

KATHOLIEKE UNIVERSITEIT
LEUVEN

Arenberg Doctoral School of Science, Engineering & Technology
Faculteit ingenieurswetenschappen
Department elektrotechniek

Detection of Epileptic Seizures based on Video and Accelerometer Recordings

Promotoren:
Prof. dr. ir. B. Vanrumste
Prof. dr. ir. S. Van Huffel

Proefschrift voorgedragen tot
het behalen van het doctoraat
in de ingenieurswetenschappen

door
Kris CUPPENS

December 2012

Detection of Epileptic Seizures based on Video and Accelerometer Recordings

Jury:
Prof. dr. ir. P. van Houtte, voorzitter
Prof. dr. ir. B. Vanrumste, promotor
Prof. dr. ir. S. Van Huffel, promotor
Prof. dr. B. Ceulemans
Prof. dr. L. Lagae
Prof. dr. ir. T. Tuytelaars
Prof. dr. ir. L. De Strycker
(KaHo Sint-Lieven)
Prof. dr. ir. H. Aghajan
(Stanford University)
Prof. dr. ir. H. Sørensen
(Technical University of Denmark)

Proefschrift voorgedragen tot
het behalen van het doctoraat
in de ingenieurswetenschappen

door
Kris CUPPENS

December 2012

© Katholieke Universiteit Leuven – Faculteit Ingenieurswetenschappen
Kasteelpark Arenberg 10, B-3001 Leuven (Belgium)

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

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

Legal depot number D/2012/7515/146
ISBN 978-94-6018-614-1

Voorwoord

Het maken van een doctoraat stopt niet bij de zoektocht naar het meest optimale algoritme of optimale meetopstelling, maar het is ook een zoektocht naar een verdere ontplooiing van jezelf. En in het bijzonder het laatste jaar van mijn doctoraat was hierin op verschillende vlakken leerrijk. Het begon op 2 januari met een vlucht naar de VS voor een verblijf aan Stanford, en eindigt met de publieke verdediging van mijn doctoraat op 21 december (symbolisch of niet, volgens sommigen het begin van een nieuw tijdperk ☺)

En hoewel de weg soms al eens wat hobbeliger verliep dan op andere momenten, ben ik blij dat ik hem niet alleen moest bewandelen. Daarom zou ik graag een aantal personen willen bedanken.

Ik zou willen beginnen met mijn dank uit te spreken aan mijn promotoren, prof. Bart Vanrumste en prof. Sabine Van Huffel. Bart, ik herinner mij het moment nog zeer goed dat je me een vacature voorstelde tijdens de proclamatie om als onderzoeker te starten aan een onderzoeksgroep in Geel. Ik had op dat moment zelfs geen idee waar Geel lag, maar kijk, ondertussen woon ik er al meer dan een jaar ☺. Bedankt voor het vertrouwen en voor de constructieve uitwisseling van ideeën die we meermaals gehad hebben. Sabine, in de zomer van 2009 stelde je voor om meer opgenomen te worden binnen het onderzoek in Leuven. Onder andere dankzij de feestjes bij je thuis leerde ik een fijne groep nieuwe collega's kennen. Ik wil jullie beiden ook bedanken voor de begrijpende en opbouwende begeleiding.

Verder wil ik mijn begeleidingscommissie danken voor de constructieve kritiek en input die ik van jullie mocht ontvangen. Prof. Berten Ceulemans en prof. Lieven Lagae om de klinische relevantie van het onderzoek niet uit het oog te verliezen. Prof. Tinne Tuytelaars voor de steeds heldere analyses en to-the-point hulp. Prof. Lieven De Strycker voor de feedback op mijn thesis en door iets verder van het topic af te staan, goede verduidelijkingen voorstelde.

Ook wil ik het Agentschap voor Innovatie door Wetenschap en Technologie (IWT) bedanken voor de financiering van mijn doctoraatsbeurs van strategisch onderzoek.

Wat me opviel tijdens de minder aangename tijden is dat deze dikwijls gepaard gaan met ook fijne momenten en dat dan toch vooral het positieve bijblijft. Hierbij wil ik dan ook iedereen bedanken voor de spontane steunbetuigingen via smsjes of mails tijdens mijn soms lange avonden van schrijfwerk aan mijn thesis. Dit heeft me de nodige inspiratie bezorgd, ook zeker voor het schrijven van dit voorwoord! ☺ Bedankt ook voor het gezelschap van de mede-nachtraven, die het avondwerk een stuk gezelliger maakte.

In the beginning of this year I had the opportunity to join the AIR lab research group at Stanford University for more than 2 months. I would like to express my gratitude to prof. Hamid Aghajan for the confidence he had in me and for the fruitful collaboration. I hope our collaboration will continue in the future.

Ik wil ook mijn Mobilabcollega's bedanken, de collega's van de P224 en mijn collega-docenten. Bert, die na meer dan 5 jaar naast of tegenover mij te zitten mijn flauwe grapjes nog steeds verdraagt (maar eigenlijk stiekem dezelfde humor heeft ☺). Eveline, die naast een verre tante, ook nog eens dezelfde rare hersenkronkels maakt als ik. Nooit gedacht dat ik dat zou tegenkomen! ☺ Inge, het is een plezier om tegenover zoveel optimisme te zitten, en je haalt meer dan eens mijn non-conformistische zelf in mij naar boven. Veerle, voor alle leuke momenten op het Studio 100 eiland en die van erna! Glen, Peter, Marc, voor de fijne samenwerking rond beeld- en signaalverwerkingen en je thesis die ik veelvuldig heb mogen raadplegen. Mijn nieuwe eilandgenoten Lore en Roy, maar uiteraard ook een enorm bedankje aan alle andere Mobilabcollega's voor de fijne samenwerking op de vloer, de Mobilabweekendjes of quizzen die we wonnen! En nog een speciaal dankwoord aan Louis voor je toewijding voor onze ploeg.

I also would like to thank my colleagues of BioMed. Although I sometimes was not there for some weeks, I always felt welcome. I had the honour of joining 'the girls office', which was also the place to be for many birthday celebrations. A special thanks to my office mates Anca, Aileen, Rosy, Nicolas, Maria and Diana. I enjoyed the great times on the pizza parties, barbecues or SISTA-weekends, so a big thanks to everyone! An additional thank you to the colleagues that joined me on the conference this summer for the funny moments, although thanks to Alex and Carolina I had to pay the high prize of not being able to listen to Lionel Richie anymore...☺ And to Katrien for the numerous times I checked your PhD thesis during writing.

Furthermore, I want to thank the colleagues that I collaborated with in the TBM-project. Ann-Sofie, Milica and Anouk, it was a pleasure to work with you. And all the people from Pulderbos who made the collection and labeling of the data possible. I couldn't have made this thesis or even do the research without your input. And not to forget the people I collaborated with at the UZ Leuven to acquire video data.

Aan mijn vrienden van thuis, van bij Maarten en van de badminton. Voor de momenten van de afgelopen maanden dat jullie mij niet gezien hebben, of dat ik er te laat was (maar dat laatste is geen fenomeen van de alleen de afgelopen maanden natuurlijk ☺), bedankt voor de momenten van ontspanning en afleiding tijdens de drukke momenten! En btw, ik zal stiekem altijd wel Limburger blijven hoor, al willen de geheime uitbreidingsplannen nog niet zo goed lukken ☺.

Aan mijn broer en aan mijn ouders, voor jullie onvoorwaardelijke steun in (bijna) alles wat ik doe, en de hulp en aanmoedigingen die ik van jullie mocht ontvangen. Ik wil de fierheid over mijn werk vandaag dan ook graag met jullie delen.

Korte inhoud

Epilepsie is één van de meest voorkomende neurologische aandoeningen, vooral bij kinderen. En hoewel het merendeel van de patiënten behandeld kan worden door middel van medicatie of een operatie (70%-75%), kan een aanzienlijke groep van patiënten niet behandeld worden. Voor deze laatste groep is het aangewezen om de evolutie van de aandoening op te volgen. Dit kan door middel van een langdurige automatische monitoring, die een objectieve maat geeft van het aantal aanvallen dat de patiënt heeft, bijvoorbeeld tijdens de nacht. Anderzijds is er 's nachts een verminderde sociale controle en kunnen de ouders of zorgverleners sommige aanvallen missen. Bij hevige aanvallen is het soms echter noodzakelijk om gevaarlijke situaties tijdens of na de aanval te voorkomen (bijvoorbeeld het gevaar op verstikking door overgeven of een positie die het ademen bemoeilijkt, of gevaar op verwonding tijdens de hevige bewegingen), en om het kind gerust te stellen na de aanval.

De huidige gouden standaard voor het monitoren van epilepsie maakt gebruik van de combinatie van video en EEG. De opname van EEG wordt echter bemoeilijkt door de elektrodes die bevestigd moeten worden op de scalp. Dit is oncomfortabel voor de patiënt om mee te slapen en bovendien is het bevestigen van de elektrodes arbeidsintensief waardoor een langdurige thuismonitoring onmogelijk wordt.

Om dit mogelijk te maken onderzoeken we in deze thesis of het gebruik van sensoren die minder ingrijpen op het comfort van de patiënt, namelijk accelerometers bevestigd in pols- en enkelbandjes en video, gebruikt kunnen worden om epileptische aanvallen te detecteren waarbij er een motorische component aanwezig is. Deze thesis focust op twee typen aanvallen bij kinderen, namelijk de hypermotore en de myoclone aanvallen.

Na een voorbewerking van de data om te komen tot individuele bewegingsevents, worden eerst de individuele modaliteiten gebruikt om aanvallen te detecteren. Hiervoor worden state-of-the-art methoden aangewend die reeds toegepast worden

in andere domeinen. Zo wordt er voor de detectie op basis van accelerometerdata de vergelijking gemaakt tussen een classificatie op basis van support vector machines en het detecteren van uitschieters. De videodetectie maakt dan weer gebruik van kenmerken rond spatio-temporele punten om bewegingspatronen te classificeren. Deze methode wordt bijvoorbeeld al gebruikt om reële acties in films te herkennen. Een extra moeilijkheid in de aanvalsdetectie uit videobeelden in vergelijking met de literatuur, is dat het onderzoek in deze thesis de normale slaap van de patiënt zo weinig mogelijk probeert te verstoren. Daarom wordt er geen gebruik gemaakt van markers of andere visuele hulp, en kan de patiënt onder de lakens blijven liggen.

In een laatste deel van deze thesis wordt onderzocht of een geïntegreerde aanpak een beter resultaat kan opleveren. Hier worden de video- en accelerometerdata samengevoegd, zowel op het kenmerkniveau, als op het niveau van de uitvoer van een algoritme op de individuele modaliteiten.

Het uiteindelijke doel van het onderzoek is om te komen tot algoritmes die bruikbaar zijn in de praktijk. Dit heeft tot slot ook geleid tot een thuismonitoringssysteem dat zich op dit moment in een testfase bevindt, waarbij patiënten telkens gedurende een maand gemonitord worden. Hopelijk kan deze thesis dan ook bijdragen tot een verbetering van de opvolging en bijgevolg de levenskwaliteit van de patiënten.

Abstract

Epilepsy is one of the most common neurological diseases, especially in children. And although the majority of patients can be treated through medication or surgery (70%-75%), a significant group of patients cannot be treated. For this latter group of patients it is advisable to follow the evolution of the disease. This can be done through a long-term automatic monitoring, which gives an objective measure of the number of seizures that the patient has, for example during the night. On the other hand, there is a reduced social control overnight and the parents or caregivers can miss some seizures. In severe seizures, it is sometimes necessary, however, to avoid dangerous situations during or after the seizure (e.g. the danger of suffocation caused by vomiting or a position that obstructs breathing, or the risk of injury during violent movements), and to comfort the child after the seizure.

The current gold standard for the monitoring of epilepsy makes use of the combination of video and EEG. The recording of EEG is however hampered by the electrodes that must be attached to the scalp. This makes it uncomfortable for the patient to sleep with, and in addition, the attachment of the electrodes is labor-intensive which makes long-term home monitoring not feasible.

In order to make this possible, we examine in this thesis whether the use of sensors that are less intrusive, namely accelerometers attached with wrist and ankle straps, and video, may be used to detect epileptic seizures with a motor component. This thesis focuses on two types of seizures in children, namely hypermotor and myoclonic seizures.

After a preprocessing of the data to create individual movement events, first the individual modalities are used for detecting seizures. For this purpose, state-of-the-art methods are used which are already applied in other fields. For example, for the detection based on accelerometer data, the comparison is made between a classification based on support vector machines and an outlier detection. The video detection on its turn uses features derived in the neighborhood of spatio-temporal

interest points to classify movement patterns. This method is for example already used to recognize real actions in movies. An additional difficulty in the seizure detection using video, is that in comparison with the literature, the setup in this thesis attempts to disrupt the normal sleep of the patient as little as possible. To this end no markers or other visual aids are used, and the patient may keep using sheets.

In a last part of this thesis it is examined whether an integrated approach can yield a better result. Here, the video and accelerometer data are joined, both at feature and at classifier level.

The ultimate goal of the research is to create algorithms that are useful in practice. This has led to a home monitoring system that is currently in a test phase, in which patients are monitored for one month each. Hopefully this thesis can therefore contribute to improve the monitoring and thus the quality of life of patients.

Nomenclature

Variables and Symbols

$a, b, ...$	scalars
$\mathbf{a}, \mathbf{b}, ...$	vectors
$\mathbf{A}, \mathbf{B}, ...$	matrices
$\mathcal{A}, \mathcal{B}, ...$	tensors
$\mathbf{A}^T$	transpose of the matrix $\mathbf{A}$
a_{ij}	element of matrix $\mathbf{A}$ at position (i, j)
$\mathcal{A}_X$	partial derivative of $\mathcal{A}$ with respect to variable x
$\sum$	sum
$\min_x f(x)$	minimization of function $f(x)$ over x, minimal function value returned
$\operatorname{argmin}_x f(x)$	minimization of function $f(x)$ over x, optimal value of x returned
$a^{(k)}$	a at the k^{th} iteration

Metrics

fps	frames per second
g	gravitational force
GB	gigabyte
Hz	Hertz
kB	kilobyte
kΩ	kiloohm
MB	megabyte
mV	millivolt
nF	nanofarad
s	second

Acronyms

ACM	accelerometer
ADL	Activities of Daily Living
AED	Antiepileptic Drugs
AIC	Akaike Information Criterion
BIC	Bayesian Information Criterion
CDF	Cumulative Distribution Function
CDM	Change Detection Map
CP	Cerebral Palsy
CT	Computed Tomography
DC	Direct Current
ECG	Electrocardiogram
EDA	Electrodermal Activity
EEG	Electroencephalogram
EMG	Electromyogram
EOG	Electrooculogram
FN	False Negative
FP	False Positive
FPGA	Field-programmable Gate Array
GMM	Gaussian Mixture Model
GPU	Graphics Processing Unit
HoF	Histogram of optical Flow
HoG	Histogram of oriented Gradients

HRV	Heart Rate Variability
ILAE	International League Against Epilepsy
LDA	Linear Discriminant Analysis
MEMS	Microelectromechanical Systems
MML	minimum message length
MNMR	Movement versus Non-Movement Ratio
MPEG	Moving Picture Experts Group
MRI	Magnetic Resonance Imaging
NICU	Neonatal Intensive Care Unit
NPV	Negative Predictive Value
PDF	Probability Density Function
PNES	Psychogenic Non-epileptic Seizures
PPG	Photoplethysmography
PPV	Positive Predictive Value
PSD	Power Spectral Density
PSG	Polysomnography
RBF	Radial Basis Function
RF	Radio Frequency
rms	root mean square
ROC	Receiver Operating Characteristic
SIFT	scale-invariant feature transform
SLE	Systemic Lupus Erythematosus
SMOTE	Synthetic Minority Over-sampling TEchnique
std	standard deviation
STIP	Spatio-temporal Interest Point
SUDEP	Sudden Unexplained Death In Epilepsy
SVM	Support Vector Machine
TN	True Negative
TP	True Positive
VNS	Vagus Nerve Stimulation

Contents

Contents **xiii**

1 Introduction **1**

1.1 What is epilepsy? 1

1.1.1 Subdivisions 3

1.1.2 Methods of treatment 5

1.1.3 Diagnostic methods 6

1.1.4 Gold standard 7

1.2 Home monitoring 7

1.2.1 Why home monitoring? 7

1.2.2 Challenges 8

1.2.3 Current state-of-the-art 8

1.3 Automatic seizure detection 9

1.3.1 Alternative methods for detection 10

1.3.2 Accelerometry 14

1.3.3 Video 16

1.4 Measures 17

1.5 Motivation of research 19

1.6 Challenges and objectives 19
1.7 Chapter-by-chapter overview 20
1.8 Personal contributions 22

2 Data acquisition **25**
2.1 Introduction 25
2.1.1 Requirements 26
2.1.2 Commercially available systems 27
2.2 Simulation data 29
2.2.1 Video data 29
2.2.2 ACM data 31
2.3 Real data 31
2.3.1 Location 1: UZ Leuven 31
2.3.2 Location 2: Rehabilitation center Pulderbos 34
2.4 Home monitoring data 38

3 Preprocessing and movement detection - Accelerometers **41**
3.1 Introduction 41
3.2 ACM based detection 42
3.3 Conclusion 44

4 Preprocessing and movement detection - Video **47**
4.1 Introduction 47
4.2 Movement detection 51
4.2.1 Video acquisition 51
4.2.2 Optical flow algorithm 51
4.2.3 Calculating $R(t)$ 54

4.2.4 Results on choosing acquisition parameters 58
4.2.5 Results on identifying τ and n_{vec} 62
4.2.6 Discussion 65
4.3 Detection of nocturnal breathing 66
4.3.1 Introduction 66
4.3.2 Method 67
4.3.3 Discussion 69
4.4 Conclusion 70

5 Unimodal Seizure detection - Accelerometers 71

5.1 Introduction 72
5.2 Thresholding 75
5.2.1 Detection algorithm 75
5.2.2 Training and validation 77
5.2.3 Results 78
5.2.4 Discussion 78
5.2.5 Conclusion 79
5.3 Support Vector Machines 79
5.3.1 Feature extraction and selection 80
5.3.2 Test outline 85
5.3.3 Results 87
5.3.4 Discussion 89
5.3.5 Conclusions 93
5.4 Outlier detection by estimating the Probability Density Function . 93
5.4.1 Model estimation and hyperparameter tuning 94
5.4.2 Results 98

5.4.3 Discussion 100
5.4.4 Conclusion 102
5.5 Comparison between methods and conclusion 102

6 Unimodal Seizure detection - Video 105
6.1 Introduction 105
6.2 Clustering flow vectors 108
6.2.1 Features used for clustering 109
6.2.2 Clustering methods 110
6.2.3 Determine number of clusters 114
6.2.4 Comparison of clustering methods 117
6.2.5 Tracking 119
6.2.6 Feature extraction 119
6.2.7 Classification 120
6.2.8 Results 121
6.3 Spatio-Temporal Interest Points (STIP) 123
6.3.1 Spatio-temporal interest points (STIP) and feature extraction 123
6.3.2 Classification 125
6.3.3 Results 125
6.3.4 Discussion 130
6.4 Comparison between both methods and conclusion 131

7 Multimodal seizure detection 135
7.1 Introduction 135
7.2 Multi-modal integration 136
7.3 Results 138
7.4 Discussion 138

8 Conclusion **141**

8.1 Concluding remarks 141

8.2 Future Research 144

8.2.1 Research directions 144

8.2.2 Integrated detection in real setting 145

A Additional results **147**

Bibliography **151**

Curriculum vitae **171**

Publications by the author **173**

Chapter 1

Introduction

This chapter aims at introducing epilepsy and the different ways of monitoring and classifying this neurological disorder. Furthermore, a general overview is given on the recognition of motion patterns on the modalities we investigate. Section 1.1 gives a general introduction stating general facts and figures and elaborates more on the patients we focus on. Also, more information is given on the subdivisions that can be distinguished in epilepsy, and different diagnostic and treating methods that are applied in clinical practice today. Furthermore, in Section 1.3 we look ahead at promising modalities that are investigated which can aid the clinicians in the nearby future to make more profound decisions on the diagnosis of epilepsy or help the patients and their environment in monitoring the seizures at home. In Section 1.3.2 and Section 1.3.3, we give an introduction on the current state-of-the-art in accelerometry and video, as these modalities are further investigated in this thesis. Section 1.4 introduces measures that are used in the next chapters for evaluating the different algorithms. The goals and challenges of the thesis are described in Section 1.5 and Section 1.6, respectively. Section 1.7 gives a chapter-by-chapter overview of this thesis, and finally Section 1.8 ends with a summary of the personal contributions.

