of Heart Rate from Noisy Multimodal Recordings. In *Proc. of the 41st Annual Computing in Cardiology (CinC2014)* (2014), pp. 253–256.

[75] 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.

[76] Greene, B., Marnane, W., Lightbody, G., Reilly, R., and Boylan, G. Classifier models and architectures for eeg-based neonatal seizure detection. *Physiological Measurement 29*, 10 (2008), 1157.

[77] Greene, B. R., De Chazal, P., Boylan, G. B., Connolly, S., and Reilly, R. B. Electrocardiogram based neonatal seizure detection. *Biomedical Engineering, IEEE Transactions on 54*, 4 (2007), 673–682.

[78] Gu, Y., Cleeren, E., Dan, J., Claes, K., Van Paesschen, W., Van Huffel, S., and Hunyadi, B. Comparison between scalp eeg and behind-the-ear eeg for development of a wearable seizure detection system for patients with focal epilepsy. *Sensors 18*, 1 (2017), 29.

[79] Güler, İ., and Übeyli, E. D. An expert system for detection of electrocardiographic changes in patients with partial epilepsy using wavelet-based neural networks. *Expert Systems 22*, 2 (2005), 62–71.

[80] Gutierrez, E. G., Crone, N. E., Kang, J. Y., Carmenate, Y. I., and Krauss, G. L. Strategies for non-eeg seizure detection and timing for alerting and interventions with tonic–clonic seizures. *Epilepsia 59* (2018), 36–41.

[81] Halford, J. J., Sperling, M. R., Nair, D. R., Dlugos, D. J., Tatum, W. O., Harvey, J., French, J. A., Pollard, J. R., Faught, E., Noe, K. H., et al. Detection of generalized tonic–clonic seizures using surface electromyographic monitoring. *Epilepsia 58*, 11 (2017), 1861–1869.

[82] Hallett, M. Myoclonus: Relation to epilepsy. *Epilepsia 26*, s1 (1985), S67–S77.

[83] Hampel, K. G., Jahanbekam, A., Elger, C. E., and Surges, R. Seizure-related modulation of systemic arterial blood pressure in focal epilepsy. *Epilepsia 57*, 10 (2016), 1709–1718.

[84] Hampel, K. G., Vatter, H., Elger, C. E., and Surges, R. Cardiac-based vagus nerve stimulation reduced seizure duration in a patient with refractory epilepsy. *Seizure 26* (2015), 81–85.

[85] Hauser, W. A. Seizure disorders: the changes with age. *Epilepsia 33*, s4 (1992), 6–14.

[86] Herman, S., Walczak, T., and Bazil, C. Distribution of partial seizures during the sleep–wake cycle. *Neurology 56*, 11 (2001), 1453–1459.

[87] Hesdorffer, D. C., Tomson, T., Benn, E., Sander, J. W., Nilsson, L., Langan, Y., Walczak, T. S., Beghi, E., Brodie, M. J., and Hauser, A. Combined analysis of risk factors for sudep. *Epilepsia 52*, 6 (2011), 1150–1159.

[88] Hirsch, M., Altenmüller, D.-M., and Schulze-Bonhage, A. Latencies from intracranial seizure onset to ictal tachycardia: A comparison to surface eeg patterns and other clinical signs. *Epilepsia 10* (2015), 1639–1647.

[89] Hirtz, D., Thurman, D., Gwinn-Hardy, K., Mohamed, M., Chaudhuri, A., and Zalutsky, R. How common are the "common" neurologic disorders? *Neurology 68*, 5 (2007), 326–337.

[90] Hjorth, B. EEG analysis based on time domain properties. *Electroencephalography and clinical neurophysiology 29*, 3 (1970), 306–310.

[91] Hoppe, C., Feldmann, M., Blachut, B., Surges, R., Elger, C. E., and Helmstaedter, C. Novel techniques for automated seizure registration: patients' wants and needs. *Epilepsy & Behavior 52* (2015), 1–7.

[92] Hoppe, C., Poepel, A., and Elger, C. E. Epilepsy: accuracy of patient seizure counts. *Archives of neurology 64*, 11 (2007), 1595–1599.

[93] Hunyadi, B., Signoretto, M., Van Paesschen, W., Suykens, J. A., Van Huffel, S., and De Vos, M. Incorporating structural information from the multichannel eeg improves patient-specific seizure detection. *Clinical Neurophysiology 123*, 12 (2012), 2352–2361.

[94] Ihle, M., Feldwisch-Drentrup, H., Teixeira, C. A., Witon, A., Schelter, B., Timmer, J., and Schulze-Bonhage, A. Epilepsiae–a european epilepsy database. *Computer methods and programs in biomedicine 106*, 3 (2012), 127–138.

[95] Jain, A. K. Data clustering: 50 years beyond k-means. *Pattern recognition letters 31*, 8 (2010), 651–666.

[96] Jansen, K., and Lagae, L. Cardiac changes in epilepsy. *Seizure 19*, 8 (2010), 455–460.

[97] 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.

[98] 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.

[99] Jeppesen, J., Beniczky, S., Fuglsang-Frederiksen, A., Sidenius, P., and Jasemian, Y. Detection of epileptic-seizures by means of power spectrum analysis of heart rate variability: a pilot study. *Technology and Health Care 18*, 6 (2010), 417–426.

[100] Jeppesen, J., Beniczky, S., Johansen, P., Sidenius, P., and Fuglsang-Frederiksen, A. Detection of epileptic seizures with a modified heart rate variability algorithm based on lorenz plot. *Seizure 24* (2015), 1–7.

[101] Jeppesen, J., Fuglsang-Frederiksen, A., Brugada, R., Pedersen, B., Rubboli, G., Johansen, P., and Beniczky, S. Heart rate variability analysis indicates preictal parasympathetic overdrive preceding seizure-induced cardiac dysrhythmias leading to sudden unexpected death in a patient with epilepsy. *Epilepsia 55*, 7 (2014), e67–e71.

[102] Johnson, A., Behar, J., Andreotti, F., Clifford, G., and Oster, J. R-Peak Estimation using Multimodal Lead Switching. In *Proc. of the 41st Annual Computing in Cardiology (CinC2014)* (2014), pp. 281–284.

[103] Jory, C., Shankar, R., Coker, D., McLean, B., Hanna, J., and Newman, C. Safe and sound? a systematic literature review of seizure detection methods for personal use. *Seizure-European Journal of Epilepsy 36* (2016), 4–15.

[104] Karayiannis, N. B., Tao, G., Frost, J. D., Wise, M. S., Hrachovy, R. A., and Mizrahi, E. M. Automated detection of videotaped neonatal seizures based on motion segmentation methods. *Clinical Neurophysiology 117*, 7 (2006), 1585–1594.

[105] Karoly, P. J., Ung, H., Grayden, D. B., Kuhlmann, L., Leyde, K., Cook, M. J., and Freestone, D. R. The circadian profile of epilepsy improves seizure forecasting. *Brain 140*, 8 (2017), 2169–2182.

[106] Kerem, D., and Geva, A. Forecasting epilepsy from the heart rate signal. *Medical and biological engineering and computing 43*, 2 (2005), 230–239.

[107] Kerling, F., Mueller, S., Pauli, E., and Stefan, H. When do patients forget their seizures? an electroclinical study. *Epilepsy & Behavior 9*, 2 (2006), 281–285.

[108] Klabunde, R. E. Normal heart rhythm. `https://www.cvphysiology.com/Arrhythmias/A002`. Accessed: 2018-07-24.

[109] Kohler, 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.

[110] Kverneland, M., Molteberg, E., Iversen, P. O., Veierød, M. B., Taubøll, E., Selmer, K. K., and Nakken, K. O. Effect of modified atkins diet in adults with drug-resistant focal epilepsy: A randomized clinical trial. *Epilepsia 59*, 8 (2018), 1567–1576.

[111] 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.

[112] Larsen, S. N., Conradsen, I., Beniczky, S., and Sorensen, H. B. Detection of tonic epileptic seizures based on surface electromyography. In *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE* (2014), IEEE, pp. 942–945.

[113] Leijten, F. S., Consortium, D. T., van Andel, J., Ungureanu, C., Arends, J., Tan, F., van Dijk, J., Petkov, G., Kalitzin, S., Gutter, T., et al. Multimodal seizure detection: A review. *Epilepsia 59* (2018), 42–47.

[114] 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.

[115] Lhatoo, S. D., Faulkner, H. J., Dembny, K., Trippick, K., Johnson, C., and Bird, J. M. An electroclinical case-control study of sudden unexpected death in epilepsy. *Annals of neurology 68*, 6 (2010), 787–796.

[116] Li, C., Zheng, C., and Tai, C. Detection of ecg characteristic points using wavelet transforms. *IEEE Transactions on biomedical Engineering 42*, 1 (1995), 21–28.

[117] Luca, S., Karsmakers, P., Cuppens, K., Croonenborghs, T., Van de Vel, A., Ceulemans, B., Lagae, L., Van Huffel, S., and Vanrumste, B. Detecting rare events using extreme value statistics applied to epileptic convulsions in children. *Artificial intelligence in medicine 60*, 2 (2014), 89–96.

[118] Markou, M., and Singh, S. Novelty detection: a review—part 1: statistical approaches. *Signal processing 83*, 12 (2003), 2481–2497.

[119] Markou, M., and Singh, S. Novelty detection: a review—part 2:: neural network based approaches. *Signal processing 83*, 12 (2003), 2499–2521.

[120] Martin, K., Jackson, C., Levy, R., and Cooper, P. Ketogenic diet and other dietary treatments for epilepsy. *Cochrane Database of Systematic Reviews*, 2 (2016).

[121] Marwan, N., Romano, M. C., Thiel, M., and Kurths, J. Recurrence plots for the analysis of complex systems. *Physics reports 438*, 5-6 (2007), 237–329.

[122] Massé, F., Van Bussel, M., Serteyn, A., Arends, J., and Penders, J. Miniaturized wireless ECG monitor for real-time detection of epileptic seizures. *ACM Transactions on Embedded Computing Systems (TECS) 12*, 4 (2013), 102.