1.1 What is epilepsy?

People are considered to have epilepsy when they suffer from more than one seizure which is not related to a clear cause such as fever, a head injury or drug use. A

seizure is caused by a sudden brief excess surge of electrical activity in the brain[1]. Epilepsy is one of the most common neurological disorders which affects almost 1% of the world's population. 25% to 30% of epileptic patients can not be treated by medication or surgery, they suffer from so-called refractory epilepsy [30, 56]. Epileptic seizures mainly occur as paroxysmal events, which means that they occur at a sudden, unexpected timing. The frequency of the seizures varies from patient to patient. Some patients suffer from multiple seizures during the day and/or night, while others only have one a month or less.

To be able to track the progress of the disease and to alarm caregivers during a seizure, these patients should be monitored. The gold standard in epilepsy monitoring uses EEG-electrodes [28] attached to the scalp. However, these electrodes are difficult to attach, hamper the patient's sleep during the night, and therefore are prohibiting long term home monitoring.

The patient group we focus on are children as the prevalence of epilepsy in children is higher and parents want to know when their child has a seizure during the night. Furthermore during sleep the convulsions occur more or less in a controlled reproducible manner without too much interference of other noise sources such as voluntary movement. The type of epileptic seizures we focus on in this thesis are hypermotor seizures and myoclonic seizures.

Hypermotor seizures manifest themselves as violent, uncontrolled movements of the arms and legs, e.g. by making a pedalling movement [6, 153]. The movements can last from a couple of seconds to multiple minutes. Due to the possible heavy movements, the patients can injure themselves or even die [89]. Patients may suffer from confusion after a seizure, and even when they don't, they often recall the seizure as a 'strange feeling' and need comforting [171, 172].

Myoclonic seizures are small jerk-like movements that can occur as a single seizure or precede others like tonic seizures. A more challenging problem with these types of seizures is the occurrence of non-epileptic jerks like the hypnic jerk that is experienced before falling asleep, as this jerk-like movement does not denote abnormal neurological activity.

Due to the different nature of both seizures, the aim for detecting both types is different. Because the hypermotor seizures are violent and the potential danger for the patient is high, it is necessary to detect (almost) all seizures of this type. False detections are less of an issue. The detection of myoclonic jerks, on the other

[1]Information coming from the Epilepsy Parent Guide from the International League Against Epilepsy (ILAE).

hand, is not critical. For this type of seizures no alarm is needed, but the aim is to get an idea of the number of seizures that occurred.

Although it is relatively rare, sometimes, epilepsy can be fatal. When a patient suffering from epilepsy dies without a clear cause for his death (such as drowning or caused by a trauma), the term SUDEP (Sudden Unexpected Death in EPilepsy) is used. Most of the time it is presumed that the death occurs during seizures, however, this is not the case in all situations and this is not a requirement for the diagnosis of SUDEP.

1.1.1 Subdivisions

Epilepsy and epileptic seizures can be subdivided in different groups, dependent on which aspect of the manifestation is considered, for example based on their cause (seizure etiology), origin in the brain or clinical manifestation.

The classification of seizures or epilepsies in different groups is not static and work in progress, it can be modified with advancing clinical insight. In the latest report of the International League Against Epilepsy (ILAE) [14] on the terminology and organization of seizures and epilepsies, some changes are proposed to the seizure classification such as a revision of concepts, terminology and approaches for classifying seizures.

Syndromes and epilepsies grouped by specificity of diagnosis

Electroclinical syndromes are groups of epilepsy with similar specifications such as a typical age at onset, a similar genetic defect, clinical or electrical manifestation. They represent a global overview of the disease. For example the West-syndrome is a type of epilepsy that first occurs during infancy. The onset of the seizure is between the third and twelfth month. It manifests itself by the infantile spasms it causes. On a later age, this syndrome can for example evolve in a Lennox-Gastaut syndrome.

Syndromes can be arranged by age, in this case we distinguish five categories: neonatal period, infancy, childhood, adolescence - adult and variable age at onset.

Classification of epilepsy based on cause

The cause of epileptic seizures, the so-called seizure etiology, can depend on several factors, or a combination of more of them. In the literature the following causes can be found:

- Vascular: e.g. caused by a stroke, the loss of blood or high blood pressure (hypertension)
- Infectious: meningitis, encephalitis (acute inflammation of the brain e.g. Rasmussen syndrome)
- Traumatic (by a sudden physical injury)
- Autoimmune: e.g. SLE (Systemic lupus erythematosus, a systematic autoimmune disease) or vasculitis (inflammatory destruction of blood vessels)
- Metabolic/toxic: electrolyte imbalance, toxins
- Idiopathic (unknown cause)
- Neoplastic: abnormal tissue growth, resulting in malignant or benign tumors
- Structural: cortical malformations, prior stroke, other causes of CP (cerebral palsy, motor condition with origin in the brain)
- Syndrome/genetic

Classification of seizures based on origin in the brain

According to the latest report of the ILAE on the terminology of seizures and epilepsies [14], seizures can be classified in three groups, generalized seizures, focal seizures and unknown.

Seizures can originate in specific parts of the brain (so-called focal seizures), where the seizure starts in one hemisphere, or they can arise within and rapidly engage both hemispheres of the brain. However, the focal seizures may evolve into bilateral (both parts of the brain) convulsive seizures (e.g. tonic, clonic or tonic-clonic).

Depending on the affected area of the brain, the clinical manifestation may be different. If the seizure can affect mainly the motor areas, or cause hallucinations or an autonomic response (e.g. change of heart rhythm, breathing or body temperature).

The generalized seizures can be divided into 6 subcategories:

- Absence: impairment of consciousness
- Clonic: repetitive jerks, rhythmic muscle contractions
- Tonic: sudden stretching of muscles
- Tonic-Clonic: the seizure starts with a tonic phase that evolves in a clonic phase
- Atonic: brief loss of muscle tension
- Myoclonic: single jerk or twitch

Classification of seizures based on clinical manifestation

The clinical manifestation of the seizures are much related to the origin of the seizures in the brain. For example seizures that affect the motor cortex result in convulsions, whereas seizures that affect the occipital lobe, which is responsible for processing visual information, can cause for example hallucinations. The list on generalized seizures in the previous section gives also an overview of different clinical manifestations of motor seizures[2].

For other seizures, such as for hypermotor seizures (as described in Section 1.1), the epileptic activity is often not very clear based on the EEG, also complicated by the influence of the muscle activity because the electrical signals associated with muscle contractions are much larger compared to the electrical discharges in the brain and therefore is a large source for noise in the EEG.

1.1.2 Methods of treatment

In almost 70% of the patients with epilepsy the seizures can be controlled by using medication. A wide range of medication is available for seizure control. When the medication doesn't have a positive effect on the epilepsy, different types of antiepileptic drugs (AED) can be employed. However, the probability of an AED to be effective decreases with the number of different AEDs tested. Therefore, after trying out for example three different types of drugs, the neurologist may decide to use different treatment strategies. Epilepsies that are resistant to drug treatment

[2] *Motor seizures* is a generic term for indicating seizures associated with movement.

are called refractory epilepsies. Some patients are treated with a combination of different drugs to control the seizures.

In case of these refractory epilepsies, the neurologist can decide to consider epilepsy surgery. This is not suitable for all patients. The epileptic origin of - only focal - seizures in the brain has to be clearly defined. During the surgery the area in the brain causing the seizures is disconnected from the rest of the brain or removed. A more recent development is the implantation of an electrode in the brain to apply deep brain stimulation, by sending electrical pulses to the part of the brain that causes the seizures. This type of therapy is still in a clinical testing phase.

A novel treatment is the vagus nerve stimulation (VNS). The vagus nerve is a cranial nerve that mainly transports information from the body organs to the central nervous system. DeGiorgio *et al.* [60] showed in their doulbe-blind study on 195 patients that a stimulation of the vagus nerve can reduce the number of seizures by more than 50% in 35% of the patients. The median reduction of the seizures was 45%, and in 20% of the patients there was even a reduction of more than 75%. Lulic *et al.* [125] give an overview of the different clinical studies that are conducted to investigate the effect of VNS treatment on epileptic patients. Cyberonics already commercialized this therapy.

Another alternative method is the ketogenic diet. In this special diet, the consumption of fat is high and carbohydrates is low. The name ketogenic means that it produces ketones in the body. These ketones are formed when the body uses fat as an energy source. In this diet the body changes from burning carbohydrates, as they are only provided in very small amount, to using fat as fuel. A higher level of ketones in the body often lead to an improved seizure control, although it is not fully clear why this helps in reducing the frequency of seizures.

1.1.3 Diagnostic methods

Before epilepsy is diagnosed, the patients are typically monitored during a 24-hour video/EEG. The exact diagnosis is based on a wide range of information of the patients, such as medical history, blood tests, the EEG-monitoring or other brain imaging such as MRI or CT scans.

When no treatment seems to be successful, the patients can be monitored to keep track of the evolution of the disease and maybe alter the medication to get a reduction in the seizure frequency anyway, and to alarm if a serious seizure occurs that needs care during or after the seizure.

1.1.4 Gold standard

The Gold Standard for the diagnosis for epilepsy is video-electroencephalogram (EEG) monitoring. In the EEG monitoring, electrodes are placed on the scalp. It measures the potential differences that occur caused by the electrical discharges in the brain related to brain activity. Only the activity of thousands or millions of parallel firing neurons can be detected. The EEG is employed to detect seizures and transients that can denote epileptic activity in the brain. To provide a good conductivity between the scalp and the electrode, the skin is first lightly abraded, to remove the dead skin cells, and then a conductive paste is applied. The electrode is glued to the scalp. This method, however, is patient-demanding as the attached electrodes are uncomfortable to wear for long-term monitoring. Moreover, it is labor-intensive for the technical staff as it takes 20 to 40 minutes to glue all the electrodes on the scalp. Furthermore if the impedance between the electrode and the scalp is too high, no good acquisition of the EEG is possible. Hence, multiple weeks of home monitoring is not feasible. The number of electrodes depend on which montage is used. In high-density arrays using a cap, it can go up to 256 electrodes. The measurements that are recording in our study use 24 electrodes including the reference. As the gathered signals are weak, they have to be amplified by specialized hardware. Typical sampling rates are around 256 to 512 Hz.

1.2 Home monitoring

1.2.1 Why home monitoring?

In recent years, a number of trends are observed that are expected to increase in importance in the near future. First, there is a tendency in society of an aging population. The related geriatric ailments and the extended institutionalized care will increase the health care costs. If elderly could live in their familiar home environment for a longer period of time, this would not only increase their independence and their quality of life but also lower the medical costs by postponing institutionalized care. A prolonged stay in the home environment is possible by improving their (medical) safety by telemonitoring their vital signs or e.g. detect fall incidents.

But home monitoring is not only applicable for elderly, as is illustrated by this thesis, focusing on pediatric patients. Thanks to the continuing technological progress, solutions become available that permit technologies to become part of the everyday

life of the patients. Sensor technologies are used in consumer electronics, such as accelerometers in cell phones. Thanks to the internet and wireless communication, data can be transmitted easily from patient's home to the caregiver's office. This brings care closer to the patient and provides for a greater involvement of the patient in the treatment strategy. In this way, home monitoring can be an answer to follow up chronic diseases, have more objective information about the evolution of the disease, and it can allow a quicker response to a sudden critical situation of the patient.

1.2.2 Challenges

One of the key challenges in home monitoring is the acceptance of the technology by the patient. Acquiring the vital signs of the patient can be an extra burden to the patient next to coping with the disease. Especially if the patient has to attach different sensors to the body for for example monitoring the heart rate. After all, most of the time the data have to be acquired all day and every day. If the home monitoring system is not easy in use, the patient will stop using it after a while.

Next to the easy applicability, it is also important that the system is comfortable for the patient. Especially if it is used during sleep, the user should not be hampered. In both issues, being easy to use and the comfort, it is advisable that the sensors are contactless or at least wireless.

Also the performance of the system has to be acceptable. If a monitoring system is designed to detect risky situations such as fall incidents, or as in this thesis epileptic seizures, it becomes useless or even annoying for the patient if there are too many false positives.

1.2.3 Current state-of-the-art

The first applications of home monitoring are already available on the market, for example push button operated alarm systems for elderly to indicate fall incidents [146], such as the Philips Life Line, LifeFone or Medical Alert by Connect America. The fall detector by Tunstall and the Philips Life Line also offer an automatic fall detection based next to the more common alarm buttons. Recently, more research is conducted about the detection of fall using camera systems [59]. Also systems that monitor patients with heart problems are already available. Although some studies indicated that for the monitoring of heart failure in specific, telemonitoring can be an additional tool on top of the usual care, without replacing the latter

one [79]. And more studies are required to investigate the cost-efficacy. Other studies, however, show that remote monitoring in patients with heart failure leads to fewer admissions and shorter hospital stays [39].

A large study in the UK, the Whole System Demonstrator Programme, evaluated the implementation of home monitoring involving 6191 patients [164]. Half of the subjects suffered from either diabetes, heart failure or COPD[3]. This study showed a general reduction in emergency and elective admissions in the hospital with 20% and 14%, respectively. The mortality rate was reduced by 45% compared to the control group. This indicates that home monitoring has an added value to the health care, and leads to a reduction in hospitalization costs. Different telehealth devices were used in this study which varied among the different trial sites, although they all used a pulse oximeter for chronic obstructive pulmonary disease, a glucometer for diabetes, and weighing scales for heart failure [164].

1.3 Automatic seizure detection

There is a clear distinction between the diagnosis of epilepsy and the detection and monitoring of epileptic seizures. The former needs no exact information about the number and timing information of the seizure occurrence. A global view on the patient is more important incorporating for example the medical background as explained in Section 1.1.3. Seizure detection for alarming and monitoring purposes, however, does need this information.

The diagnosis of epilepsy is only possible in a clinical setting, due to the different aspects of the disorder that are investigated. Whereas seizure detection is, next to the detection in the traditional clinical setting, also feasible in a home situation if based on the alternative methods explained in the next section.

The next sections give an overview of the different alternative methods for seizure detection. The detection based on accelerometry (ACM) and video are discussion in separate sections (Section 1.3.2 and Section 1.3.3, respectively) as these modalities are the main focus of this thesis.

[3] *Chronic Obstructive Pulmonary Disease* is a term to indicate diseases in which the airways become narrowed which makes breathing becomes more difficult.

1.3.1 Alternative methods for detection

Numerous researchers have come up with different solutions for the monitoring of patients, based on different types of sensors, e.g. Schepers *et al.* [158] use 3D ACMs and 3D gyroscopes in combination with 3D magnetometers to track human motion. Suzuki *et al.* [165] monitor human respiration and heart beat in a contactless way by means of microwave radars. Especially thanks to the development of MEMS[4], the usage of for example ACMs became more interesting in human motion detection, as these systems are smaller, lighter and less power consuming than before. The usage of MEMS became also popular for biomedical applications as described by Polla [143] and Petersen [139].

Some sensors may be better suited for detecting specific types of seizures [181]. ACMs are well suited to detect clonic seizures because of their jerky movements. Gyroscopes are more useful for versive[5] seizures as they measure the rotational acceleration. Magnetometers are interesting for detecting tonic seizures, as in this type of seizures slow movements take place [13].

Note that the performance measures used in the next sections are explained in more detail in Section 1.4.

Detection based on motor information

- **ACMs or video:** As the main modalities that are investigated in this thesis are ACMs and video, we will discuss them more extensively in two separate sections.
- **Bed sensors:** EmFit *bed sensors* (Emfit Ltd.) are already used for the detection of clonic and tonic-clonic seizures. In the literature, studies can be found that use the sensor for sleep monitoring [151] [104] and as a sensor for cardiac measurements [103], but to our knowledge there are no papers about the performance of seizure detection with the EmFit system. Other systems where no clinical studies are available are for example the Epi-Care by Danish Care (Denmark), Sensalert by Sensorium (UK), the Epi-Watcher by Vahlkamp (The Netherlands) and Ep-It by Alert-It (UK). Van de Vel *et*

[4] *Microelectromechanical systems* are systems integrating miniaturized mechanical and electromechanical parts, that lead to the development of a very broad range of miniaturized sensors (from pressure sensors, to radiation, chemical or temperature sensors). Sensors, actuators and systems can be integrated on a silicon substrate along with integrates circuits. This leads to small, light systems that consume less power than the traditional systems.

[5] *Versive seizures* are a type of partial seizures that are associated with the movement of the head and eyes to one side.

al. [181] published an overview of these different systems. One other system is evaluated in a published clinical study. The MP-5 monitor of Medpage (UK) makes use of sensors that are attached to the bed. In a study [27], 64 subjects were monitored for a total of 1528 hours, having eight tonic-clonic seizures in total. The MP-5 monitor was able to detect five of these eight seizures, generating 269 false positives (during 146 hours). This resulted in a sensitivity of 62.5% and a specificity 90.4%. Due to the relative high number of false positives compared to the number of seizures, the positive predictive value had a low value of 3.3%.

- **EMG (electromyogram):** *EMG* measures the muscle tension, and it is therefore also a good modality for measuring the motor activity. This is also very suited for the detection of tonic seizures, as in this type of seizure the muscles contract. Adriaans *et al.* [2] investigated the EMG signals to detect tonic and clonic seizures. In another study of Conradsen et al. [42], the zero-crossing rate of the EMG signal from the deltoid muscle was used to detect tonic-clonic seizures with a sensitivity of 100.00% and a median latency of 13.7 seconds. Furthermore, they also investigated the dynamics of the EMG activity during tonic-clonic seizures in a clinical study [43].

Detection based on other modalities

Not only the detection and classification of seizures based on the manifestation of the movements is investigated as an alternative for the video/EEG monitoring. Also other modalities are investigated, including several autonomic signals such as heart rate, temperature, skin conductivity (sweating) and respiration. Some of these modalities are, for example, already used in the polysomnography (PSG) for the study of sleep, such as EOG, ECG and respiration monitoring. Below, a list of the different modalities that are used or investigated is given.

- **Audio signals:** *Audio signals* can indicate the occurence of a seizure, as the patients may utter sounds during seizures such as stereotyped screams, singing or humming, (autonomic) laughing or weeping, lip smacking or bed noises as a result of movement. The advantage of recording audio is that the equipment has a low cost and is easy to install in the patient's room. Furthermore, the sensor is non-invasive and even contactless, so there is a maximal comfort for the patient. However, detection systems based on audio generally perform poorly. It is hard to distinguish the specific sound during seizures from normal vocalizations such as speech or snoring or noises that

originate from the background. Sometimes these types of noise can also be part of a seizure manifestation. Tonic-clonic seizures for example are often followed by stertorous breathing (snoring) in the post-ictal period [18].

- **Thermal cameras:** *Thermal cameras* are used in the recognition of human motions [81]. Capturing infrared light, they can overcome the problem of normal cameras as they may detect movement under sheets, if they do not block the heat too much. Furthermore, they can give extra information about the body temperature of the patients which may be a distinctive feature for separating seizures from normal movement. However, the resolution is lower than in conventional cameras and most of all, the cost of these cameras is much higher which make them less applicable in a home situation.

- **Radar systems:** *Radar systems* are already developed to detect heart beats, respiration and movements in patients. This contactless way of measuring these vital signs is more comfortable for the patient than the sensors that normally have to be attached. Suzuki *et al.* [165] used this approach for detecting vital signs (ECG and respiration) in an ambulance. Another advantage is that the RF[6] waves pass through the blankets so they are able to detect movement even if the patients are covered in their bed.

- **The electrocardiogram (ECG):** *The electrocardiogram (ECG)* shows the electrical activity of the heart. The heart rate, the heart rate variability (HRV) which indicate the changes in the beat-to-beat interval and the ECG morphology can be indicators of seizures. Changes in the heart rate might precede the clinical manifestations of the seizures or even the electrographic onset [7] of the seizures in some cases. Because of this, ECG can be interesting for early seizure detection. And although the ECG electrode still have to be attached to the patient, it is still easier to attach to the body and easier to interpret the signal compared to the EEG electrodes and signals.

- **Respiration monitoring:** Next to arousals or sleep apnea, *respiration monitoring* can also detect sighs or yawns which may be an indication for seizures. Also, a low arousability can be a sign for near-SUDEP (sudden unexplained death in epilepsy) [85, 140]. Respiration can be monitored in different ways, for example by sensing the airflow temperature, pressure or velocity. These methods need a mask covering the nose and mouth of the patient. And although this is a common way to monitor the respiration, it

[6] RF or radio frequency often denotes wireless and high-frequency signals, comprising the frequency range between a few kHz to 300GHz.

[7] Start of the seizure in the EEG signal, this usually happens before the seizure manifests itself clinically.

is uncomfortable for the patient. Another way for respiration monitoring is by making use of belts attached around the chest or abdomen, to record the respiratory rate, depth and effort. Alternatively this motion can also be monitored by EMG-sensors, ACMs, video or radar, as discussed above.

Pulse oximetry measures the concentration of oxygen in the blood. The sensor can be easily attached to the patients finger or ear, and by measuring the reflected infrared light from two different wave lengths it determines the levels of oxygenated and deoxygenated hemoglobin. The reflected light differs from the emitted light, as for example the skin tissue and blood absorb part of it. From this signal, different characteristics can be derived, such as heart rate and respiration. There is a complex interaction between the brain, heart and respiration [181], and data on oxygenation is crucial in addition to respiration and ECG monitoring [140, 154]. Related to epilepsy this is mainly important for SUDEP prevention as the oxygenation level can indicate for example if there is a blockage of the airway in the patient. Additional to (or combined with) the pulse oximetry, *finger plethysmography* can be used for the monitoring of blood pressure [181] or blood flow [77]. Finally, respiration can also be detected via audio, either by using a microphone at a certain distance from the patient, which has the advantage of being contactless. Another possibility is to attach a small microphone to the suprasternal notch[8] in the neck as done by Corbishley and Rodríguez-Villegas [44]. This device is able to detect even shallow breathing, by recording the sound created by the turbulence in the human respiratory system.

Combining the oxygen saturation and chest and abdominal movement, extra information can be generated in the detection of seizures. A fall in both pressure and saturation combined with a slower, more irregular or absence of respiration and with tachycardia [9], indicates an epileptic seizure [140].

- **The electrooculogram:** *The electrooculogram* records both eyelid and ocular movements. It is measured by electrodes, the same as for EEG, only the placement is such that the electrical muscle activity of the eyes and eyelids is recorded. Eye movement can be a discriminating factor in differentiating epileptic seizures from non-epileptic seizures (PNES), which are events that look like epileptic seizures but do not have the same EEG manifestation. At the onset of a convulsive epileptic seizure, the eyes of the patient tend to be open in contrast to convulsive PNES where the eyes are typically closed [63].

[8] The *suprasternal notch* is a small dip at the base of the neck, at the top of the sternum.
[9] Fast heart rate, typically over 100 beats per minute.

- **The electrodermal activity:** *The electrodermal activity* (EDA) is a measure to define the ability of the skin to conduct electricity. During seizures, patients might start sweating, which influences the conductivity of the skin. The changes in this conductivity can simply be measured by an Ohmmeter. EDA reflects sympathetic activity[10], in contrast to blood pressure, heart rate or respiration which are ortho- and parasympathetic activities[11]. For example tonic-clonic seizures show an increase in especially sympathetic discharges [141]. A humimeter can also be used for the detection of sweating, but also for other symptoms that can be related to epileptic seizures such as salivation, vomiting or incontinence, although they are not related to seizures only. Some commercial systems are already available such as the Comf-It by Alert-It and the Epi-Wet by Danish Care [181].
- **The body or skin temperature:** *The body or skin temperature* may be an indicator for seizures. It is a measure that can be used in the detection of febrile seizures [12], although the relationship between this type of seizure and epilepsy is not exactly known [181].

1.3.2 Accelerometry

The ACMs in this study make use of a capacitive sensing cell for measuring the mechanical acceleration in an electrical way. Three conductive elements form two back-to-back capacitors. The central element is attached to a central mass that moves with the acceleration, changing the capacitance of both capacitors. This enables sensing the acceleration.

One of the early studies in the detection of activities and the evaluation of gait using ACMs is the work of Veltink *et al.* [185, 186]. In their research they discriminate postures and movements using uniaxial ACMs.

As already mentioned in Section 1.3.1, the use of ACMs became increasingly popular thanks to the miniaturization thanks to the usages of MEMS. Nowadays, a lot of the electronic devices as laptops or cell phones even have integrated ACMs. Kwapisz *et al.* [109] used ACMs from cell phones to detect walking, jogging, ascending and

[10]The *sympathetic activity* is activity initiated by the sympathetic nervous system, which is a part of the autonomic nervous system. It is related to stress or activity (so called fight-or-flight responses) and leads for example to an increased heart rate, blood pressure and a lower digestion activity.

[11]The *parasympathetic activity* is activity initiated by the parasympathetic nervous system, which is also a part of the autonomic nervous system. It is related to activities that occur when the body is in rest, allowing proper functioning of the internal organs.

[12]A convulsion related to a significant rise in body temperature.

descending stairs, sitting and standing. Brezmes *et al.* [24] and Yang [196] also used cell phone ACMs for similar movements as Kwapisz *et al.* but Brezmes *et al.* also including falling and Yang included driving and bicycling in his study. Also Gyorbiro *et al.* [80] and Miluzzo *et al.* [132] also used the internal ACMs of a smart phone for activity recognition.

In most studies, common activities such as walking, sitting, standing, running, lying, ascending and descending stairs are classified [7, 106, 107, 114, 130], using different classification methods such as decision trees, linear discriminant analysis or support vector machines. Mathie *et al.* [129] included falling events in their study and investigated movements after the fall incident. Ravi and Dandekar [152] included, in turn, brushing teeth and vacuum cleaning.

Some research groups also included other sensors in their system. Choudhury *et al.* [38] propose a platform which is able to collect data from for example gyroscopes next to ACM data and they proposed different pilot studies where the system could be used. Cho et al. [37] combine ACMs with a camera attached to the subject, to detect other types of movement such as taking the elevator and walking forwards and backwards which is well distinguishable in combination with the camera. Tapia *et al.* [166] made use of the heart rate next to 5 ACMs for recognizing activities in the gym. Inooka et al. [90] integrated ACMs with a barometer, which permitted a better detection of vertical movement such as stair climbing.

Accelerometry is widely used in the classification of movements and daily activities. Fudurić et al. [71] use a k-nearest neighbor approach for identifying activities such as falling, walking, running, sitting and standing. But it is also used for the detection of hand gestures as elaborated by for example Pylvanainen [145].

The more common applications for ACMs in human motion analysis is the analysis of gait. The work of Derawi gives an elaborated overview of the groups working on the extraction of gait parameters using ACMs [61].

The use of ACMs for the detection of human movement and epileptic seizure is already studied by Nijsen *et al.* [135]. Five 3D ACMs were attached to the four limbs (ankles and wrists) and one was attached to the chest.

During the last decade researchers started to investigate ways of detecting seizures with a motor component in a less intrusive way by means of ACMs or video. Last year, the first detectors that are built in in a wrist-watch were presented by BioLert (the EpiLert watch) and Smart Monitor Company (the SmartWatch). Both systems have recently been validated in clinical studies [105, 120], mainly on generalized tonic-clonic patients. In the study of Lockman *et al.* [120], the SmartWatch was

succesfully tested on a group of 40 patients (6 with tonic-clonic seizures). Seven of eight tonic-clonic seizures, and 204 non-seizures events were detected, with only one false detection during sleep. Kramer *et al.* [105] validated the EpiLert in a study on 31 patients. 20 of the 22 seizures were detected with a total of 8 false alarms during the 1 692 hours of monitoring. These sensors are ambulant and can therefore be worn during the day.

Jallon *et al.* [94], Cuppens et al. [48] and Nijsen et al. [136] used *multiple ACMs* attached to the extremities for seizure detection. Nijsen et al. [136] already investigated the detection of seizures by means of ACMs, used on myoclonic seizures. They tested various methods, resulting in a maximal sensitivity of 80%, but with a low PPV. The best PPV reached was 16%. Jallon [94] also investigated epileptic seizure detection based on ACMs. His algorithm is based on a Bayesian approach using hidden Markov models, resulting in a sensitivity of 88% and 89% in two patients. The according PPV was 75% and 55%, respectively. Conradsen *et al.* [41] used a multimodal approach for seizure detection, mainly on tonic seizures. They obtained a sensitivity between 91% and 100% with a specificity of 100%, this means that negative events are all correctly classified. A general conclusion from these papers is that large and intense seizures (tonic-clonic, hypermotor) can be detected with a high sensitivity and a low number of false positives, and that smaller seizures (myoclonic jerks) can be detected but with a higher number of false positives. Several systems are already developed (and some already commercialized) that integrate multiple modalities, a combination of ACMs, gyroscopes and magnetometers, to localize the sensor nodes. When these sensors are attached to the patients' extremities, they allow for a detailed movement analysis [20, 124, 158]. The advantage over for example video detection, is that ACMs can measure body movement more easily under blankets and can better separate the movements of the individual limbs, but the sensors have to be attached to the body.

1.3.3 Video

Human motion analysis has already been studied extensively in other applications such as gesture recognition or surveillance [87] [176]. Bonroy *et al.* [21] monitor the visual expression of humans using a camera system for pain assessment. This automatic vision-based monitoring can also be a solution for monitoring of epileptic patients.

Vision based human action recognition has been widely studied, and various methods have been proposed [144]. Based on representations, these approaches

can be broadly categorized into global representations and local representations. In the former, the entire body or its articulated poses are encoded in the model. In this work, because the body is usually occluded and not fully observed, local representations are used, where the observation is described as a collection of local descriptors. In one of the investigated approaches, we use space-time interest point detectors and descriptors proposed by Ivan Laptev in [112], which achieves state-of-the-art performance in action recognition in video on real life actions. In that study, different realistic actions from movies (such as kissing, answering the phone and getting out of a car) are learned. This method outperforms the other algorithms on the KTH actions dataset [159] and reaches an accuracy of 91.8%.

Video based detection is part of the gold standard for seizure detection and it is the most common way (and most of the time the only way) for clinicians to judge the detected event visually. Next to these diagnostic purposes, this modality can also be used for detection. It has the advantage of being non-invasive and contactless, that is, in the case that no (reflective) markers attached to the patient are used. A downside here is that most video based detection approaches make use of markers [33,115,153] or other ways to track limbs, like using colored pyjamas [122]. Another remark is that in most of the studies no real seizure detection algorithm is developed. The movement quantification is mainly used to investigate the behavior and motion pattern of the specific seizures and some features to quantify these patterns are proposed. Although Rémi *et al.* [153] investigated different features to identify hyperkinetic[13] seizures from non-hyperkinetic seizures. Karayiannis *et al.* [96] did not use any markers, but the moving limbs of the patients were clearly visible as they were monitored in the Neonatal Intensive Care Unit (NICU) of the hospital. The best obtained result had a sensitivity above 90% and a specificity above 85%, in patients with myoclonic and focal seizures. Another disadvantage of using video for seizure detection is the difficulty of detecting movement under the sheets or blanket, and that the system is not ambulant as the patient should be in the view of the camera at all time.

1.4 Measures

In the following chapters of this thesis, we evaluate the performance of the proposed algorithms. To this end we introduce different measures in this section, namely the

[13] Hyperkinetic seizures are marked by intense movements of the proximal segments of the limbs and trunk [191]. The term hyperkinetic seizure was proposed instead of the term hypermotor seizure by the ILAE task force on seizure classification and terminology [19, 191]

sensitivity, Positive Predictive Value (PPV), specificity and the Negative Predictive Value (NPV).

$$\begin{aligned}\text{sensitivity} &= \frac{TP}{TP+FN} \\ &= \frac{\text{number of correctly classified positive events}}{\text{total number of positive events}}\end{aligned} \tag{1.1}$$

$$\begin{aligned}\text{specificity} &= \frac{TN}{TN+FP} \\ &= \frac{\text{number of correctly classified negative events}}{\text{total number of negative events}}\end{aligned} \tag{1.2}$$

$$\begin{aligned}\text{PPV} &= \frac{TP}{TP+FP} \\ &= \frac{\text{number of correctly classified positive events}}{\text{total number of positive detections by the algorithm}},\end{aligned} \tag{1.3}$$

$$\begin{aligned}\text{NPV} &= \frac{TN}{(TN+FN)} \\ &= \frac{\text{number of correctly classified negative events}}{\text{total number of negative detections by the algorithm}}\end{aligned} \tag{1.4}$$

where TP, FN, FP and TN stand for true positives, false negatives, false positives and true negatives, respectively. Table 1.1 gives a clearer overview of the interpretation of these values. Depending on the application, a positive event can be, for example, a movement epoch or a seizure. Whereas a negative event can be, for example, a non-movement epoch or a normal movement. In the movement detection part, the *sensitivity* gives the fraction of the correctly detected epochs from all labeled movements, it is a measure that indicates the percentage of movements that is detected, also indicated as recall. The *PPV* gives the fraction of the detections (TP + FP) which are TP, this measure is also indicated as precision. The *specificity* is a measure that indicates the number of TN over all negative events. The NPV is a measure to indicate the correctly classified negative events (TN) over the total events that were not detected (TN + FN) by the algorithm. Ideally all these measures are one. In case of unbalanced data (we assume the case when there are more negative than positive examples), a high value for the specificity may lead to the incorrect interpretation that the algorithm performs well, as a large number of true negatives can predominate the results. In this case it is better to use the PPV.

Table 1.1: *Classification class outcomes.*

		Detection algorithm outcome	
		True	False
Annotated real data	True	TP	FN
	False	FP	TN

1.5 Motivation of research

Having epilepsy can have a major impact on the quality of life of the patients but also on their environment. This research was initiated based on the need from parents with epileptic children to have a better tool for monitoring the epilepsy. Especially during the night when they cannot keep an eye on their child. The parents actively asked for such a system as the commercially available systems often do not suffice in sensitivity or generate too much false alarms. Also, for many of them the clinical impact is not proven in a clinical study.

The current state-of-the-art for seizure detection is only applicable in a clinical setting as discussed in the previous sections. Therefore we want to develop a system that allows for home monitoring to detect seizures as an alarm system for the parents, but developed based on data that is labeled according to the gold standard. But added to that, we also want to keep track of the seizures, to have an objective measure of their number as an extra input to the neurologist. As the child can be monitored for a long period of time, more reliable and detailed information on the disease can be generated.

1.6 Challenges and objectives

In this thesis we focus on two types of seizures, *hypermotor* and *myoclonic* seizures. We want to detect them in an automated way based on two modalities that are less intrusive compared to the gold standard, namely using ACMs and video. We aim at detecting nocturnal seizures as during sleep the seizures occur without too much interference of for example voluntary movement. Furthermore, as stated above, during the night, the parents are more likely to miss a seizure.

Four objectives are defined in this thesis:

1. Extract from the raw video and ACM data segments in which movement occurs
2. The separated events should be classified as epileptic or normal movement based on the recorded modalities. We restrict the research to the acceleration of limbs and video.
3. The approaches based on both modalities should be compared and integrated to obtain a higher performance
4. From the detected seizures, information about the occurrence and type of seizure should be extracted

To be able to accomplish the preset objectives, the used approach should address the following issues:

1. The sleep situation of the patients should be influenced as little as possible, hence no usage of markers in video approach. Furthermore the usage of sheets and blankets is allowed
2. The number of sensors should be as low as possible in order not to hamper the patient during the night
3. The hardware should be easily installable in the patient's room
4. To be of good practical use, the used algorithms should be able to work in real-time

1.7 Chapter-by-chapter overview

Figure 1.1 gives a schematic overview of the different chapters and the relation between them. There are different steps involved in the detection of seizures. First the data have to be acquired, this is explained in Chapter 2 and indicated in blue in Figure 1.1. In a next step, the acquired data is preprocessed, and the relevant movements are extracted, this is discussed in Chapter 3 and Chapter 4 and indicated in green in Figure 1.1. An integration of both approaches in a mobile system is elaborated in the future work section in Chapter 8. As a last step, the different movement epochs are classified as seizure or normal movement. Chapters 5, 6 and 7 discuss this topic for the different modalities, in the overview figure, these chapters are shown in yellow.

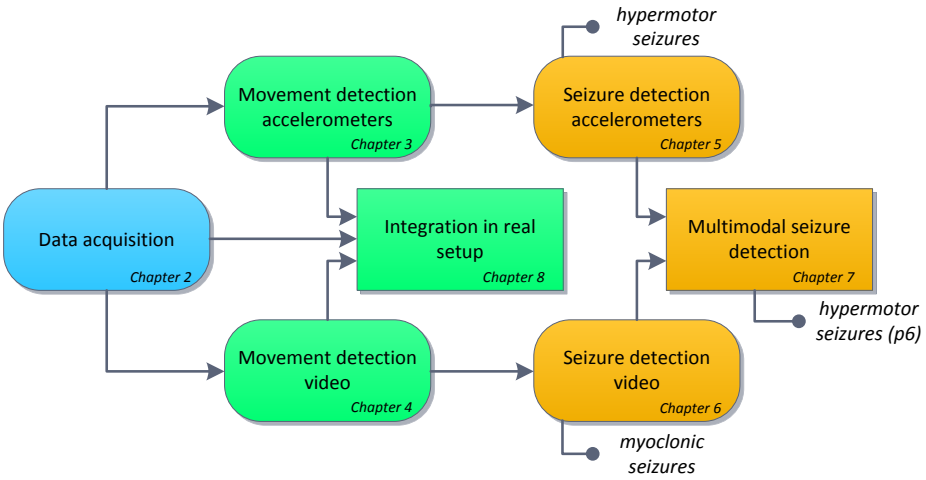


Figure 1.1: *Schematic overview of the different chapters*

Chapter 1 introduces the reader to the background of epilepsy, both in the electroclinical causes and in the different ways of monitoring and treatment. Furthermore, it elaborates on the current state-of-the-art in action recognition in video and ACM signals. Finally, it summarizes the motivation, the challenges and objectives of the research and it gives a perspective to the following chapters of the thesis.