[123] 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.

[124] Mayer, H., Benninger, F., Urak, L., Plattner, B., Geldner, J., and Feucht, M. Ekg abnormalities in children and adolescents with symptomatic temporal lobe epilepsy. *Neurology 63*, 2 (2004), 324–328.

[125] Meritam, P., Ryvlin, P., and Beniczky, S. User-based evaluation of applicability and usability of a wearable accelerometer device for detecting bilateral tonic–clonic seizures: A field study. *Epilepsia 59* (2018), 48–52.

[126] Milošević, M., Van de Vel, A., Bonroy, B., Ceulemans, B., Lagae, L., Vanrumste, B., and Van Huffel, S. Automated detection of tonic–clonic seizures using 3-d accelerometry and surface electromyography in pediatric patients. *IEEE journal of biomedical and health informatics 20*, 5 (2016), 1333–1341.

[127] Moeyersons, J., Varon, C., Testelmans, D., Buyse, B., and Van Huffel, S. Ecg artefact detection using ensemble decision trees. In *Proceedings of the 44th Computing in Cardiology Conference (CinC 2017)* (2017), vol. 44, pp. 1–4.

[128] Moody, G. B., and Mark, R. G. The impact of the mit-bih arrhythmia database. *Engineering in Medicine and Biology Magazine, IEEE 20*, 3 (2001), 45–50.

[129] Morrell, M. J. Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. *Neurology 77*, 13 (2011), 1295–1304.

[130] Murthy, J. N., van Jaarsveld, J., Fei, J., Pavlidis, I., Harrykissoon, R. I., Lucke, J. F., Faiz, S., and Castriotta, R. J. Thermal infrared imaging: a novel method to monitor airflow during polysomnography. *Sleep 32*, 11 (2009), 1521–1527.

[131] Neal, E. G., Chaffe, H., Schwartz, R. H., Lawson, M. S., Edwards, N., Fitzsimmons, G., Whitney, A., and Cross, J. H. The ketogenic diet for the treatment of childhood epilepsy: a randomised controlled trial. *The Lancet Neurology 7*, 6 (2008), 500–506.

[132] Nei, M., Ho, R. T., Abou-Khalil, B. W., Drislane, F. W., Liporace, J., Romeo, A., and Sperling, M. R. Eeg and ecg in sudden unexplained death in epilepsy. *Epilepsia 45*, 4 (2004), 338–345.

[133] Nei, M., Ho, R. T., and Sperling, M. R. Ekg abnormalities during partial seizures in refractory epilepsy. *Epilepsia 41*, 5 (2000), 542–548.

[134] Neura. Temporal lobe. `https://library.neura.edu.au/browse-library/physical-features/brain-regions/temporal-lobe/`. Accessed: 2018-06-19.

[135] Nijsen, T. M., Aarts, R. M., Arends, J. B., and Cluitmans, P. J. Automated detection of tonic seizures using 3-d accelerometry. In *4th European Conference of the International Federation for Medical and Biological Engineering* (2009), Springer, pp. 188–191.

[136] Nijsen, T. M., Aarts, R. M., Cluitmans, P. J., and Griep, P. A. Time-frequency analysis of accelerometry data for detection of myoclonic seizures. *IEEE Transactions on Information Technology in Biomedicine 14*, 5 (2010), 1197–1203.

[137] Nijsen, T. M., Arends, J. B., Griep, P. A., and Cluitmans, P. J. The potential value of three-dimensional accelerometry for detection of motor seizures in severe epilepsy. *Epilepsy & Behavior 7*, 1 (2005), 74–84.

[138] Nilsen, K. B., Haram, M., Tangedal, S., Sand, T., and Brodtkorb, E. Is elevated pre-ictal heart rate associated with secondary generalization in partial epilepsy? *Seizure 19*, 5 (2010), 291–295.

[139] Onorati, F., Regalia, G., Caborni, C., Migliorini, M., Bender, D., Poh, M.-Z., Frazier, C., Kovitch Thropp, E., Mynatt, E. D., Bidwell, J., et al. Multicenter clinical assessment of improved wearable multimodal convulsive seizure detectors. *Epilepsia 58*, 11 (2017), 1870–1879.

[140] Opherk, C., Coromilas, J., and Hirsch, L. J. Heart rate and ekg changes in 102 seizures: analysis of influencing factors. *Epilepsy research 52*, 2 (2002), 117–127.

[141] O'Regan, M. E., and Brown, J. K. Abnormalities in cardiac and respiratory function observed during seizures in childhood. *Developmental medicine and child neurology 47*, 1 (2005), 4–9.

[142] Osorio, I. Automated seizure detection using EKG. *International Journal of Neural Systems 24*, 02 (2014), 1450001.

[143] Pan, J., and Tompkins, W. J. A real-time qrs detection algorithm. *Biomedical Engineering, IEEE Transactions on 3* (1985), 230–236.

[144] Patel, A. D., Moss, R., Rust, S. W., Patterson, J., Strouse, R., Gedela, S., Haines, J., and Lin, S. M. Patient-centered design criteria for wearable seizure detection devices. *Epilepsy & Behavior 64* (2016), 116–121.