Chapter 2 gives an overview of the acquired data during the research. Data is acquired at different locations and under different circumstances such as real and simulation data. Also in most datasets the hardware that was available in the hospital was used, but one dataset containing data from the home monitoring system was collected. Furthermore, this chapter gives a full list of the acquired datasets.

Chapter 3 elaborates on the preprocessing steps that are performed on the raw ACM data. After this part the movement detection is explained. The movement detection is an important part for reducing the quantity of data and segment the continuous data in separate segments (events).

Chapter 4 gives, similarly to Chapter 3, an overview of the preprocessing and movement detection of the video data. In this chapter the optical flow technique

is also introduced, which is also used in the further chapters. Furthermore, the detection of the respiration signal from the video recordings is discussed.

Chapter 5 compares three classification methods used for the ACM signals, each with their own advantages and disadvantages. We start with simple thresholding of intuitive features, going to a state-of-the-art *support vector machine* (SVM) approach and we finish with an approach based on outlier detection by determining the probability density function of the normal movement. The last part of this chapter gives a comprehensive comparison of the different approaches.

Chapter 6 proposes two different methods in the detection of seizures in video. One method makes use of the optical flow calculation, already used in Chapter 4. The second method uses the detection based on *spatio-temporal interest points* (STIP) which is the state-of-the-art in action recognition in video.

Chapter 7 integrates both modalities, ACM and video, in a combined detection, making use of the findings of the previous two chapters. The idea here is to improve the performance of the classifier by exploiting the complementarity of both modalities. Both an early integration (combining features) as a late integration (combining classifiers) is investigated in this chapter.

Chapter 8 Concludes this thesis with the main findings and an outlook to the future. It mentions the main conclusions of each chapter and gives the most important items to focus on for future research.

1.8 Personal contributions

My own contribution to this work ranges from supporting the data acquisition (by taking care of the well functioning of the ACM acquisition hardware), to the application of different methods for seizure detection and finally their integration in a multi-modal detection.

I started working on the detection of seizures using video data, in collaboration with prof. Lieven Lagae from the University Hospital in Leuven. This resulted in different papers in conference proceedings and a journal paper publication on the detection of body movement in epileptic children in *Medical & Biological*

Engineering & Computing [49]. Chapter 4 addresses this topic. Further steps were taken, also thanks to valuable feedback from prof. Tinne Tuytelaars, which led to a first detection of myoclonic seizures in video. These results were published in the proceedings of the Biosignals conference of 2012 [50].

A side topic of this research was the detection of respiration which was feasible with the collaboration of Katrien Janssen and prof. Lieven Lagae from the University Hospital in Leuven. The results from this research were published in a conference paper of the *International Conference on Positioning and Context-Awareness* (PoCa) [182].

In 2008 a IWT TBM project started which aimed at the detection of epileptic seizures using ACMs, and built on our early work of the ACM detection, in collaboration with prof. Bart Vanrumste and prof. Sabine Van Huffel, my promotors, and prof. Berten Ceulemans, prof. Lieven Lagae, Bert Bonroy, Ann-Sofie Decaigny, Anouk Van de Vel, Milica Milosevic. Here I improved the detection of hypermotor seizures using ACMs, which was published in several conference proceedings and is now submitted for publication in *IEEE Transactions on Information Technology in Biomedicine* [53]. Futhermore, with the input from Peter Karsmakers, I investigated the influence of the imbalance of the data on the classification results which was submitted for publication in the proceedings of the *International Conference on Pattern Recognition Applications and Methods* (ICPRAM) [54].

During a two months stay at the Stanford University I had the opportunity to collaborate with the Ambient Intelligence Research Lab of prof. Hamid Aghajan. Together with my colleagues there, Kevin Wong and Louis Chen, I worked on the detection of seizures using the STIP method. This is published at the IEEE EMBC 2012 conference [51]. This collaboration also contributed to the integration the ACM and video modality. A first study on this topic is presented at the *ACM International Conference on Multimodal Interaction* (ICMI) and published in the proceedings [52].

Chapter 2

Data acquisition

This chapter gives an overview of the acquisition of the datasets. First, a short overview on the requirements and the specific issues for gathering the data is given in Section 2.1. Also, a list of systems that are already available on the market for recording seizures or nocturnal movement is given. After this introduction, two types of datasets are discussed, the acquisition of simulation data in Section 2.2 and real data in Section 2.3. Both sections discuss the patients or subjects, the setup of the acquisition system and an overview of the acquired data.

2.1 Introduction

Epileptic patients are traditionally monitored with video, EEG and often also with ECG and EMG. Therefore, the acquisition hardware for these modalities is already available in the hospital setting. Furthermore, many datasets of patients are already gathered and stored. However, most of the time these datasets are used for clinical purposes only for the diagnosis of epilepsy, rather than seizure detection for inventorying an extensive list of all seizures the patient has. Also the normal movements are normally not maintained. For developing an automated seizure detection algorithm, however, it is important to have enough positive and negative examples. Therefore, the existing databases with seizure data, are useful to a certain extend.

For monitoring the movement of the limbs, we make use of an extra set of ACMs on top of the existing measurement equipment. To measure these signals we also

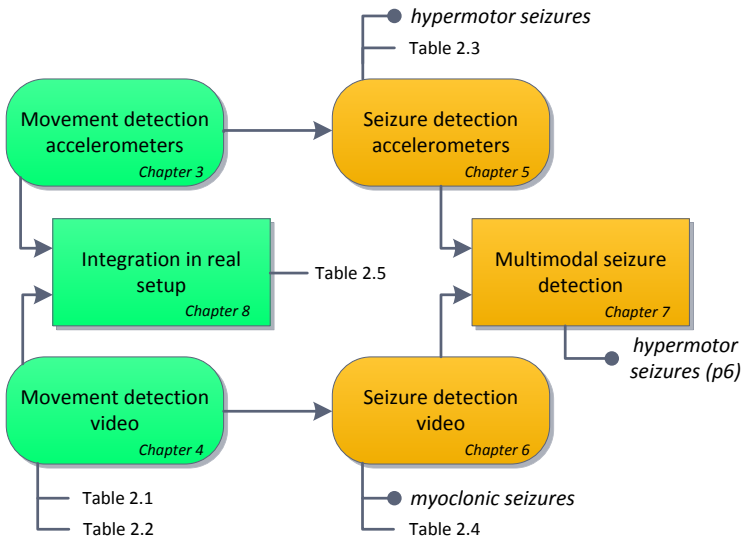


Figure 2.1: *Schematic overview of the datasets that are used over the different chapters.*

have built and integrated an ACM acquisition system with the existing monitoring hardware of the hospital.

Figure 2.1 gives an overview of the type of seizures and according datasets that are used in each chapter. The number of the table in this figure corresponds to the tables in this chapter.

2.1.1 Requirements

When developing a detection system that is usable in a home situation, we need to keep in mind the following limitations that apply, that were already mentioned in Section 1.6:

- The sleep situation of the patients should be influenced as little as possible, hence no usage of markers in video approach, and usage of sheets and blankets is allowed
- The number of sensors should be as low as possible in order not to hamper the patient during the night
- The hardware should be easily installable in the patient's room

However, when gathering data in the hospital setting for developing the detection algorithms, some of these restrictions are not as strict. For example, during the data acquisition the limitation of the number of sensors is not an issue. In this phase it is more important to gather data from different modalities and determine which work better under which circumstances.

When selecting a system for the wireless acquisition of the ACM signals, the following characteristics should be present:

- Autonomy of battery at least 12 hours (one full night)
- Packaging should be small enough to fit around children's (and baby's) limbs
- Data transmission should be robust (minimize lost data)
- Raw data should be available (not present in some commercial products)
- Mechanical robustness (throwing, biting, water resistant)

The hardware in the hospital for monitoring the traditional modalities in epilepsy was readily available and no elaborate analysis of the system requirements had to be made.

2.1.2 Commercially available systems

Keeping in mind the requirements that were determined in the previous section, a hardware platform was selected that met these requirements to be integrated in our home monitoring system. Some of these systems include only sensors for motion detection, whereas others provide for a multimodal integration.

There are products available that can monitor different modalities such as ECG, audio or respiration. Some of these modalities have proven a correlation with epileptic seizures, as already discussed in Section 1.3.1. For example an ECG bracelet with chest electrodes by Holst Centre/IMEC and Hobo Heeze B.V. was tested in Massé *et al.* [128]. In fields other than epilepsy, these devices have already proven their value, for example for cardiac monitoring and they are commonly used. Their reliability for the detection of epileptic seizures, however, is most of the time unproven [181] and further research is required [161].

In motion detection systems, the pedometers, such as the Orthocare Innovations Stepwatch Activity Monitor or activity monitors, such as the McRoberts Dynaport ADL Monitor, are wildly spread and available. They can store the data locally or

send it to a base station or computer. Most of these systems include processing of the data, such as calculating activity classes or energy consumption. As the data is stored locally or only the processed data is transmitted or stored, the battery autonomy can be multiple weeks in some systems depending on the sample frequency. There is a wide range in system specifications that can be applied in different situations. However, most of the time, these systems do not have the opportunity to collect the raw data. As we focus on a different application, we want to have access to the raw data and extract different features from it. Furthermore, these systems usually are only limited to one application and lack a certain flexibility to be adapted by the user.

Another group of systems is more focused on the research applications. These systems typically have the opportunity to be modified to the users wishes and the raw data can be accessed easily. They typically have multiple sensors, a shorter autonomy but a larger flexibility. Examples of those kind of systems are the MTx sensors of Xsens Technologies, the ETB PEGASUS-G, the Philips PI-node, Wireless Tracking system or the Shimmer Wireless Sensor Platform[1].

In our application we first looked for a wired setup for collecting ACM data from the patients. The acquisition hardware in the hospital provides an easy connection for extra analogue inputs. Therefore, the ACM data that is acquired in Section 2.3 are all recorded with a wired setup. To make the sensor system more comfortable for the patients, in a next step, we want to do all the recordings with wireless sensors. To this end we started to use wireless nodes from Shimmer which starts to become increasingly popular in biomedical research. We chose this platform for its open architecture, the nodes are programmable by the user, and user-made sensors can easily be attached to the nodes. Furthermore, the system meets the requirements listed in Section 2.1.1, and the sensors are relatively cheep (250 euro) compared to most medical sensors. Two different setups are currently under development and are tested. The first one is the integration of the wireless nodes in the BrainRT[2] system from OSG[3], which provides the epilepsy monitoring system in the hospital. In a collaboration with OSG, the wireless nodes will be integrated in their system so that ACM data, or other modalities that can be recorded with the Shimmer nodes, can synchronously be acquired with the video/EEG data directly. The second system is a full stand-alone system that integrates video and ACM recordings (also combined with a radar sensor) in a mobile setup that is easily installable in the patient's room in a home situation. In Section 8.2.2 we will elaborate more on the first findings of the use of this system.

[1] Shimmer Research: http://www.shimmer-research.com
[2] (http://www.brainrt.com/)
[3] http://www.osg.be



Figure 2.2: *Two camera view points during recording of the simulation videos. The left picture shows the foot end view, the right one shows the top-view. Both images were taken at the same time as both video recordings were synchronized. The shown images were recorded during daylight with an adult simulator not covered with a blanket.*

2.2 Simulation data

Simulation data was recorded to be used in the movement detection step, which performs a data reduction and segmentation. Part of the data was simulated in these tests to define precisely the smallest movement we still want to be able to detect. Furthermore, also epileptic seizures and frequently occurring normal nocturnal movements such as turning over in bed were simulated to get a first idea of the influence of the movement on the signals.

2.2.1 Video data

The simulation video was recorded at the UZ Leuven hospital at the department of child neurology. The simulated movements included 12 normal nocturnal movements such as turning in bed, tucking legs, turning the head, eye blinking and getting out of bed, and 11 simulated seizure-like movements such as a small jerk with the shoulders, stretching the arms, large random movements with arms and legs and overstretching the body. We also recorded the movements with and without the blankets, and recorded movements using a sheet with a raster that reflects infrared light. The video recordings of the infrared reflecting raster were acquired to be able, in future work, to reconstruct the movement under the blankets in a better way. Furthermore, thanks to the pattern, the optical flow calculations, which are discussed in Section 4.2.2, in this setting worked better. Data from one adult simulator was used.

Table 2.1: *Overview of the used video sequences of the simulated video data.*

Data	Simulator	Lighting	Type of movement	Actions
Sequence 1	Adult	Daylight	Normal	Movement with head, scratching the nose
Sequence 2	Adult	Daylight	Epileptic	Myoclonic jerk, overstretching the arm
Sequence 3	Adult	Dark	Normal	Movement with head, scratching the nose
Sequence 4	Adult	Dark	Epileptic	Myoclonic jerk, overstretching the arm

This data was used to evaluate different settings (frame rate, compression, lighting conditions, resolution, nature of movement) on the movement detection algorithm.

The video was acquired in the presence of an experienced neurologist who gave input on the performed movements, in a room which could be darkened. Two different camera setups were used in for recording the simulation videos. The first one was a Cypress Stereo XAECK10005 camera system with two cameras that could take images synchronously at 30 frames per second. The images are monochrome and have a resolution of 640x480 pixels. One camera acquired a top-view of the bed, the other acquired a view from the foot end towards the head, as also visible in Figure 2.2. The recordings in this figure were from the adult simulator not covered with a blanket during daylight. The both cameras from this first setup were no infrared cameras. In the second setup, infrared videos were recorded at the rear end of the bed, to investigate the influence of the infrared reflecting grid.

Each video record is about 20s long. Initially the subject is in rest on the bed, then a certain movement is performed as described above and finally the subject is again in a resting state.

With this camera system a sequence of movements was taken in full daylight (27.5

lux at the head of the bed and 16.8 lux at the foot end) and in a darkened room (4.7 lux at the head as well as the foot end). A subset of the recorded video sequences is used in this thesis in Chapter 4. An overview of the used video sequences is given in Table 2.1.

2.2.2 ACM data

The ACM simulation data was acquired to determine the thresholds for the movement detection algorithm, as a data reduction and segmentation, which is elaborated in Chapter 3. In this dataset different small movements were simulated, some of them we want to detect and others not. For example a sequence of 5 minutes was recorded where the arms of the subject rested on the chest to obtain noise from the breathing of the subject. This kind of movement is not relevant. Other small movements, such a subtle movement of the fingers or toes were also recorded. These movements are considered as the smallest movements we want to detect in the seizure detection application. Other influences of noise that could be picked up by the ACM acquisition hardware were also recorded, e.g. turning off and on the fluorescent lighting and the influence of an active cell phone next to the signal amplifiers.

2.3 Real data

In this research, we collaborated with two institutions for the acquisition of real data from epileptic children. We could make use of the datasets that were already acquired at the Child Neurology department of the UZ Leuven, containing the traditional recordings of video/EEG data. Later on in the project, we started to collect data in the Pulderbos Rehabilitation Center for Children and Youth, where we combined the video/EEG measurements with ACM data.

2.3.1 Location 1: UZ Leuven

The department of child neurology at the university hospital of Leuven has a large collection of nocturnal recordings of epileptic patients with EEG and video data. A list of patients with according seizures was made available to select appropriate fragments from the collection. Over 80 single seizures and 5 full nights, containing normal and seizure data, were used from this collection. These full nights were

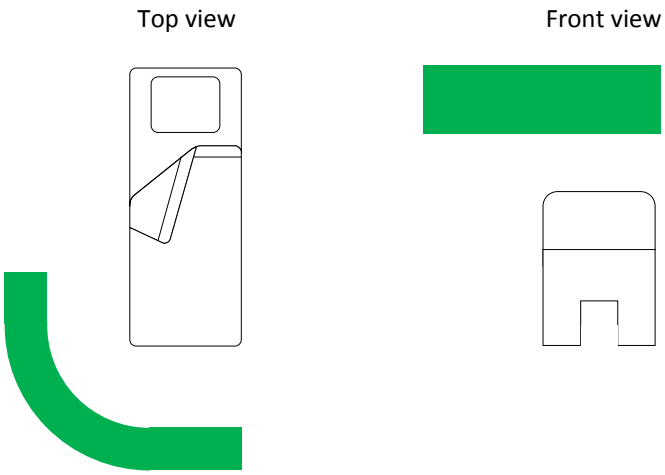


Figure 2.3: *A sketch of the camera positions in the different video sequences. The marked areas indicate the place where the camera is placed.*

especially maintained for our research to have a set of normal movements included in our dataset.

Patients

The patients in the dataset suffered from different types of epileptic syndromes such as Lennox-Gastaut, West-syndrome, myoclonic seizures, frontal lobe seizures and absences. The youngest patient in this dataset was 21 days old, and the oldest one 25 years. Most of the patients take at least one anti-epileptic drug. But not all of them showed seizures during their time of monitoring.

Setup

In almost all the recordings, which happened at different locations in the hospital, the camera was placed in a top view position, sometimes only the movements of the torso and head were recorded, in the other cases, the whole body was captured. Combined with the video, also the EEG is recorded.

The video sequences were compressed with an MPEG-1 compression before they were stored. The resolution of the used video data was 320x240 with a frame rate of 25 fps (frames per second). The video data was acquired by BrainRT-software.

The data was labeled, i.e. the starting and ending points in time of the movement epochs (including epileptic seizures) were marked by an expert. Each sequence (epoch) contains a movement event. The length of the events ranges from 1 to 304 seconds.

The infrared camera was placed according to the sketch in Figure 2.3. The video was zoomed in on the patient in a way that the head, the arms and the torso was visible in videos 1,3,4,5,7,8 and 11 (an overview with the numbers of these videos is given in Table 2.2). In the other videos the whole patient was visible.

In one dataset, the breathing of patients was also measured using a respiration monitoring belt which measures the thoracic and abdominal movements in addition to the video recordings.

Datasets

From this setting, one subset was manually annotated by selecting all the movement events that were visible in the video. Table 2.2 gives an overview of the annotated data.

The inclusion criteria for the sequences to be a part of the dataset are the following:

- The recorded patients should be maximum 18 years of age.
- The recordings are taken in a nocturnal setting, a blanket is allowed.
- The video data is recorded with an infrared camera, hence grayscale images.
- The sequences should contain at least one movement epoch, which does not necessarily have to be of epileptic nature.

This means that some of the video sequences that were recorded during the morning when the patient was awake and out of bed were not retained.

The dataset consisted of a 1 hour 48 minutes and 44 seconds recording with 73 movement events from which 11 were convulsions. This dataset is a sequence of 11 recordings from 8 patients between the age of 9 months and 16 years.

The dataset that is collected using data from video and a respiration belt included four patients who were monitored during one night each. A subset of this data is used in our tests.

Table 2.2: *Overview of the obtained video sequences containing movement*

Sequence	Length	Movements	Epileptic movements	Patient
1	15 min	9	1	a
2	15 min	17	0	b
3	15 min	7	0	c
4	15 min	4	0	c
5	15 min	3	0	c
6	30 min	16	0	d
7	30 s	4	3	e
8	34 s	2	1	f
9	1 min 55 s	7	4	g
10	26 s	1	1	h
11	19 s	3	1	f
Total	*1 h 48 min 44 s*	*73*	*11*	

2.3.2 Location 2: Rehabilitation center Pulderbos

Patients

In the Pulderbos rehabilitation center for children and youth, children up to 21 years old with epilepsy are treated and monitored. The epilepsy center monitors and treats children with difficult-to-treat epilepsies. Children can be observed there for a longer period of time. During the last years we have been collaborating with the center within an IWT TBM[4] project to collect data from epileptic patients using the common modalities combined with ACMs.

So far we have collected a large dataset with 196 nights of recording of 64 patients, from which 34 have seizures during the recordings. In this dataset different types of seizures with a motor component are collected, such as myolonics, clonic seizures, tonic seizures, spasms, hypermotor seizures, versive seizures and tonic-clonic seizures.

[4] *Toegepast Biomedisch onderzoek met een primair maatschappelijke finaliteit*, Applied biomedical research

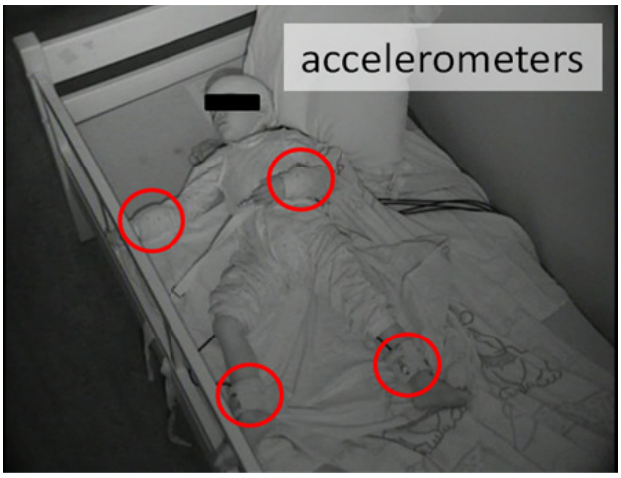


Figure 2.4: *This figure shows the acquisition setup. The placement of the ACMs is indicated by the circles.*

Setup

The ACM data was recorded synchronously with video, EEG-, ECG-, EOG- and EMG-data. The synchronization of the signals was done by BrainRT software from OSG, and a Schwarzer interface box which converted the analog input from the sensors to a digital signal for the connection with the recording computer.

To convert the physical measurement (acceleration) to an electrical measurement (voltage) we use the ADXL330 accelerometer. The ADXL330 is a small, low power, 3-axis acceleration measurement system on a single monolithic IC. It has a range of at least -3.0 g[5] to 3.0 g, and a typical range of -3.6 g to 3.6 g. With an applied voltage supply of 3.3 V, the ADXL330 delivers an output voltage of approximately 1.6 V by 0 g acceleration. This voltage increases or decreases with 320 mV per +1 g respectively -1 g.

In the developed acquisition hardware system, we make use of four 3 dimensional ACMs, attached to the extremities with comfortable wrist and ankle bands. The setup is shown in Figure 2.4. So in total we have 12 output channels (3 per ACM). For each output channel, a 32 kΩ resistor is integrated. The purpose of this resistor is to set the bandwidth by an external capacitor. In our setup, a 100 nF capacitor is used.

[5] With g being the unit gravitational force: $9.81 \frac{m}{s^2}$

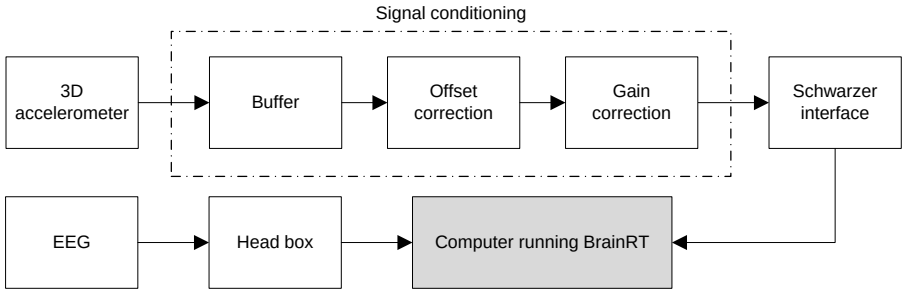


Figure 2.5: *This figure shows a block diagram of the acquisition circuit. The signal conditioning part is custom made hardware for our setup.*

Figure 2.5 illustrates a block diagram of the data acquisition circuit that is used. The 3D ACM's output voltage contains an offset and a low voltage deviation per g, as described above. *The buffer* converts the relative high impedance of the ACM output to an impedance of only a few Ohms. It has a high input impedance and a low output impedance, to prevent the ACM circuit from being overloaded by the following block. *The offset correction* eliminates the 1.6 V offset on the ACM output signal by subtracting an adjustable DC-voltage. This results in a signal of 0 V by an acceleration of 0 g, increased or decreased by 320 mV/g. *The gain correction* amplifies the signal by an adjustable factor to create a signal of 1 V/g. This results into a minimum range of -3 V to 3 V for an acceleration of -3 g to 3 g. The gain correction improves the accuracy after digitalization by the next block. *The Schwarzer interface box*[6] contains 12 analog input channels with a range of -5.5 V to 5.5 V and a 12-bit analog-to-digital converter with a resolution of 2.7 mV and a maximum sample rate of 1 kHz. The interface box sends the acquired data to BrainRT which also gathers the EEG and video data. In this way we log our ACM data synchronously with our EEG and video data, the Gold Standard data.

The video data is compressed using Microsoft's WMV compression. The video has a resolution of 720 x 576 with a frame rate of 25 frames per second. The camera is positioned in the upper corner of the room, monitoring a child sleeping in a bed. The image is recorded in grayscale using an infrared camera, no RGB information is available. Every night, approximately 5 gigabytes (GB) of video data and 800 MB of ACM data is collected.

[6] http://www.schwarzer.net

Table 2.3: *Dataset overview of hypermotor seizures acquired at Pulderbos rehabilitation center used for the ACM detection.*

Patient number	Nights of monitoring	Hypermotor seizures	Normal movements
1	1	2	117
2	2	9	287
3	2	2	439
4	1	2	239
5	5	26	784
6	2	7	381
7	2	3	468
Total	*15*	*51*	*2715*

Datasets

The annotation of the data was performed by an EEG expert based on the clinical manifestations in the video data and the patterns in the EEG data. When it was not clear whether an event was epileptic or not, a group of two EEG experts and two neurologists agreed about the type in a consensus.

As mentioned in the patient overview, patients with different manifestations of seizures were monitored while collecting the dataset. In this thesis however, we focus on the hypermotor and myoclonic seizures. Table 2.3 and Table 2.4 give an overview of the collected data that was used to this end. The hypermotor seizures were used in the ACM detection and the multi-modal approach combining ACM and video data. The myoclonic seizures were used in the video detection part.

One of the groups of patients we studied are patients, between the age of 5 and 16 years, with hypermotor seizures. Here we collected data from 7 patients, having a total of 51 seizures, and 2715 normal movements. The recordings took place during the night starting typically at 20h00 in the evening and finishing at 8h00 in the morning the next day. Table 2.3 gives an overview of the patients and the number of seizures and normal movements they had during the monitoring.

The other group of patients suffered from myoclonic seizures. The patients in this group are between the age of 3 and 8 years. Data from 3 patients is used, having a total of 27 myoclonic seizures and 82 normal movements. Table 2.4 gives an

Table 2.4: *Dataset overview of myoclonic seizures acquired at the Pulderbos rehabilitation center used for the video detection algorithm*

Patient number	Nights of monitoring	Myoclonic seizures	Normal movements
3	1	9	41
8	1	9	14
9	1	9	27
Total	*3*	*27*	*82*

overview of this dataset.

2.4 Home monitoring data

Setup

In the home monitoring system which is elaborated in Chapter 8, three modalities were combined, namely ACMs, video and a radar sensor. For this radar sensor, a module from Hygrosens Instruments was used [70]. This sensor is typically used for alarm and safety systems or for door openers. As it works through many materials such as plastic, it was integrated in a box of the monitoring system. The sensitivity of the sensor can be adjusted, in our setup it was tuned in a way to detect the movements of the patient when sitting upright in bed, or give detections when other people such as the nursing staff entered the room. In this way, it has an added value to the ACM and video signals.

The camera that is used is an AXIS M1011 IP-camera that used an external infrared light source to be able to record video during the night. The video recordings with a resolution of 640x480 pixels at 30 frames per seconds in RGB[7] are compressed with an MJPEG compression. This compression used a JPEG compression at each frame without taking into account temporal information to compress the data. This results in a lower compression rate, but the video can be cut at any frame without a significant loss in quality.

[7] *Red Blue Green* color model.

Table 2.5: *Data acquired using the home monitoring system. The number of seizures are the ones that are observed by the clinical staff, not the ones labeled based on the EEG/video gold standard, as these were not included in these recordings. Patients that were selected for these measurements suffered from hypermotor seizures.*

Patient number	Nights of monitoring	Total movements	Observed seizures
1	30	3952	3
10	3	2418	0
11	24	5557	0
Total	*57*	*11927*	*3*

To record the ACM data, we made use of the wireless Shimmer nodes as introduced in Section 2.1.2. These sensors were integrated in comfortable wrist bands, and are connected to a host laptop via a Bluetooth connection. The four ACMs are synchronously measured at 50 Hz with a range between -6 and 6 *g*. Next to the connection with the Shimmers, the laptop also records the other modalities using a LabView interface.

Datasets

Table 2.5 shows the data acquired with the home monitoring system. The aim of this research was to record data for one month in each patient, to evaluate if home monitoring would be feasible. The number of movements per night for patient 10 is high compared to the other patients. This is due to the fact that the third night of monitoring, 2141 movement events were recorded. During this night, the patient was awake almost all night. But he is known to have difficulties during sleep (e.g. biting and tearing sheets or clothes). He was very active and also took off the wrist and ankle bands multiple times, therefore the measurements on this patient were stopped.

Chapter 3

Preprocessing and movement detection - Accelerometers

The purpose of the preprocessing and movement detection step is threefold. a) It removes noise and other irrelevant information from the signals. b) It segments the data in separate movement events. And c) it reduces the amount of data hence speeding-up the post-processing. This processing can be done based on the ACM or the video signal. In this chapter, the processing based on the ACM signals is discussed. In Section 3.2, the methods and the results are covered, respectively. As a last part, this chapter ends with a conclusion that summarizes the main findings in Section 3.3.

3.1 Introduction

The raw data is collected and labeled by the neurologists. The starting points of the seizures or abnormal behavior is marked. But for an algorithm to be able to extract features from the data, the long data stream should be divided in shorter fragments, each representing one movement event. Another issue is the amount of collected data each night. As indicated in the previous chapter, approximately 5 gigabytes (GB) of video data and 800 MB of ACM data is collected. This amount of data should be reduced allowing for a quicker processing time and to avoid memory issues when processing. The preprocessing and movement detection step elaborated in this section addresses these issues.

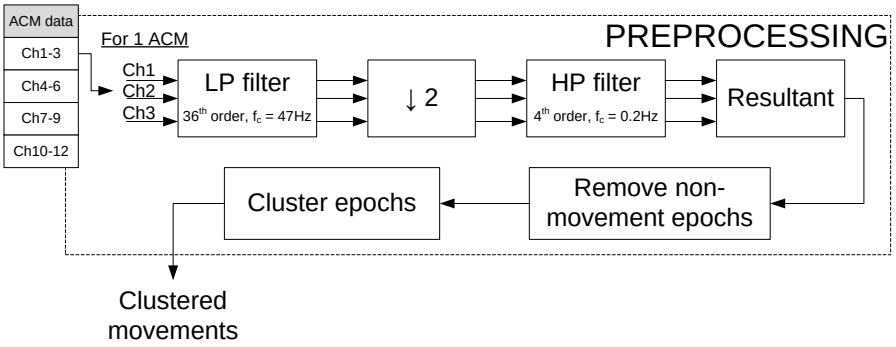


Figure 3.1: *Schematic overview of preprocessing step that detects the movements events.*

Another issue is the influence of noise sources in the signal. The ACM signal for example is distorted by the 50 Hz power noise. Also, the higher frequencies seem not to add any relevant information for the human movements. In the preprocessing step, these components of the signals will be removed.

3.2 ACM based detection

A preprocessing step is executed on the raw data. This step has three goals: first it removes the noise in the signals, second, it reduces the amount of data significantly, and third, it divides the data into separate movement epochs which are further used to classify as a seizure or a normal nocturnal movement. An overview of the different steps in the preprocessing algorithm is given in Figure 3.1.

The first step consists of the ACM data being filtered by a digital low pass 36th order type 2 Chebyshev filter with a cut-off frequency of 47 Hz. To avoid the phase distortion introduced by the filter, we processed the input data in both the forward and reverse directions, which results in a zero-phase filtering. The cut-off frequency is chosen such that the filter cancels out the aliasing effect when the data is downsampled. Moreover, a possible influence of the net frequency (50 Hz, and the first harmonic at 100 Hz) is also eliminated. We can easily downsample the data by a factor of two to 125 Hz because the frequency content above 40 Hz is negligible.

After the downsampling, the influence of the static earth gravity field (represented

by a DC component[1]) is eliminated by a high pass filter with a cut-off frequency of 0.2 Hz. Here, a fourth order type 2 Chebyshev filter is used. As with the previous filtering, here also a zero-phase filtering is used.

Next, the resultant of each ACM (the norm or the 3D vector amplitude) signal is calculated. As a result it gives the size of the dynamic acceleration of the sensor, inducing a reduction from twelve to four channels. The total data reduction in these first steps is a factor of six (from 12×250 samples per second to 4×125 samples per second).

After this data reduction step, the epochs without movement are discarded. To do this, the standard deviation of a sliding window of two seconds is calculated. If the standard deviation crosses a certain threshold, the epoch is considered as a movement, otherwise the epoch is discarded. The threshold is determined by a simulation, of which an overview is given in Section 2.2.2, where a simulator lies in bed during five minutes, after which he makes small movements of the fingers and toes which should be detected by the algorithm. This resulted in a threshold of 10 mg for the arms and 5 mg for the legs. This difference can be explained by the fact that toe movements have a lower influence on the ACM signal than the finger movements. Figure 3.2 shows the boxplots of the recorded data during the simulation. The plots of a sequence of no relevant movement is shown in Figure 3.2A, here the simulator was in rest with the hands on the chest. In Figure 3.2B, the simulator made small finger movements, and in Figure 3.2C, small movements with the toes were made. After this step, the data is clustered. Epochs within 30 seconds are merged and are considered as a single movement epoch, because the duration of the epileptic activity we want to monitor is typically long (tens of seconds) and sometimes, during one hypermotor seizure the patient stops moving (order of seconds). This 30 second duration is chosen specifically for the dataset with hypermotor seizures, as these seizures last typically long in time. If shorter seizures are to investigated, a shorter duration for the clustering is advisable.

Note that after this clustering step, the movement events have a variable length in time. In the remainder of the ACM research, these events with variable length are used.

Ten to twenty percent of the datasets are considered as movement events. Considering all the preprocessing steps, this means that only about 1.5% to

[1]However, when an ACM rotates, the influence on the different channels of the gravity field, which is visible as a static acceleration, changes. The resulting vector stays the same but its decomposition on the three channels is different. This often results in a block-shaped pattern in the ACM channels.

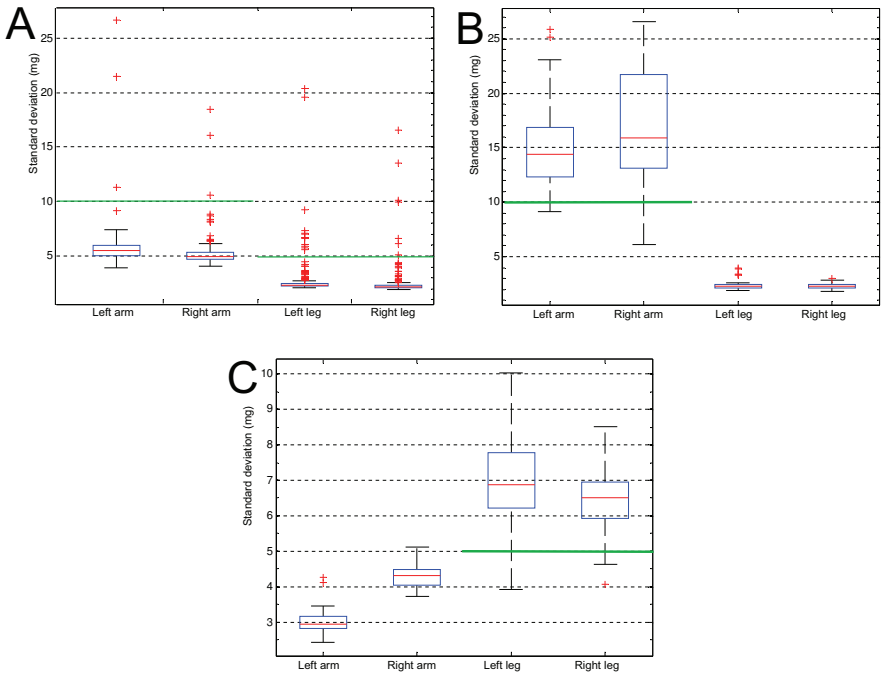


Figure 3.2: *Overview of the standard deviations within a sliding window of the simulated movements to determine the threshold for movement detection. The upper left figure (A) shows the recorded values when no movement occurs (hands on chest in rest). The upper right (B) and lower figure (C) show the values while making subtle movements with the fingers and toes, respectively. The thresholds are indicated in green on the three figures.*

3% of the original raw data has to be processed, which is a considerable reduction.

3.3 Conclusion

In this chapter the preprocessing and the movement detection based on the ACM signals is discussed. This led to a total data reduction of 97% to 98.5%. Furthermore, the segmented epochs can be used in the further feature extraction and classification of movements. To this end, the segmentation is used in Chapter 5. But it is also employed in the multimodal approach in Chapter 7 as the segmentation in the ACM data and the video data should be the same.

Furthermore, a movement detector can already be used to do a first investigation

on the data. For example the events that are abnormally long in duration or have a large maximum or average amplitude can indicate an event that needs further investigation. Also the number of movements during the night can provide information on the sleeping behavior of the patient. To this end, it is implemented in the stand-alone home monitoring system that is proposed in Section 8.2.2.

Chapter 4

Preprocessing and movement detection - Video

In the previous chapter the preprocessing and movement detection based on ACM signals was discussed. This chapter serves the same purpose, except that video data is used instead. Furthermore, as extension to the movement detection, it is investigated whether the respiratory signal can be derived from the video images. This chapter is organized as follows: first an introduction is given about the background of movement detection in video with a short overview of the literature in Section 4.1. Then, the approach is explained in Section 4.2 including the calculation of the optical flow and an overview of the different measures to analyze the performance of the proposed algorithms. Section 4.2.4, 4.2.5 and Section 4.2.6 show the results and give the interpretation of these results in the discussion, respectively. To conclude, Section 4.3 elaborates a little more on the detection of respiration signals in video. Section 4.4 finishes with a conclusion.

4.1 Introduction

Section 1.3.3 already introduced the advantages and challenges of video detection. In this chapter, we go more into detail in the motion detection step of the video processing. Motion detection or segmentation in video can be done in different ways. Some of the most used methods are *temporal differencing*, *background*

modeling and subtraction and the calculation of the motion vector field by *optical flow* [87, 95, 101, 134, 189].

Temporal differencing, as for example used by Dai and Leng [55], is essentially calculating the difference image of video frames at two different time instances, it shows the change in intensity in each pixel between these two time instances. It is computationally fast, but does not derive extra information from the video sequence. The method is adaptive to slow changes in illumination and dynamic environments [87], but typically only the edges of the subjects are detected which generates holes inside moving entities [95] and non-moving objects remain in the background.

Background modeling and subtraction is a popular method for motion segmentation, especially in situations with a more or less static background. Several different approaches are proposed in literature, mainly with different ways of modeling the background. These backgrounds range from temporally averaging the images, to more sophisticated methods as temporal median filtering [119] or modeling every pixel in the background by a mixture of Gaussians [163].

Other methods can be based on *point trajectories* or *spatio-temporal filter responses* [176]. Wang *et al.* [189] mentioned in their survey some techniques which could be found in the literature such as an extended *Expectation-maximization* algorithm, and a morphological algorithm for scene change detection.