[145] Patterson, A. L., Mudigoudar, B., Fulton, S., McGregor, A., Van Poppel, K., Wheless, M. C., Brooks, L., and Wheless, J. W. Smartwatch by smartmonitor: assessment of seizure detection efficacy for various seizure types in children, a large prospective single-center study. *Pediatric neurology 53*, 4 (2015), 309–311.

[146] Penzel, T., Kantelhardt, J. W., Grote, L., Peter, J.-H., and Bunde, A. Comparison of detrended fluctuation analysis and spectral analysis for heart rate variability in sleep and sleep apnea. *IEEE Transactions on biomedical engineering 50*, 10 (2003), 1143–1151.

[147] Picard, R. W., Migliorini, M., Caborni, C., Onorati, F., Regalia, G., Friedman, D., and Devinsky, O. Wrist sensor reveals sympathetic hyperactivity and hypoventilation before probable sudep. *Neurology 89*, 6 (2017), 633–635.

[148] Poggio, T., and Cauwenberghs, G. Incremental and decremental support vector machine learning. *Advances in neural information processing systems 13* (2001), 409–415.

[149] Poh, M.-Z., Loddenkemper, T., Reinsberger, C., Swenson, N., Goyal, S., Madsen, J., and Picard, R. W. Autonomic changes with seizures correlate with postictal eeg suppression. *Neurology 78*, 23 (2012), 1868–1876.

[150] Poh, M.-Z., Loddenkemper, T., Reinsberger, C., Swenson, N. C., Goyal, S., Sabtala, M. C., Madsen, J. R., and Picard, R. W.

Convulsive seizure detection using a wrist-worn electrodermal activity and accelerometry biosensor. *Epilepsia 53*, 5 (2012), e93–e97.

[151] Quint, S. R., Messenheimer, J. A., Tennison, M. B., and Nagle, H. T. Assessing autonomic activity from the ekg related to seizure onset detection and localization. In *Computer-Based Medical Systems, 1989. Proceedings., Second Annual IEEE Symposium on* (1989), IEEE, pp. 2–9.

[152] Ramgopal, S., Thome-Souza, S., Jackson, M., Kadish, N. E., Fernández, I. S., Klehm, J., Bosl, W., Reinsberger, C., Schachter, S., and Loddenkemper, T. Seizure detection, seizure prediction, and closed-loop warning systems in epilepsy. *Epilepsy & behavior 37* (2014), 291–307.

[153] Raya, M. A. D., and Sison, L. G. Adaptive noise cancelling of motion artifact in stress ecg signals using accelerometer. In *Proceedings of the Second Joint 24th Annual Conference of Engineering in Medicine and Biology and the Annual Fall Meeting of the Biomedical Engineering Society* (2002), vol. 2, pp. 1756–1757.

[154] Ryvlin, P., and Beniczky, S. Seizure detection and mobile health devices in epilepsy: Update and future developments. *Epilepsia 59* (2018), 7–8.

[155] Ryvlin, P., Nashef, L., Lhatoo, S. D., Bateman, L. M., Bird, J., Bleasel, A., Boon, P., Crespel, A., Dworetzky, B. A., Høgenhaven, H., et al. Incidence and mechanisms of cardiorespiratory arrests in epilepsy monitoring units (mortemus): a retrospective study. *The Lancet Neurology 12*, 10 (2013), 966–977.

[156] Ryvlin, P., Nashef, L., and Tomson, T. Prevention of sudden unexpected death in epilepsy: a realistic goal? *Epilepsia 54* (2013), 23–28.

[157] Sabesan, S., and Sankar, R. Improving long-term management of epilepsy using a wearable multimodal seizure detection system. *Epilepsy & Behavior 46* (2015), 56–57.

[158] Scheffer, I. E., Berkovic, S., Capovilla, G., Connolly, M. B., French, J., Guilhoto, L., Hirsch, E., Jain, S., Mathern, G. W., Moshé, S. L., et al. Ilae classification of the epilepsies: position paper of the ilae commission for classification and terminology. *Epilepsia 58*, 4 (2017), 512–521.

[159] Schernthaner, C., Lindinger, G., Pötzelberger, K., Zeiler, K., and Baumgartner, C. Autonomic epilepsy–the influence of epileptic

discharges on heart rate and rhythm. *Wiener Klinische Wochenschrift 111*, 10 (1999), 392–401.

[160] Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., and Williamson, R. C. Estimating the support of a high-dimensional distribution. *Neural computation 13*, 7 (2001), 1443–1471.

[161] Schölkopf, B., Williamson, R. C., Smola, A. J., Shawe-Taylor, J., and Platt, J. C. Support vector method for novelty detection. In *Advances in neural information processing systems* (2000), pp. 582–588.

[162] Schulze-Bonhage, A., Sales, F., Wagner, K., Teotonio, R., Carius, A., Schelle, A., and Ihle, M. Views of patients with epilepsy on seizure prediction devices. *Epilepsy & behavior 18*, 4 (2010), 388–396.