Finally, the optical flow algorithm calculates a motion vector in each pixel of the video image with the according estimated movement. Optical flow based motion segmentation makes use of the characteristics of these flow vectors to detect moving regions in the image sequence [87, 175]. The optical flow calculation is computational intensive and sensitive to noise. Furthermore it is important to supply video with rigid objects and a good intensity of the image. Otherwise the algorithm will respond less accurately [194], but the optical flow calculation provides more information than the former methods, such as the direction and intensity of the motion in each pixel. This information can be used in the later processing steps by investigating the direction and amplitude of the moving parts in the image, therefore we apply this method on our data.

The most common implementations of the optical flow algorithm are the *Horn-Schunk* method [86] which makes use of a smoothness constraint to solve the aperture problem and the *Lucas-Kanade* [123] method which assumes that the flow in small sections of pixels is the same. Barron *et al.* [9, 10] give an overview of different optical flow implementations and their performance in terms of angular error, standard deviation and the density of measurements. The different techniques

can be divided into *differential* techniques, *region-based* matching, *energy-based* methods, *phase-based* techniques. The Horn-Schunck and the Lucas-Kanade method are both categorized as differential techniques. These techniques compute the velocity from the spatio-temporal derivatives of the image intensity [10, 98]. But also other research groups such as Uras *et al.* [179] have implemented variations of this technique.

Region-based techniques are more appropriate when the calculation of the differentiation is inaccurate for example because of noise in the image. These techniques define the velocity by finding the best match between image patches, similar to block matching as used by Burt *et al.* [26], Little and Verri [117], Anandan [5] and Fuh and Maragos [72].

Energy-based methods use velocity-tuned filters in the frequency domain. By tuning these different filters according to each movement direction, the output energy of each filter can be used to determine the flow vectors in the video, as for example used by Adelson and Bergen [1], Barman *et al.* [8] and Heeger [83].

Phase-based methods, finally, make also use of the filtering in the frequency domain like the energy-based methods, but this method uses phase information instead of the amplitude of the filter response [67].

The optical flow calculation is not yet a solved problem [174], although the last two decades more accurate methods of calculating the motion flow were developed, such as Black and Anandan [16] and Brox *et al.* [25]. The approach of Brox *et al.* produces a highly accurate optical flow field which is robust to noise and which preserves discontinuities at motion boundaries. But all the optical flow methods still make certain assumptions such as a uniform illumination, so in situations where these assumptions would not hold, the correct calculation becomes difficult [174]. Also, large or fast motions are difficult to track. Recent years also motion estimation methods that differ from the optical flow were developed which look promising, such as *Optical Snow* [110] or *Particle Video* [157]. The Optical Snow method tries to model multiple motions at each pixel whereas the Particle Video combines particle based motion estimation with dense optical flow.

Although movement detection is most studied in surveillance, some progress has been made towards the detection of epileptic seizures in adults [118] and neonates [96]. Liu *et al.* [118] detected changes between video frames in a certain sequence of video data from epileptic seizures. This is comparable with the detection of movement. They developed an accurate and robust algorithm for change detection. To visualize this, a *Change Detection Map* (CDM) was constructed which shows the detected motion in the subsequent images. The algorithm was tested both on

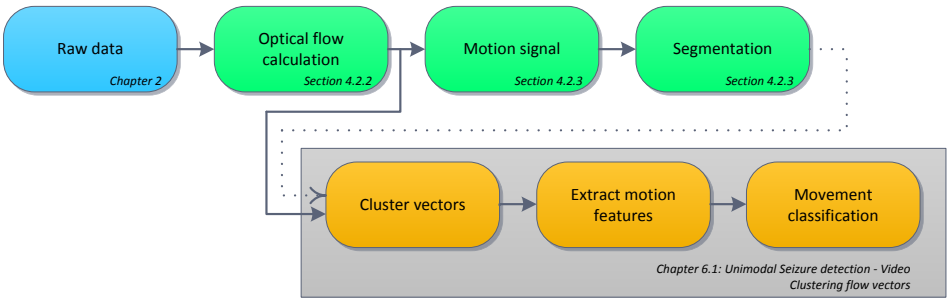


Figure 4.1: *Schematic overview of the different parts of the algorithm.*

MATLAB generated images as well as on real data of epileptic patients, however no information about sensitivity and specificity of their method are given in their paper. The CDM computed in this algorithm does not bear any information about the intensity of the movement. In contrast, the optical flow does give this information in our approach.

Figure 4.1 gives an overview of the different parts of the algorithm that are described in this chapter and the relationship with other parts of the thesis. From the raw video data, the *optical flow* field is calculated. This calculation of the optical flow vectors is also used later on in Chapter 6 when epileptic seizure events are distinguished from normal movement using video. In each pixel, the amplitude and direction of the motion vector is returned. From these vectors *a motion signal* is derived which reduces as much as possible noise related to non-movement such as camera noise or compression artefacts. This motion signal gives an indication of the intensity of the captured movement. This signal is used to determine the motion segments, which primarily reduces the data, but also delimits the movement epochs. Generating this information can also be medically relevant, as a movement profile of the patient can be drawn up.

Another interesting research question is whether we can derive breathing information from the video signals. This would provide the clinical staff with extra information which could aid the seizure detection, or it could be used for more general sleep monitoring. In Section 4.3.2 this part is more elaborated. First tests point out that in some cases the breathing of the patient can be monitored by using video data. The results are compared with measured data from respiration belts as a ground truth.

4.2 Movement detection

In this part of the data processing we make use of the optical flow algorithm. However, the results of this method are not only used for the detection and segmentation of movement epochs, but also to gather more information about the size and direction of the movement for classification of movement into normal or seizure movement.

4.2.1 Video acquisition

In Chapter 2, an overview of the different acquired datasets was already given. For this part of the research we make use of the datasets that were acquired in the University Hospital Leuven (Belgium). Both simulation data, described in Section 2.2.1, and real data, described in Section 2.3.1, from a database of epileptic seizures from the Department of Child Neurology, were used. The length of the events from the real data ranges from 1 to 304 seconds. An overview of the video sequences was already given in Table 2.2.

The real dataset consisted of a 1 hour 48 minutes and 44 seconds recording with 73 movement events from which 11 were convulsions. This dataset is a sequence of 11 recordings from 8 patients.

4.2.2 Optical flow algorithm

The optical flow algorithm calculates the velocity vectors in each pixel of every frame of a video sequence, based on the change in the pixel intensities over time. It gives a 2D representation of 3D motion [99]. However, Beauchemin and Barron [12] already stated in their work that the optical flow is only an *"approximation to the image motion"*. Or using a more formal phrasing, the computation of the optical flow is ill-posed [174], as no unique solution exists.

For the calculation of the velocity vectors, first the assumption is made that the intensity of a pixel from the image of an object does not change as the pixel moves from frame to frame, this is the so called *brightness constraint* [68]. This means that we assume that the brightness of a pixel does not change as it is tracked from frame to frame [22, 100].

The image brightness at point (x, y) in the image at time t can be denoted as $\mathcal{I}(x, y, t)$. The brightness at a certain point is considered constant in time as given

in the brightness constraint. If the point moves dx in the x-direction and dy in the y-direction in a time dt, the brightness of the point stays constant [86], by which

$$\mathcal{I}(x, y, t) = \mathcal{I}(x + dx, y + dy, t + dt). \tag{4.1}$$

The right-hand side can be expressed as a Taylor series [162]:

$$\mathcal{I}(x + dx, y + dy, t + dt) = \mathcal{I}(x, y, t) + \mathcal{I}_X dx + \mathcal{I}_Y dy + \mathcal{I}_T dt + O(\delta^2). \tag{4.2}$$

$O(\delta^2)$ represents the second and higher order terms, and $\mathcal{I}_X$, $\mathcal{I}_Y$ and $\mathcal{I}_T$ are the partial derivatives of $\mathcal{I}$. If dx, dy and dt are very small, the higher order terms in equation 4.2 vanish and this equation can be rewritten as

$$\mathcal{I}_X \frac{dx}{dt} + \mathcal{I}_Y \frac{dy}{dt} + \mathcal{I}_T = 0, \tag{4.3}$$

with $u = \frac{dx}{dt}$ and $v = \frac{dy}{dt}$ are the velocities in the horizontal and vertical direction, respectively. These values represent the movement we want to calculate. For this equation to be solved, extra constraints have to be determined, because it is ill-posed; u and v cannot be computed from this one equation. This is known as the aperture problem [178]. This means that only the velocity component perpendicular to the iso-brighness contours can be calculated (in the direction of the brightest gradient). Figure 4.2 gives an intuitive representation of the aperture problem. As the line in the left part of the image moves from the left to the right, locally, it is undetermined to which exact location every point on this line moves. In the right part of the figure, the square moves to the right. However, locally the movement direction of the sides of the square are unknown and only the movement perpendicular to the sides is seen [12]. Only in the corners of the square, the correct movement can be calculated as there is enough structure at this point.

To solve this problem the Horn-Schunk method [86] makes use of a *smoothness constraint*. It is presumed that the velocity of neighboring points is more or less the same, not every point in the image can move independently. This is easy to understand because normally the moving points belong to larger objects that undergo rigid motion or deformation. The smoothness constraint is introduced in the algorithm by minimizing Equation 4.4 (the squared error quantity),

$$E^2(x, y) = (\mathcal{I}_X u + \mathcal{I}_Y v + \mathcal{I}_T)^2 + \alpha^2(u_X^2 + u_Y^2 + v_X^2 + v_Y^2), \tag{4.4}$$

where u_X^2, u_Y^2, v_X^2 and v_Y^2 are the square of partial derivatives of the optical flow component. The first term is the brightness constraint from Equation 4.3, the

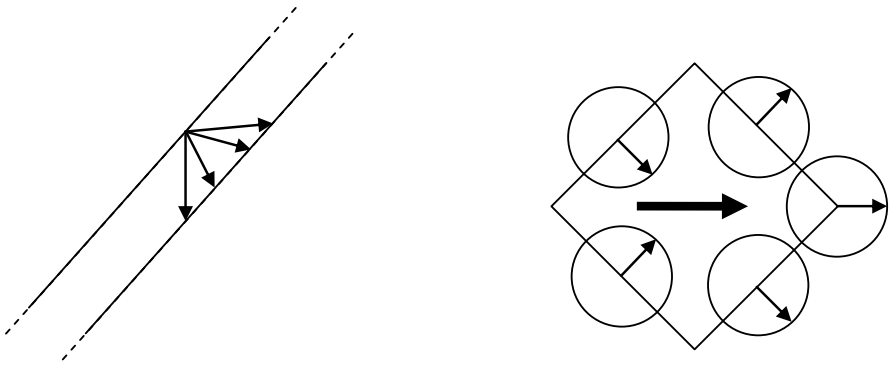


Figure 4.2: *This figure visualizes the aperture problem. The left image shows a line that moves from the left to the right. Locally, it is undetermined to which exact location every point on this line moves. The right images shows a square that moves to the right. Also in this figure, due to the aperture problem, the movement the square can only be determined perpendicular to the edges of the square. In the corners however, enough structure is available to determine the correct motion.*

second term is the smoothness criterion. Alpha is the regularization parameter to specify the influence of the smoothness term, compared to the brightness constraint.

Horn and Schunck suggested to use the Gauss-Seidel method [84] for calculating the motion in each point and minimizing Equation 4.4. This is an iterative method which gives the following equations:

$$u^{(k+1)} = \bar{u}^{(k)} - \mathcal{I}_X \frac{\mathcal{I}_X \bar{u}^{(k)} + \mathcal{I}_Y \bar{v}^{(k)} + \mathcal{I}_T}{\alpha^2 + \mathcal{I}_X^2 + \mathcal{I}_Y^2}, and \tag{4.5}$$

$$v^{(k+1)} = \bar{v}^{(k)} - \mathcal{I}_Y \frac{\mathcal{I}_X \bar{u}^{(k)} + \mathcal{I}_Y \bar{v}^{(k)} + \mathcal{I}_T}{\alpha^2 + \mathcal{I}_X^2 + \mathcal{I}_Y^2}, \tag{4.6}$$

where k denotes the iteration number. $\bar{u}^{(k)}$ and $\bar{v}^{(k)}$ are the averages in the neighborhood around $u^{(k)}$ and $v^{(k)}$, respectively. The starting values (before the first iteration) for u and v are zero. In the iterative process, uniform regions of the image are filled in [174]. The values of the boundary region are propagated inwards with each iteration. The number of iterations needed should thus be larger than the width or height of the largest uniform region [174]. However, this depends on the application of course; in this chapter a lower number of iterations (10) could be used because for the movement detection it was not necessary to have such an accurate calculation of the flow vector field.

4.2.3 Calculating $R(t)$

Figure 4.1 already gave an overview of the different steps to deploy a movement detection algorithm. After the calculation of the optical flow field in Section 4.2.2 the obtained motion vectors are used to generate a motion signal.

The goal of the detection algorithm is to distinguish epochs with movement from epochs without movement. The optical flow algorithm calculates a vector field with a motion vector in each pixel, as shown in Figure 4.3(b) considering Figure 4.3(a) as one frame of the video sequence of a child that moves its left arm. For the optical flow calculation for the detection of movement it suffices to take a lower number of iterations compared to the seizure detection. This produces a less accurate vector field, but motion is still detected and the processing time is reduced. For the further calculations, the magnitude of each motion vector is used. This gives video images with each pixel representing the intensity of the movement (the ℓ^2-norm of the vector). The result of such an image is shown in Figure 4.3(c). As visible in this figure, several spurious pixel intensities are shown which are noise which may be introduced by the camera or by the MPEG-compression. Note that Figure 4.3(c) is amplified with a factor of 15 to emphasize these spurious patches. To eliminate the noise, only a small number of pixels is taken into account. The pixels with the highest movement intensity are maintained. The result of this image is shown in Figure 4.3(d). The number of pixels with the highest amplitude is a free parameter that needs to be determined as explained in Section 4.2.5.

The next step takes the mean of the selected vector magnitudes which results into the 1D output signal (further referred to as $R(t)$) indicating the motion in the video sequence. This signal is further examined to find the epochs with movement and it is investigated what the influence of different settings is on $R(t)$. Figure 4.5 gives an example of an $R(t)$ signal of a movement. In the first second of this signal, there is no movement. At 1.5 seconds a myoclonic jerk occurred. To detect motion, we can easily threshold this 1 dimensional signal.

One part of this chapter studies the effects of a change in the image acquisition settings (spatial resolution, frame rate and video compression) and environmental settings (change in illumination and camera point of view) on the derived motion signal. The following enumeration gives an overview of the different effects we tested:

- **Impact of downsampling in space:** the datasets with scratching the nose (small movement) and a myoclonic jerk are used in this investigation. The



Figure 4.3: *Image from acquired video data(a) Vector field calculated by the optical flow algorithm (b) Optical flow image where each pixel indicates an optical flow vector's amplitude with an amplification of 15 to emphasize the spurious pixels (c) Image after taking 0.06% pixels with highest intensity (thus vectors with highest amplitude) (d).*

original dataset of 640×480 pixels is downsampled to 320×240, 160×120 and 80×60 pixels.

- **Impact of downsampling in time:** for this investigation, the dataset with the myoclonic jerks (fast movement) is used as well as the dataset with scratching the nose and moving the head (small movements). The original dataset with 30 fps is downsampled in time to 15, 10, and 5 fps.
- **Impact of compression:** to study the impact of the compression, the dataset with the scratching of the nose was used to investigate the impact of the compression on small movements.
- **Impact of illumination:** the dataset with the stretching of the arms is used in this comparison. This includes large (movement of arms) and small movements (during the stretching of the arms).

- **Impact of camera viewpoint:** here the dataset with the simulated overstretch of the body is used. This motion is significantly more intense from one camera point of view compared to the other.

For analyzing the performance of the algorithm under these different acquisition and environmental settings, we introduce the *Movement vs. Non-Movement Ratio* (MNMR) and $max(R(t))$.

The MNMR is defined as

$$MNMR = \frac{max(R(t))_{movement}}{max(R(t))_{non-movement}}, \tag{4.7}$$

where $max(R(t))_{movement}$ and $max(R(t))_{non-movement}$ represent the maximum value of $R(t)$ in a movement epoch and in a non-movement epoch, respectively. A high MNMR for movement epochs facilitates the discrimination by means of thresholding between movement and non-movement epochs. Therefore this measure should be as large as possible and at least be larger than 1.

This movement detection algorithm is a data reduction step, and the sensitivity should be as close to 1 as possible because we do not want to lose significant data. Because we want as less unnecessary data as possible, the PPV should also be as high as possible.

Temporal clustering

Next to the classification by the algorithm, the video epochs are all manually labeled by a human labeler, which counts as a ground truth. For the comparison of the events, the automatic detection and the manually labeled events are first clustered separately. If two events lie less than 4 seconds apart, the events are merged. The next step adds 2 seconds before and after each event to overcome the inaccuracy of the human labeler. The clustering is shown in Figure 4.4(b).

We compare the labeled dataset with the result of the automatic detection algorithm to calculate the TPs, FNs and FPs. When the algorithm detects movement during a labeled event, this epoch is considered as a true positive. If we detect no movement during a labeled event, the epoch is considered as a false negative. This is visualized in Figure 4.4(c).

Sometimes one detected event can last for much longer than the human labeler observed. It would be preferable to detect such an event as a false positive. That

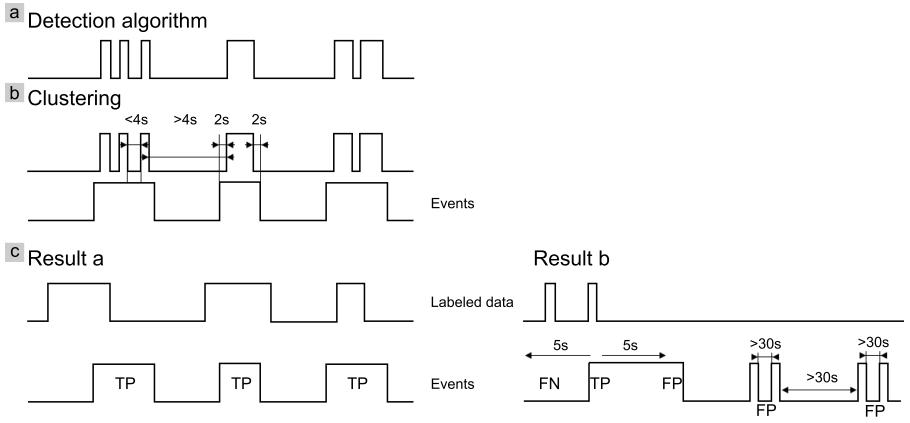


Figure 4.4: *Clustering of the events and indicate them as TP, FP or FN. (a) shows the events detected by the algorithm, (b) shows the clustering which is done afterwards. (c) gives an overview of how the events are labeled in two examples.*

is why false positives are calculated in the following way. First the events labeled by the human labeler are expanded by 5 seconds in the beginning and at the end. Every event detected by the algorithm outside these regions is considered as a false positive.

In one of the sequences a CRT computer screen was visible in the background, causing a flickering motion in the video. This led to an increased movement signal resulting in false positives. When the threshold is set higher in this or other similar cases, it is possible that only a few spikes cross the threshold resulting in several short false positives in stead of one large false positive. This leads to a counterintuitive situation where the number of false positives increases when the threshold is set higher. To overcome this problem, false positives within 30 seconds are merged.

Parameter optimization

We need to optimize two parameters. The first one is the percentage of highest movement vectors n_{vec} (or fraction of highest movement vectors) that we apply to calculate $R(t)$.

The second parameter we want to optimize is the threshold τ, we have to set on $R(t)$ to get a good distinction. We have worked out two approaches for the

threshold calculation, as will be explained in the following paragraphs.

First approach: Global threshold

In the first approach we try to find an optimal fixed threshold to distinguish between movement and non-movement for all the different video sequences. If such a threshold exists, it is not necessary to find an optimal threshold for each sequence individually, which is the most convenient solution.

Second approach: Variable threshold

A second approach to determine the optimal threshold is to make the threshold dependent of the noise (i.e. non-movement epochs) in the sequence. Therefore a short epoch, which does not contain movement, at the start of the sequence is chosen. If we take Figure 4.5 as an example, the first second could be used for calculating the variable threshold, as no movement occurs in this part. From this epoch the mean μ and standard deviation σ are calculated. This gives a representation of the non-movement. The threshold τ_{var} is determined by adding a multiple α of the standard deviation to the mean of the epoch, which gives: $\tau_{var} = \alpha \times \sigma + \mu$. This way an adaptive threshold is created. We use the same multiplication factor for each sequence. This approach is tested on the same datasets as the first approach.

4.2.4 Results on choosing acquisition parameters

In a first step, we investigate $R(t)$ in a video sequence under different environmental and acquisition parameters. To this end, we recorded simulation videos representing different epileptic and non-epileptic movements, of which an overview is given in Table 2.1. The influence of the image acquisition settings (spatial resolution, frame rate and video compression) and environment settings (change in illumination and camera point of view) on the optical flow algorithm is already shortly introduced in Section 4.2. In these tests we set the n_{vec} parameter to 0.06%. In Section 4.2.5 we go more into detail on the selection of the optimal value of this parameter. In this section we discuss the results and draw conclusions regarding the most optimal acquisition settings.

Impact of downsampling in space: When the video sequences are downsampled in space both $R(t)$ and the MNMR become smaller. This means that the vectors

Table 4.1: *MNMR and R(t) at different downsample ratios in time expressed in frames per second (fps), with the percentage of pixels fixed at 0.06.*

MNMR	30fps	15fps	10fps	5fps
Scratching nose	2.4572	3.5291	3.6661	2.3860
Myoclonic jerk	82.0994	90.3868	103.0590	32.9629
Moving head	64.0910	30.2057	28.6357	22.8782
$\boldsymbol{max(R(t))_{non-movement}}$				
Scratching nose	0.0363	0.0479	0.0729	0.0804
Myoclonic jerk	0.0350	0.0436	0.0455	0.1339
Moving head	0.0143	0.0398	0.0598	0.0540
$\boldsymbol{max(R(t))_{movement}}$				
Scratching nose	0.0892	0.1692	0.2674	0.1919
Myoclonic jerk	2.8766	3.9406	4.6922	4.4143
Moving head	0.9146	1.2031	1.7111	1.2358

originating from a movement are decreasing more than the noise. We investigated two sequences, namely scratching the nose and a myolonic jerk. The downsampling to 160×120 and 80×60 of the scratching nose sequence resulted in a reduction in the MNMR of 0.04% and 29.26%, respectively. For the myoclonic jerk the reduction was 16.25% and 83.62%. For the scatching nose sequence, the reduction becomes considerably larger only at a downsampling to 80×60 pixels. A greater downsampling in space than 320×240 pixels, however, resulted in a considerable lower MNMR for the myoclonic jerk. For this reason a resolution of 320×240 is used in this paper.

Impact of downsampling in time: When the signal is downsampled in time, $R(t)$ increases. This is visible in Table 4.1. In this table the values for $R(t)$ during movement and non-movement and the MNMR are shown. Notice however that this trend is not always true when the signal is downsampled to 5 fps. When the difference between an object is too large between two consecutive frames, the optical flow calculation can be inaccurate. The increase of $R(t)$ can be explained by the fact that if an object, or body part moves at a certain speed, it has moved a longer distance between each frame if we sample at a lower frequency. Notice however that aliasing effects can have an influence on the correctness of the optical flow calculation. Also, due to aliasing, specific patterns in the $R(t)$ curve as shown

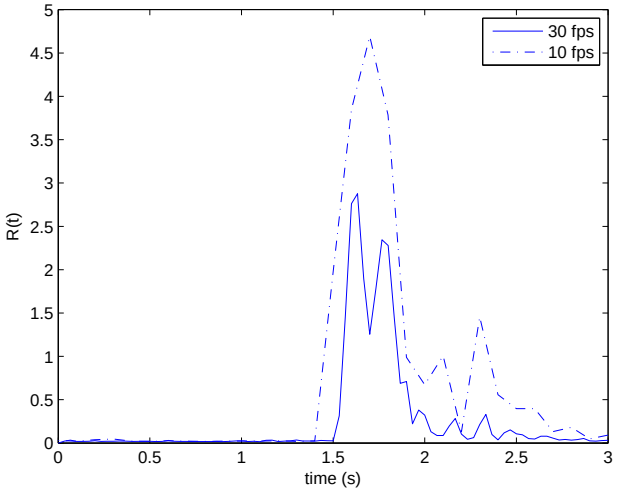


Figure 4.5: *Downsampling in time of myoclonic jerk.*

in Figure 4.5 may be lost. The two spikes in the signal represent the fast movement of the arms in the upward and downward direction. This pattern is not visible in the downsampled signal. Such a pattern could be interesting for the detection of certain epileptic seizures, but does not play a role in the discrimination between movement and non-movement.

As we want a sufficiently high MNMR, a downsampling to 5 fps may be unwanted, as the MNMR drops substantially in the different sequences for this frame rate, also shown in Table 4.1.

Another observation is that the MNMR value for subtle movements is clearly smaller than for bigger movements, but no clear trend can be observed regarding the MNMR when the video is downsampled in time.

As a conclusion we can say that the temporal sample frequency should not be lower than 10 fps. In our dataset this means that we can for example downsample by a factor of 2 to 12.5 fps.

Impact of illumination: When the illumination is lowered, $R(t)$ stays more or less the same for movement epochs but $R(t)$ increases for the non-movement epochs, which can mainly be explained by the increased noise in the video. In this test the recordings were more difficult to compare as they were two separate recordings taken under different circumstances of illumination.

Impact of compression: As the video data that was recorded in UZ Leuven was compressed with an MPEG1 compression, we applied the same compression to our simulated data. A typical dataset of 300 MB was compressed to 650 kB. We noticed that the compression had little influence on $R(t)$ during large movement epochs, but $R(t)$ was significantly increased during non-movement epochs. This can also be explained by the increased noise that has a relative higher influence when no movement occurs in the scene. During very small movements, some information can be lost.

Impact on camera viewpoint: The amplitude of the signal is larger if the movement lies in the plane parallel with the camera. The different view points of the cameras investigated are shown in Figure 2.2. One camera view in Figure 2.2 is taken from the foot end of the bed towards the head, and the other view is a top-view of the bed. When, for example, investigating $R(t)$ during the overstretching of the arms, more pixels are involved from the viewpoint from the foot end, hence $R(t)$ is larger.

As a conclusion on choosing the acquisition parameters, we can say that the amplitude of $R(t)$ is influenced by several parameters. Although a downsampling in time and space makes the algorithm less demanding with respect to computational load, there are some disadvantages to it. The MNMR can decrease which makes it more difficult to discriminate between movement and non-movement epochs. Furthermore the specific wave shape of $R(t)$ can be lost due to downsampling in time, as was illustrated in Figure 4.5 for a myoclonic jerk.

Furthermore we observed that the MPEG-compressed video recordings mainly increased $R(t)$ of non-movement epochs. Hence it becomes more difficult to distinguish between small movements and non-movement epochs. The same holds for the illumination: the $R(t)$ measure increases for non-movement epochs due to a reduction in contrast. Again, it becomes more difficult to distinguish between small movements and non-movements epochs. As the movement detection is mainly carried out during the night, we have to take these poor contrast images in consideration. One way to partly overcome this problem is to use good quality infra-red cameras and infra-red light sources. This will yield more contrast in the video-recordings which again leads to better MNMR.

With respect to the camera position we have observed that in plain movement yields a higher $R(t)$ than out of plain movement. We need to be aware that the camera position is often not free to choose and an optimal position for one type of movement may not be the optimal position for another type of movement.

To conclude, we can say that with the above mentioned findings we can make recommendations for the detection of movement during sleep with the optical flow

algorithm from video recordings. We believe that a reduction to 320x240 pixels and 10 fps is acceptable for the detection of movement. This will significantly reduce the computational load which allows for real-time execution of the algorithm. Note however, to distinguish between normal movement and epileptic seizures the video should not be downsampled in time. Furthermore we recommend to use lossless compressed video recordings or video compression with a minimal loss of image quality and to use camera systems which allow good contrast images even in dark settings. Another strategy is to compress the video data after the algorithm is executed when the movement detection algorithm is executed in real-time.

4.2.5 Results on identifying τ and n_{vec}

The performance of the algorithm is calculated for different values of the two parameters we want to optimize, namely τ and n_{vec} as introduced in Section 4.2.3. For the different combinations of these parameters, we have evaluated the performance to be able to select the parameter combination that performs best. Therefore we tested different values of n_{vec}, between 0.0013 (one vector) and 100 (all the vectors). The tested values for τ were chosen out of a range between 0.005 and 1.5. The multiplication factor α of the standard deviation, calculated on a sequence of 50 samples, for the variable threshold ranges between 1.5 and 30. For a proper evaluation, the dataset is split up into three subsets containing all more or less the same number of movement events. These subsets are used to perform a 3 fold cross-validation. From each training set the most optimal parameter combination is determined, and used in the test set. The performance of each test set is then given by the sensitivity and PPV. The results for both approaches are discussed below and can be found in Figure 4.6 and Table 4.2.

First approach: Global threshold

On each training set the Precision-Recall curves (PPV-sensitivity) are calculated as shown in Figure 4.6(a) and (e). The most optimal point on each Precision-Recall curve has a sensitivity of one, as we do not want to miss important events, with a maximal PPV. For the first training set for the fixed threshold, this point is indicated with a green circle on the ROC-curve shown in Figure 4.6(a). The matching parameter combinations of this point are shown in Figure 4.6(b) as white pixels. Each point in this figure represents a combination of a certain τ (y-axis) and n_{vec} (x-axis). Notice that multiple combinations of the parameters yield optimal performance. The axes origin for Figure 4.6b-d and f-h is situated in the left upper

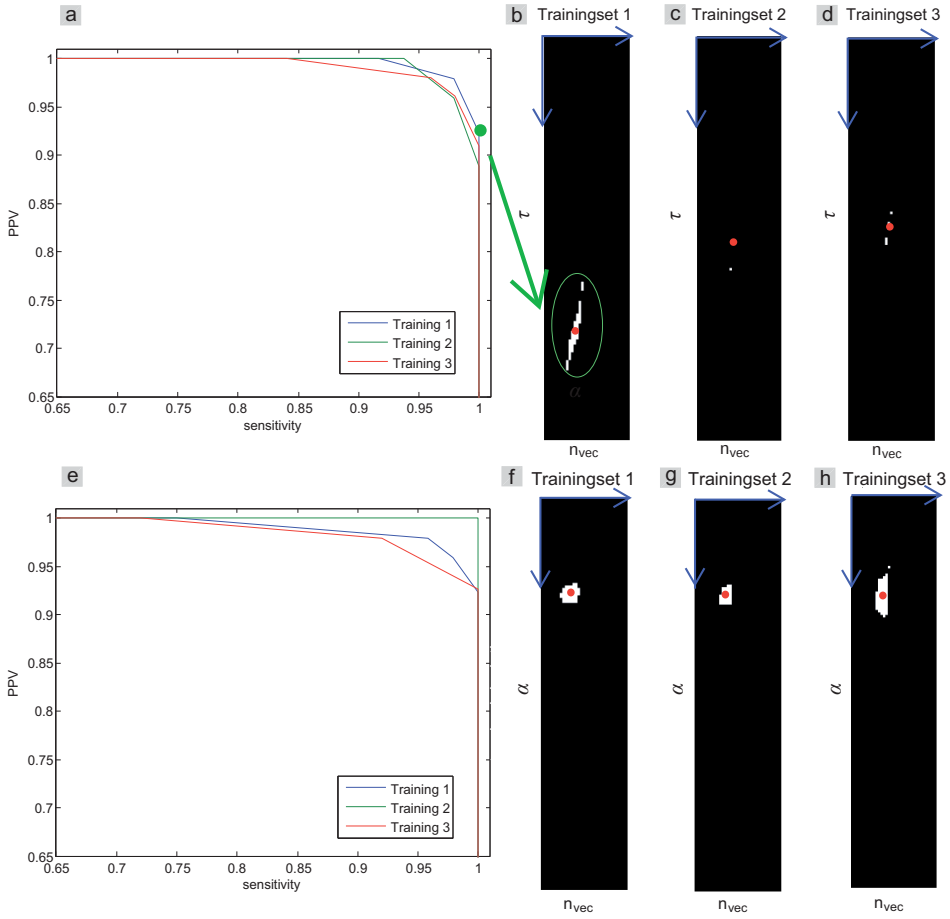


Figure 4.6: *Evaluation of three training sets for the fixed (in a-d) and variable (in e-h) threshold. Figure a and e show the envelope for the Precision-Recall curves (where the precision and recall are synonyms for the PPV and sensitivity, respectively) for the different training sets. Notice that the x- and y-axis are zoomed in for a better visualization. For each optimal point in the Precision-Recall curve (a sensitivity of one and the highest possible PPV) multiple parameter combinations are possible. The optimal point in the first training set indicated by a green mark in (a) corresponds to the white pixels in (b). The parameter combinations used in the test set are marked with a red circle in figure b-d and f-h.*

Table 4.2: *Performance of the training and test sets of the 3-fold cross-validation. The parameters used in the test set are obtained from the combination in the training set with a sensitivity of 1 with the highest obtainable PPV, which is also indicated in this table.*

Fixed threshold	n_{vec}	τ	Test set	
			sens	PPV
Training 1 (PPV: 92.00%)	0.06	0.6050	92.00%	100.00%
Training 2 (PPV: 89.00%)	0.20	0.4200	100.00%	89.29%
Training 3 (PPV: 91.00%)	0.30	0.3850	100.00%	79.31%
Overall	/	/	97.33%	89.53%
Variable threshold	n_{vec}	α	Test set	
			sens	PPV
Training 1 (PPV: 92.00%)	0.06	5.4	100.00%	100.00%
Training 2 (PPV: 100.00%)	0.06	5.4	100.00%	86.21%
Training 3 (PPV: 93.00%)	0.06	5.6	100.00%	100.00%
Overall	0.06	5.5	100.00%	95.40%

corner. The parameter combination used in the test set is indicated with a red circle in this figure, and the corresponding values can be found in Table 4.2. When we look at the overall sensitivity and PPV we observe a value of 97.33% and 89.53%, respectively.

Second approach: Variable threshold

The evaluation of this approach is similar to the first approach, but the values for τ in Figure 4.6 are replaced by α. For each training set, multiple points are obtained with a sensitivity of one with a maximal PPV. Furthermore multiple of these parameter combinations are the same in the three training sets, so the highlighted points in Figure 4.6 overlap. As shown in Table 4.2 the best values for each training set are nearly the same. They give good results as in two test sets a sensitivity and PPV of 100.00% is reached and in the other test set a PPV of still 86.21% is reached (when the sensitivity is 100.00%). The false positives here are caused by a flickering computer screen and a glitch in the video. The average

sensitivity and PPV over the three folds is 100.00% and 95.40%, respectively, which is an increase in performance over the first approach.

4.2.6 Discussion

The goal was to develop a detection algorithm based on video recordings to discriminate between movement and non-movement in nocturnal recordings from patients with epilepsy. We found that the usage of a variable threshold leads to better results in comparison to a fixed threshold, as the sensitivity in the three test sets when using the variable threshold was 100.00%, with a corresponding PPV of 100.00%, 86.21% and 100.00%. Furthermore, there was an overlap in the optimal points in the three test sets, which indicates that this value is robust for different datasets.

In Figure 4.6(b) and also to some extent in (c) and (d), a linear trend is visible. This can be explained as follows. When the percentage of the highest motion vectors increases (more vectors are taken into account) the resulting signal $R(t)$ decreases as more vectors with a lower amplitude are preserved. This will lower the optimal threshold and results in the correlation visible in Figure 4.6(b)-(d).

A different camera setup may result in a different ideal fixed threshold. The variable threshold overcomes this problem. It will contain less false positives because the threshold is more adapted to the noise. A way to anticipate to changing environmental parameters during the acquisition (e.g. sunrise, which may influence the noise level) when using the algorithm in a real setting, an epoch without movement can be taken each hour and the threshold can be adapted accordingly. However the approach with the variable threshold requires that the selected reference epoch is a good representation of the non-movement epochs in the sequence.

For a further reduction of the false positives we could integrate additional intelligence such as a human body detector, and only inspect $R(t)$ in this region. This would eliminate the false positive due to the flickering screen.

Our algorithm gives a higher PPV compared to the movement detection based on ACMs by Nijsen *et al.* [135]. However, it should be noted that some very subtle motions which occur under the blankets may not have been noticed by the expert in the video data, and hence ACM data can be useful in addition to the video data. A combination of both detection methods may result in a system with a sensitivity of one and a higher PPV than the one reached in each system separately for the detection of epileptic seizures.

A real-time implementation of the optical flow algorithm is necessary in the future, for example when an alarm has to be generated when a severe epileptic seizure occurs. Rowekamp *et al.* [155] already implemented the Horn-Schunck optical flow algorithm in hardware. Their implementation reached a throughput of 50 frames per second for images with a resolution of 128×128. Chase *et al.* [31] compared a tensor-based implementation of the optical flow on an FPGA with the implementation using a GPU. On the Yosemite sequence with a resolution of 640×480, both implementations had a similar performance of about 300 frames per second. We have applied our method also in real-time, to integrate it in a stand-alone detector, which is discussed in Section 8.2.2. To calculate the optical flow, however, we reduced the number of iterations of this calculation to 10 to make the real-time detection possible.

We can conclude that nocturnal movement detection in epileptic pediatric patients using the optical flow algorithm is possible and a variable threshold seems to give a considerable better result than a fixed threshold, resulting in a sensitivity of 100.00% in the test sets and a corresponding PPV between 86.21% and 100.00%.

4.3 Detection of nocturnal breathing

4.3.1 Introduction

During the night when the observed person is not moving, still subtle motions can be observed, such as the beating of the heart or respiration. In this part of the thesis we investigate if it is possible to detect the breathing signal during sleep in certain circumstances making only use of the video data.

A contactless way of monitoring respiration can have a substantial added value in the monitoring of epilepsy or sleep, as it is an important physiological parameter as discussed in Chapter 1 with some correlation with epilepsy.

Other research groups have already monitored respiration with video, for example to detect sleep apnoea [190]. However in this paper, only the apnoea were detected without any knowledge of the normal respiration frequency. Recently, Poh *et al.* [142] used the color information from the face acquired by a webcam to derive several physiological parameters such as the heart rate variability (HRV) and respiratory rate. Kuo *et al.* [108] used a near infrared camera for their respiration measurement. From the recordings, the segments without body movement were

analyzed on respiration making use of the Horn-Schunk optical flow calculation. In AL-Khalidi *et al.* [3] thermal images are used.

4.3.2 Method

As in the previous sections, we make use of the Horn-Schunk optical flow calculation for determining the movement in the video sequences. After the motion vector field is determined, the motion vectors with a sufficiently high amplitude are clustered. In this section a clustering of the motion vectors is employed to derive a respiration signal from the video. We refer to Chapter 6 for a full overview and elaborative study of the different clustering techniques and segmentation of the motion vectors. In this section we only apply a clustering method for generating the respiration signal.