[163] Silva, I., Moody, B., Behar, J., Johnson, A., Oster, J., Clifford, G. D., and Moody, G. B. Robust detection of heart beats in multimodal data. *Physiological measurement 36*, 8 (2015), 1629.

[164] Silva, I., and Moody, G. An open-source toolbox for analysing and processing physionet databases in matlab and octave. *Journal of Open Research Software 2*, 1 (2014), e27.

[165] Silvetti, M. S., Drago, F., and Ragonese, P. Heart rate variability in healthy children and adolescents is partially related to age and gender. *International journal of cardiology 81*, 2-3 (2001), 169–174.

[166] Smith, P., Howell, S., Owen, L., and Blumhardt, L. Profiles of instant heart rate during partial seizures. *Electroencephalography and clinical neurophysiology 72*, 3 (1989), 207–217.

[167] Smithson, W. H., Colwell, B., and Hanna, J. Sudden unexpected death in epilepsy: addressing the challenges. *Current neurology and neuroscience reports 14*, 12 (2014), 502.

[168] Solomon, N., and McHale, K. An overview of epilepsy in children and young people. *Nursing Children and Young People (through 2013) 24*, 6 (2012), 28.

[169] Son, W. H., Hwang, W. S., Koo, D. L., Hwang, K. J., Kim, D. Y., Seo, J.-H., Na, G.-Y., Joo, E. Y., Hong, S. B., and Seo, D.-W. The difference in heart rate change between temporal and frontal lobe seizures during peri-ictal period. *Journal of epilepsy research 6*, 1 (2016), 16.

[170] Stefanidou, M., Carlson, C., and Friedman, D. The relationship between seizure onset zone and ictal tachycardia: an intracranial eeg study. *Clinical Neurophysiology 126*, 12 (2015), 2255–2260.

[171] Sun, F. T., Morrell, M. J., and Wharen Jr, R. E. Responsive cortical stimulation for the treatment of epilepsy. *Neurotherapeutics 5*, 1 (2008), 68–74.

[172] Surges, R., Scott, C. A., and Walker, M. C. Enhanced qt shortening and persistent tachycardia after generalized seizures. *Neurology 74*, 5 (2010), 421–426.

[173] Suykens, J. A., Van Gestel, T., and De Brabanter, J. *Least squares support vector machines*. World Scientific, 2002.

[174] Szabó, C. Á., Morgan, L. C., Karkar, K. M., Leary, L. D., Lie, O. V., Girouard, M., and Cavazos, J. E. Electromyography-based seizure detector: Preliminary results comparing a generalized tonic–clonic seizure detection algorithm to video-eeg recordings. *Epilepsia 56*, 9 (2015), 1432–1437.

[175] Teo, S., Yang, B., Huan, D abd Christopher, M., and Su, Y. Robust Detection of Heart Beats using ECG, BP and EEG signals. In *Proc. of the 41st Annual Computing in Cardiology (CinC2014)* (2014), pp. 565–568.

[176] Tzallas, A. T., Tsalikakis, D. G., Karvounis, E. C., Astrakas, L., Tzaphlidou, M., Tsipouras, M. G., and Konitsiotis, S. *Automated epileptic seizure detection methods: a review study*. INTECH Open Access Publisher, 2012.

[177] Ulate-Campos, A., Coughlin, F., Gainza-Lein, M., Fernández, I. S., Pearl, P., and Loddenkemper, T. Automated seizure detection systems and their effectiveness for each type of seizure. *Seizure-European Journal of Epilepsy 40* (2016), 88–101.

[178] Ungureanu, C., Bui, V., Roosmalen, W., Aarts, R. M., Arends, J. B., Verhoeven, R., Lukkien, J. J., et al. A wearable monitoring system for nocturnal epileptic seizures. In *Medical Information and Communication Technology (ISMICT), 2014 8th International Symposium on* (2014), IEEE, pp. 1–5.

[179] Van Andel, J., Leijten, F., Van Delden, H., and van Thiel, G. What makes a good home-based nocturnal seizure detector? a value sensitive design. *PloS one 10*, 4 (2015), e0121446.

[180] van Andel, J., Ungureanu, C., Aarts, R., Leijten, F., and Arends, J. Using photoplethysmography in heart rate monitoring of patients with epilepsy. *Epilepsy & Behavior 45* (2015), 142–145.

[181] van Andel, J., Ungureanu, C., Arends, J., Tan, F., Van Dijk, J., Petkov, G., Kalitzin, S., Gutter, T., de Weerd, A., Vledder, B., et al. Multimodal, automated detection of nocturnal motor seizures at home: Is a reliable seizure detector feasible? *Epilepsia open 2*, 4 (2017), 424–431.

[182] Van de Vel, A. *Epileptic seizure detection in children and young adults in their home replacement environment.* PhD thesis, University of Antwerp, 2016.

[183] Van de Vel, A., Cuppens, K., Bonroy, B., Milosevic, M., Jansen, K., Van Huffel, S., Vanrumste, B., Cras, P., Lagae, L., and Ceulemans, B. Non-eeg seizure detection systems and potential sudep prevention: State of the art: Review and update. *Seizure 41* (2016), 141–153.

[184] 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.

[185] Van de Vel, A., Cuppens, K., Bonroy, B., Milosevic, M., Van Huffel, S., Vanrumste, B., Lagae, L., and Ceulemans, B. Long-term home monitoring of hypermotor seizures by patient-worn accelerometers. *Epilepsy & Behavior 26*, 1 (2013), 118–125.

[186] Van de Vel, A., Smets, K., Wouters, K., and Ceulemans, B. Automated non-eeg based seizure detection: Do users have a say? *Epilepsy & Behavior 62* (2016), 121–128.

[187] Van de Vel, A., Verhaert, K., and Ceulemans, B. Critical evaluation of four different seizure detection systems tested on one patient with focal and generalized tonic and clonic seizures. *Epilepsy & Behavior 37* (2014), 91–94.

[188] van Elmpt, W., Wouter, J., Nijsen, T., Griep, P., and Arends, J. A model of heart rate changes to detect seizures in severe epilepsy. *Seizure 15*, 6 (2006), 366–375.

[189] van Mierlo, P., Papadopoulou, M., Carrette, E., Boon, P., Vandenberghe, S., Vonck, K., and Marinazzo, D. Functional brain connectivity from eeg in epilepsy: Seizure prediction and epileptogenic focus localization. *Progress in neurobiology 121* (2014), 19–35.

[190] Van Paesschen, W. The future of seizure detection. *The Lancet Neurology 17*, 3 (2018), 200–202.

[191] Vandecasteele, K., De Cooman, T., Gu, Y., Cleeren, E., Claes, K., Van Paesschen, W., Van Huffel, S., and Hunyadi, B. Automated epileptic seizure detection based on wearable ecg and ppg in a hospital environment. *Sensors 17*, 10 (2017), 2338.

[192] 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.

[193] Vapnik, V. *The nature of statistical learning theory*. Springer science & business media, 2013.

[194] Varon, C. *Mining the ECG: Algorithms and Applications*. PhD thesis, KU Leuven, 2015.

[195] 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 *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE* (2014), IEEE, pp. 3789–3790.

[196] 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 62*, 9 (2015), 2269–2278.

[197] 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.

[198] Varon, C., Jansen, K., Lagae, L., and Van Huffel, S. Can ecg monitoring identify seizures? *Journal of electrocardiology 48*, 6 (2015), 1069–1074.

[199] Varon, C., Testelmans, D., Buyse, B., Suykens, J. A., and Van Huffel, S. Robust artefact detection in long-term ECG recordings based on autocorrelation function similarity and percentile analysis. In *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE* (2012), IEEE, pp. 3151–3154.

[200] Vollmer, M. Robust Detection of Heart Beats using Dynamic Thresholds and Moving Windows. In *Proc. of the 41st Annual Computing in Cardiology (CinC2014)* (2014), pp. 569–572.

[201] Vonck, K., and Boon, P. Epilepsy: Closing the loop for patients with epilepsy. *Nature Reviews Neurology 11*, 5 (2015), 252–254.

[202] Welch, J., Ford, P., Teplick, R., and Rubsamen, R. The massachusetts general hospital-marquette foundation hemodynamic and electrocardiographic database–comprehensive collection of critical care waveforms. *Clinical Monitoring 7*, 1 (1991), 96–97.

[203] Wikilectures. Unipolar and bipolar connection. `https://www.wikilectures.eu/w/Unipolar_and_bipolar_connection`. Accessed: 2018-07-24.

[204] Wilder, R. The effects of ketonemia on the course of epilepsy. *Mayo Clin. Proc. 2* (1921), 307–308.

[205] Williams, C. K., and Seeger, M. Using the nyström method to speed up kernel machines. In *Advances in neural information processing systems* (2001), pp. 682–688.

[206] Yadollahpour, A., and Jalilifar, M. Seizure prediction methods: a review of the current predicting techniques. *Biomedical and Pharmacology Journal 7*, 1 (2015), 153–162.

[207] Yang, J., Yan, R., and Hauptmann, A. G. Adapting svm classifiers to data with shifted distributions. In *Seventh IEEE International Conference on Data Mining Workshops (ICDMW 2007)* (2007), IEEE, pp. 69–76.

[208] Yeh, Y.-C., and Wang, W.-J. QRS complexes detection for ECG signal: The Difference Operation Method. *Computer methods and programs in biomedicine 91*, 3 (2008), 245–254.

[209] Zavar, M., Rahati, S., Akbarzadeh-T, M.-R., and Ghasemifard, H. Evolutionary model selection in a wavelet-based support vector machine for automated seizure detection. *Expert Systems with Applications 38*, 9 (2011), 10751–10758.

[210] Zibrandtsen, I., Kidmose, P., Christensen, C. B., and Kjaer, T. Ear-eeg detects ictal and interictal abnormalities in focal and generalized epilepsy–a comparison with scalp eeg monitoring. *Clinical Neurophysiology 128*, 12 (2017), 2454–2461.

[211] Zibrandtsen, I., Kidmose, P., and Kjaer, T. Detection of generalized tonic-clonic seizures from ear-eeg based on emg analysis. *Seizure 59* (2018), 54–59.

[212] 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

Thomas De Cooman was born in Dendermonde, Belgium on September 17, 1989. He received the B.Sc. degree in computer science from the UGent, Gent, Belgium, in 2013. His master thesis was on "*Seizure detection based on ECG changes in patients with epilepsy*" under supervision of Prof. dr. ir. Stefaan Vandenberghe and Prof. dr. Christian Vanhove.

From October 2013, he 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. dr. ir. Sabine Van Huffel and Prof. dr. Lieven Lagae on the topic of automated seizure detection in a home environment. His strongly interdisciplinary research was in collaboration with both national partners (UZ Leuven, UAntwerp & Pulderbos revalidation center) and international partners (DTU, Zealand University Hospital, Academic Centre for Epileptology Kempenhaeghe center and Technological University Eindhoven). In 2014, his submission for the 2014 Computing in Cardiology Challenge was awarded with the first prize during the second phase and the third prize in the final phase of the competition.

List of publications

Papers in international journals

1. van Westrhenen A., **De Cooman T.**, Lazeron R.H.C., Van Huffel S., Thijs R.D., "Ictal autonomic changes as a tool for seizure detection", Clinical Autonomic Research, accepted for publication in Sep. 2018 *(Invited contribution)*.

2. **De Cooman T.**, Varon C., Van de Vel A., Jansen K., Ceulemans B., Lagae L., Van Huffel S., "Adaptive Nocturnal Seizure Detection Using Heart Rate and Low-Complexity Novelty Detection", Seizure, vol. 59, Jul. 2018, pp. 48-53.

3. **De Cooman T.**, Kjaer T., Van Huffel S., Sorensen H., "Adaptive Heart Rate-Based Epileptic Seizure Detection Using Real-Time User Feedback", Physiological Measurement, Focus on wearables and mHealth in mental health and neurological disorder, vol. 39, Jan. 2018, 014005 p.

4. Vandecasteele K., **De Cooman T.**, Gu Y., Cleeren E., Claes K., Van Paesschen W., Van Huffel S., Hunyadi B., "Automated Epileptic Seizure Detection Based on Wearable ECG and PPG in a Hospital Environment", Sensors, Special Issue on Wearable and Ambient Sensors for Healthcare and Wellness Applications, vol. 17, no. 10, Oct. 2017, pp. 1 - 12.

5. **De Cooman T.**, Varon C., Hunyadi B., Van Paesschen W., Lagae L., Van Huffel S., "Online Automated Seizure Detection In Temporal Lobe Epilepsy Patients Using Single-Lead ECG", International Journal Of Neural Systems, vol. 27, No. 7, Mar. 2017, pp. 1750022.

6. **De Cooman T.**, Goovaerts G., Varon C., Widjaja D., Willemen T., Van Huffel S., "Heart Beat Detection in Multimodal Data Using Automatic Relevant Signal Detection", Physiological Measurement, Focus issue on

robust detection of heart beats in multimodal data, vol. 36, no. 8, Aug. 2015, pp. 1691-1704 *(Invited contribution, Winner of phase II and third place in phase III of the Physionet/CinC challenge 2014.)*

Papers submitted to international journals

7. **De Cooman T.**, Vandecasteele K., Varon C., Hunyadi B., Cleeren E., ´ Van Paesschen W., Van Huffel S., "Personalizing Heart Rate Based Seizure Detection Using Supervised SVM Transfer Learning", IEEE Journal of Biomedical and Health Informatics, submitted in September 2018.

Submitted patents

8. **De Cooman T.**, inventor; KU Leuven, assignee. Identifying Seizures Using Heart Rate Increase. GB1607646.5, May 03 2016.

Papers in proceedings of international conferences

9. **De Cooman T.**, Varon C., Van Paesschen W., Van Huffel S., "Supervised Transfer Learning for Personalized Heart Rate Based Epileptic Seizure Detection", in Proceedings of the 45th Annual Computing in Cardiology Conference (Cinc 2018), Maastricht, The Netherlands, Sep. 2018. *In press.*

10. **De Cooman T.**, Varon C., Van de Vel A., Ceulemans B., Lagae L., Van Huffel S., "Comparison and Combination of Electrocardiogram, Electromyogram and Accelerometry for Tonic-Clonic Seizure Detection in Children", in Proceedings of the 2018 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI) (BHI 2018), Las Vegas, US, Mar. 2018, pp. 438-441.

11. **De Cooman T.**, Varon C., Van de Vel A., Ceulemans B., Lagae L., Van Huffel S., "Semi-Supervised One-Class Transfer Learning For Heart Rate Based Epileptic Seizure Detection", in Proceedings of the 44th Annual Computing In Cardiology Conference (CinC 2017), Rennes, France, Sep. 2017, pp.1-4.

12. **De Cooman T.**, Van de Vel A., Ceulemans B., Lagae L., Van Paesschen W., Vanrumste B., Van Huffel S., "Estimation of the Maximal Heart

Rate to Improve Online Tonic-Clonic Seizure Detection using ECG", in Proceedings of the 42nd Annual Computing in Cardiology (CinC2015), Nice,France, Sep. 2015, pp. 977-980.

13. **De Cooman T.**, Van de Vel A., Ceulemans B., Lagae L., Vanrumste B., Van Huffel S., "Online Detection of Tonic-Clonic Seizures in Pediatric Patients using ECG and Low-Complexity Incremental Novelty Detection", in Proceedings of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society of the IEEE (IEEE EMBS), Milan, Italy, Aug. 2015, pp. 5597-5600.

14. **De Cooman T.**, Goovaerts G., Varon C., Widjaja D., Van Huffel S., "Heart Beat Detection in Multimodal Data Using Signal Recognition and Beat Location Estimation", in Proceedings of the 41st Annual Computing in Cardiology (CinC2014), Cambridge, USA, Sep. 2014, pp. 257-260. *(Winner of phase II and third place in phase III of the Physionet/CinC challenge 2014.)*

15. **De Cooman T.**, Carrette E., Boon P., Meurs A., Van Huffel S., "Online Seizure Detection in Adults with Temporal Lobe Epilepsy Using Single-Lead ECG", in Proceedings of the 22nd European Signal Processing Conference (EUSIPCO 2014), Lisbon, Portugal, Sep. 2014, pp. 1532-1536.

Abstract in proceedings of (inter)national conferences

16. Vandecasteele K., **De Cooman T.**, Gu Y., Vandeput S., Cleeren E., Claes K., Dan J., Van Paesschen W., Van Huffel S., Hunyadi B., "Automated Multimodal Epileptic Seizure Detection Using EEG and ECG", In proceedings of the American Epilepsy Society Conference 2018, New Orleans, United States of America, Dec. 2018.

17. **De Cooman T.**, Varon C., Van Paesschen W., Van Huffel S., "From a Patient-Independent to a Patient-Specific Seizure Detection Algorithm Using Limited Heart Rate Data", In proceedings of the 1st International Congress on Mobile Health Devices and Seizure Detection in Epilepsy, Copenhagen, Denmark, July, 2017.

18. **De Cooman T.**, Van de Vel A., Varon C., Ceulemans B., Lagae L., Van Huffel S., "Comparison of Tonic-Clonic Seizure Detection Using Heart Rate, Accelerometer and EMG", In proceedings of the 15th Belgian Day on Biomedical Engineering, Brussels, Belgium, Dec. 2016.

19. Deviaene M., **De Cooman T.**, Lagae L., Van de Vel A., Ceulemans B., Van Huffel S., "Incremental Learning for Online Epileptic Seizure Detection Using Heart Rate", In proceedings of the 15th Belgian Day on Biomedical Engineering, Brussels, Belgium, Dec. 2016.

20. Vandecasteele K., **De Cooman T.**, Lagae L., Van de Vel A., Ceulemans B., Van Huffel S., "Online Epileptic Seizure Detection at Home: A Multimodal Approach", In proceedings of the 15th Belgian Day on Biomedical Engineering, Brussels, Belgium, Dec. 2016.

21. **De Cooman T.**, Van de Vel A., Ceulemans B., Lagae L., Van Huffel S.,"Automated Epileptic Seizure Detection Using Accelerometry, Heart Rate And Electromyogram", In proceedings of the American Epilepsy Society Conference 2016, Houston, United States of America, Dec. 2016.

22. **De Cooman T.**, Varon C., Van de Vel A., Ceulemans B., Lagae L, Vanrumste B., Van Huffel S., "Detection of Nocturnal Convulsive Seizures Using an Automatically Updated Heart Rate-Based Algorithm", In proceedings of the 14th International Child Neurology Congress, Amsterdam, The Netherlands, Mar. 2016.

23. **De Cooman T.**, Van de Vel A., Ceulemans B., Lagae L., Vanrumste B., Van Huffel S., "Online Detection of Tonic-Clonic Seizures in Pediatric Patients Using ECG and Low-Complexity Incremental Novelty Detection", in proceedings of the 14th Belgian Day on Biomedical Engineering, Brussels, Belgium, Nov. 2015.

24. **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", In proceedings of the 5th Dutch Bio-Medical Engineering Conference 2015, Egmond aan Zee, The Netherlands, Jan. 2015.

25. **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", In proceedings of the 13th Belgian Day on Biomedical Engineering, Brussels, Belgium, Nov. 2014.

26. **De Cooman T.**, Jansen K., Lagae L., Carrette E., Van Huffel S., "Online Seizure Detection Using ECG in Patients with Temporal Lobe Epilepsy", in proceedings of the 12th Belgian Day on Biomedical Engineering, Brussels, Belgium, Dec. 2013.



FACULTY OF ENGINEERING SCIENCE
DEPARTMENT OF ELECTRICAL ENGINEERING
STADIUS CENTER FOR DYNAMICAL SYSTEMS, SIGNAL PROCESSING AND DATA ANALYTICS
Kasteelpark Arenberg 10 box 2446
B-3001 Leuven
thomas.decooman@esat.kuleuven.be