Data description

The amplitude and direction of the motion of the cluster in each frame is analyzed and compared with the ground truth of respiration monitoring. The four subjects we analyzed also wore a belt to monitor the thoracic and abdominal movements. The signals were sampled at 25 Hz, next to the video recordings. With this ground truth, we are able to objectively compare the obtained motion (respiration) signals. In the sequences we tested, no large nocturnal movements occurred (except the breathing), hence only one cluster was present in each frame.

Movement analysis

The motion of the chest is slow and hence the frame-to-frame displacement is small. Therefore we calculated the flow vectors over multiple frames, to get a larger movement. However, this resulted in a time delay in the derived respiration signals, but this did not influence the calculated respiration rate. In our study we used a calculation of the flow over 10 frames, sampled at 12.5 Hz. For a better detection, a region of interest focusing on the chest was chosen. Figure 4.7 shows the setup and the detected motion of the patient.

From the detected cluster in each video frame, the average movement direction and amplitude is calculated. This produces two signals in time from which we hope to be able to extract the respiration.



Figure 4.7: *This figure shows the setup used for the breathing monitoring in the left picture. The right picture below shows the detected movement indicated in the green color. This means that the torso is moving upwards, as visualized by the color direction indication in the upper right picture, where the hue of the colors outside the center indicate the direction of the movement.*

Results

Figure 4.8 shows the ground truth signal in the upper plot for both the thoracic and abdominal measurements. The lower plot shows the velocity amplitude and direction derived from the video data. The plot shows 30 seconds or 375 frames of the original video. In the plots of both measurements (video and belt) the reoccurring pattern of the in- and exhaling is visible. In the angular signal we notice a block-shaped pattern from a motion that is shifted 180°. In the amplitude signal two undulations are visible during each period. The phase change in the angular signal coincides with the local minimum in the amplitude signal. This denotes the transition from inhaling to exhaling. We further notice that the number of periods of the respiration belt signals and the video signals are the same. From this we can conclude that it is possible to derive the respiration rate from the video signal only, at least under the right circumstances (chest is visible, no other movements).

In Figure 4.8 it is also visible that there is a transition time of about one second to produce the correct angular and amplitude signals from the video. This is due

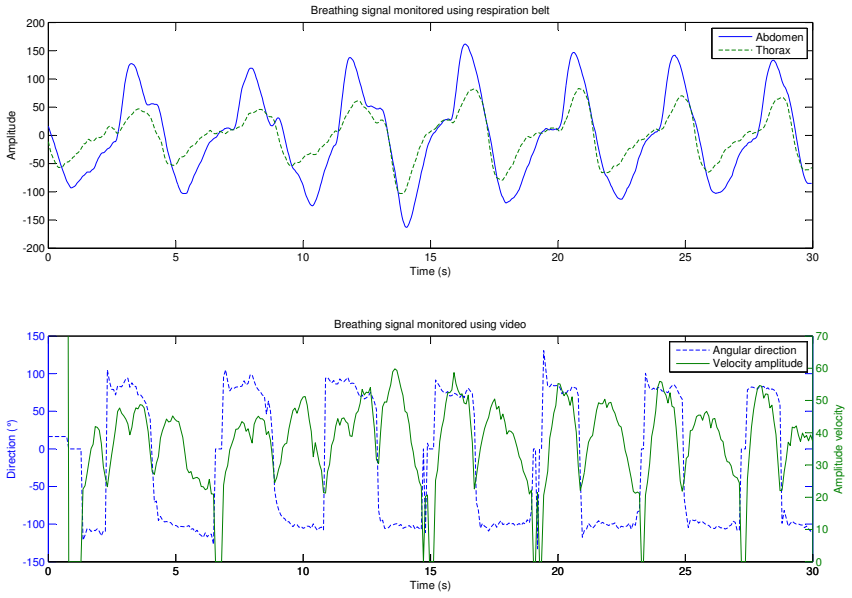


Figure 4.8: *The upper graph shows the respiration signals acquired with respiration belt that monitors the thoracic and abdominal motion. The lower graph shows the respiration signals expressed in angular direction and velocity amplitude using video data. The breathing pattern is clearly visible in both graphs. Notice that the phase change in the angular signal and the local minimum in the velocity amplitude coincide. This is the transition from inhaling to exhaling.*

to the fact that we calculate the optical flow over 10 frames. Furthermore we can see that in between two breathing cycles the amplitude and direction is zero for a short time. Because of the lack of motion between the cycles, no large enough motion vectors are calculated hence the signal becomes zero.

4.3.3 Discussion

The detection of the respiration signal only based on the video data is a valuable additional application to the motion detection. This element can be integrated in the system as described in Section 8.2.2 which can provide additional information about the patient.

In future work the monitoring of respiration should be combined with nocturnal movement. This will not be easy as the magnitude of the respiration is much smaller compared to the magnitude of movement. However, as already described in Kuo *et al.* [108], using only video for monitoring the respiration is not possible during voluntary nocturnal movement. In their system, they use a preceding movement detection step, and only calculate the breathing signal when no other movement occurs. This also fits well in our approach, as the respiration monitoring can be integrated with our movement detection algorithm.

4.4 Conclusion

In this chapter we discussed the reduction of video data by preprocessing and the detection of movements by making use of optical flow. The most important recommendations on the acquisition parameters are that a spatial to 320x240 and a temporal reduction to 10 frames per second is acceptable for movement detection. This reduces the computational load. However, we do not recommend this for seizure detection, as specific motion patterns may be lost due to the temporal downsampling.

Regarding the movement detection algorithm, the variable threshold based on noise in the video leads to better results compared to the fixed threshold, with a sensitivity of 100.00% compared to 97.33% and a PPV of 95.40% compared to 89.53%. Furthermore, the obtained parameter values for the three folds using the variable threshold are similar. This indicates they are robust for different datasets.

Furthermore, from the optical flow vectors, we can also derive extra information such as the respiration signal. And from the reduced data, we can start to extract features for the classification of sequences in epileptic and non-epileptic movement. This will be elaborated in the following chapters.

Chapter 5

Unimodal Seizure detection - Accelerometers

The previous two chapters discussed the preprocessing and the movement detection of both ACM and video data. This is only a preparation of the data to be able to do a further processing for classifying the detected movement into epileptic or non-epileptic events. The seizure detection algorithms will also be discussed separately based on the used modalities. This chapter elaborates on algorithms based on the ACM signals, whereas Chapter 6 goes more into detail on the video processing and Chapter 7 discusses an integrated detection based on both ACMs and video. In this chapter three different techniques are proposed. We start with a low-complex and straightforward technique of thresholding extracted features in Section 5.2. After that, we discuss a classification based on support vector machines (SVM) in Section 5.3 which is considered as the state-of-the-art. Finally in Section 5.4, we propose a method where the classification model is built based on only information from normal nocturnal movements. In this section the probability density function of the normal movements is estimated using a Parzen window approach and outliers are classified as seizures. Finally, in Section 5.5 we formulate our conclusions and give a comparative overview of the different methods.

5.1 Introduction

In this part of the thesis we investigate different methods for the detection of hypermotor seizures using ACMs. The group we studied consisted of 7 patients with hypermotor seizures, all between the age of 5 and 16 years. An overview of the dataset that is used is already given in Table 2.3. In total, the dataset contain 51 epileptic and 2715 normal movements, but as is visible in the overview, the distribution of seizures over the patients is not equal. Note that the movement epochs that were caused by attaching or removing the ACMs to the body were removed from the dataset in advance. These events occur at the start and end of the night and during the night when the patient has to go to the restroom for example.

The thresholding method and the method based on outlier detection on the probability density function make use of the full dataset. However, some of the patients in the dataset show too little seizures to train a *Support Vector Machine* (SVM) classifier. Therefore, the SVM approach is only applied to three patients, more specifically patients 2, 5 and 6, which had at least 7 seizures during the recording period.

When making use of the given datasets, two concerns arose: a) much more normal movements occur compared to seizure data which leads to very unbalanced data when modeling the movements of these patients and b) some nights of recording produced only one seizure (and we even excluded the ones without seizures). Because we want the proposed approaches to work in a real setting as easy as possible, we want to be able to train our algorithms with as few seizures as possible.

In the *thresholding method* we therefore make use of only one seizure to train our decision model. The imbalance doesn't have so much influence on the model generation in this approach. The *SVM approach* is much more sensitive for unbalanced data. Therefore we compare state-of-the-art techniques compensating for the highly unbalanced seizure data, to generate a classification model that is able to detect epileptic seizures. These techniques are a) the usage of a weighting factor to give a lower or higher weight to the error costs of the positive (seizure) and negative (normal) class, b) resampling the minority class (seizure data) by estimating the probability density function, and c) resampling the minority class based on the 'Synthetic Minority Over-sampling TEchnique' (SMOTE) technique [32]. The method based on the *outlier detection using the probability density function* doesn't need any seizures at all for the patient specific classification modeling. However seizure information is used for determining the optimal group specific hyper-parameters, but when implementing the system in a home situation, this step is

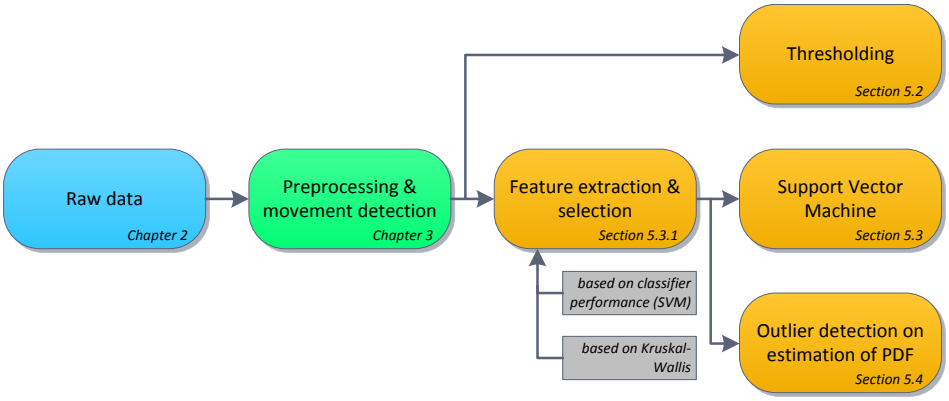


Figure 5.1: *Schematic overview of the different approaches for seizure detection based on ACMs addressed in this chapter.*

not needed anymore. Furthermore, in contrast to other methods, this approach just needs unbalanced data to work as will be explained in Section 5.4. So it is well suited for this problem.

Figure 5.1 shows an overview of the different parts of this chapter. After the preprocessing and movement detection step based on ACMs which was elaborated in Chapter 3, different approaches for seizure detection are proposed. The thresholding method, that is explained in Section 5.2, makes use of intuitive features namely the energy in the four limbs, so no extensive feature extraction is used here. The SVM approach and the classification based on the outlier detection both use features that are extracted and selected in Section 5.3.1. For the feature selection, two different approaches are considered. The first one makes use of the Kruskal-Wallis test, which is a statistical test to investigate if samples come from the same distribution. We use this for a univariate feature selection, by ranking the features according to their p-value[1]. The second approach uses a forward feature selection by using the classification performance of an SVM classifier as a measure for finding the best features. After adding features in the forward selection, the features that contribute the least to improving the performance are pruned in a backward elimination.

[1]Probability value for the null-hypothesis that all samples are drawn from the same population (or different population with the same distribution). A low p-value makes the null-hypothesis less likely and suggests that the samples are taken from different distributions and thus that the investigated feature holds a better discriminative power.

Table 5.1: *Overview of the different tests that are performed in this chapter.*

Method	Features	Patients	Tests
Thresholding	maximum per limb (4 features)	7 hypermotor patients	• n folds with $n = \#$ seizures • 1 fold = training, $n-1$ folds = test • **Table 5.2**
SVM	selection of features using: • Kruskal-Wallis test (used on 7 patients) • SVM classifier integration: on 3 patients having the most seizures, combined in 1 test	3 hypermotor patients with largest number of seizures	• 10 fold randomization with 2/3 train (67n, 4s) and 1/3 test (33n, 2e) • Also determine optimal method for coping with imbalace • Results using Kruskal-Wallis features: **Table A.1** • Results using SVM features: **Table 5.5** • 2 extra tests: • first: influence of class imbalance (**Table 5.7**) • second: influence of training set size (**Table 5.6**)
outlier detection	selection of features using: • Kruskal-Wallis test (used on 7 patients) • SVM classifier integration: on 3 patients having the most seizures, combined in 1 test	7 hypermotor patients	• Group specific determination of band width (β) • 10 fold randomization with 2/3 train (2/3 of all normal movements resampled to 1000) and 1/3 test (33n, 2e) • Results using Kruskal-Wallis features: **Table 5.8** • Results using SVM features: **Table A.2**

After the feature selection step, the classification method using SVM is explained in Section 5.3. As discussed in the previous paragraph, different methods to overcome the imbalance of the data are proposed. In a last proposed method, discussed in Section 5.4, the probability density function of the normal movements is estimated using a non-parametric Parzen window approach, where outliers are considered to be seizures.

In this chapter, a broad set of tests are performed. To get a better and more structured view of these tests, Table 5.1 gives an elaborated overview on the methods, the features that are used, the patients and the test outline.

5.2 Thresholding

As a first approach we tested a basic ad hoc method by thresholding intuitive features. This approach is considered as a baseline to identify the obtained performance when little intelligence is added to the classifier. When three of the signals of the limbs cross a predefined threshold a seizure is detected, so this algorithm focuses on seizures with a large motor component, like the hypermotor seizures. Therefore, data that is used comes from 7 patients with hypermotor seizures, as already introduced in Section 5.1.

5.2.1 Detection algorithm

In Chapter 3 the segmentation of the raw ACM data is discussed. From the epochs that were identified there, features are extracted to discriminate between epileptic and normal movements. As hypermotor seizures contain a large movement component we mainly investigate the energy in the epochs in this approach.

First the energy of the signal of each ACM is calculated in a sliding window of half a second. Thanks to this sliding window, sharp peaks from a single or a few samples are filtered out. These short peaks are most likely not originating from a hypermotor seizure as these last longer, but they can be caused by epileptic or non-epileptic jerks. As we only want to detect hypermotor seizures in this algorithm, those events should not be detected. The algorithm calculates the average energy in the sliding window and keeps the maximum of this value for each epoch. Thus for each epoch, four values are retained, namely the maximum average energy in each ACM.

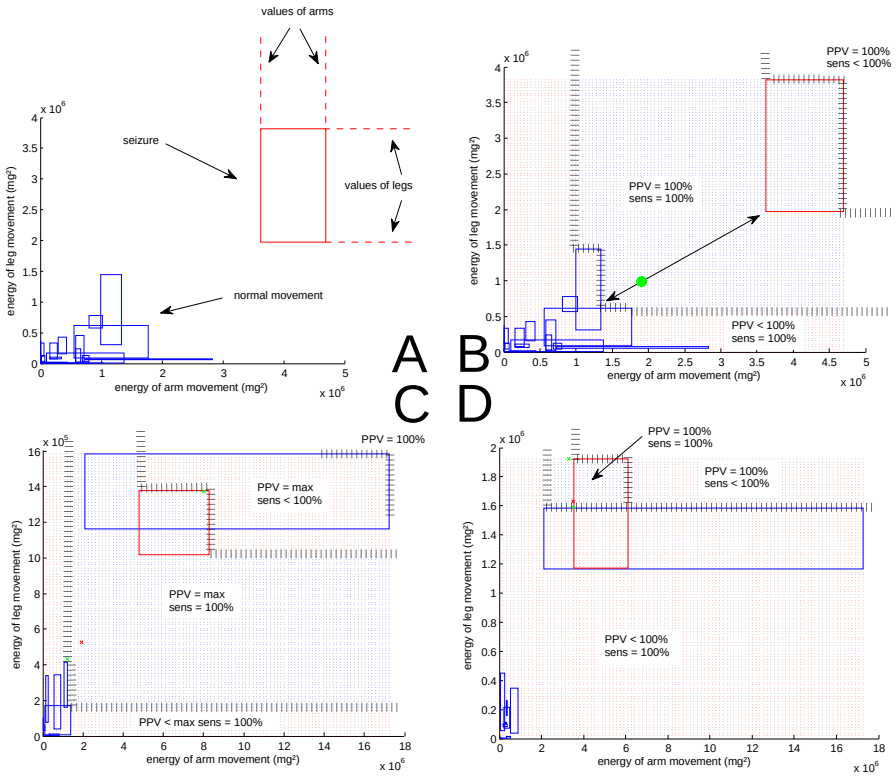


Figure 5.2: *Determining ideal point for thresholding in three different cases (B-C-D). The edges of the rectangles indicate the maximum squared acceleration of each limb. The red and blue rectangles indicate a seizure and normal movement, respectively.*

With these straightforward features we want to investigate if a detection is possible. To visualize the different movement epochs, both normal and epileptic, we plot rectangles where the edges represent the maximum squared acceleration of a limb in one epoch. This is shown in Figure 5.2a. Normal and epileptic movement is shown in blue and red, respectively. As visible in this figure, epileptic and normal movement is well-separable with a large margin between the two classes. Determining an optimal point for separating the two classes for different situations is discussed in the next section.

5.2.2 Training and validation

We want to investigate if it is possible to configure our algorithm patient-specifically with as few prior knowledge from the seizures as possible. In the training phase we train our algorithm with only one seizure and a number of normal movements, depending on the dataset. This number varies between 24 and 220 (depending on the number of normal movements in the dataset) to maintain the same normal movement/seizure balance in the tested subsets. A more elaborated clarification of this approach is given later on in this section.

A high sensitivity is more important in our study than a high PPV, because we do not want to miss a seizure as the patient can injure himself but a few false positives are tolerable. To determine the optimal threshold for the arms and for the legs based on this one seizure we first estimate the region in which the sensitivity is 100% with the largest possible value for the PPV in the training phase. In Figure 5.2 this is indicated as the blue marked area (also indicated as PPV = max; sens = 100% or PPV = 100%; sens = 100%).

When the two classes, epileptic and non-epileptic, are well-separable (Figure 5.2b), the optimal point is found in the following way. The point will be located on the line between the lower left corner of the seizure and the point in the blue area closest to the origin. Because a high sensitivity is more important than the number of false positives, we locate the point closer to the normal movement epochs than to the seizure. In this way, seizures with a lower amplitude are still detected.

When a PPV of 100% is not reachable, e.g. in Figure 5.2c, the same method can be used, but applied to the area with a sensitivity of 100% and the highest possible PPV.

When the lower left corner of the seizure is situated outside of the blue area, the ideal point lies in this blue area, collinear with the lower left corner of the seizure and the point in the blue area closest to this lower left corner, shown in Figure 5.2d.

Notice that in Figure 5.2c there is a region where a sensitivity and PPV of 100% is reached, but it is situated above the topmost normal movement. If the ideal point would be in this area it would be less sensitive for other seizures. So the possible region where the optimal point is located is bounded by the maximal value for arm and leg movements found in the training set for normal and epileptic movement.

To validate our method with a training set containing only one seizure, we cut the data from each patient in equal parts, all containing one seizure. For example, the data from patient 2 contains 9 seizures. This dataset is split up in 9 parts, all

Table 5.2: *Results of k-fold cross validation, with k being the number of seizures of the according patient. The results are averaged out over the k randomizations. The mean sensitivity and mean PPV for each patient is given in column 1 and 3. The second and last column give the standard deviation over the calculated sensitivity and PPV values.*

patient	sens.(%)	std.(%)	PPV(%)	std.(%)
1	100.00	0.00	66.67	47.14
2	98.61	4.17	66.13	19.97
3	50.00	70.71	1.72	2.44
4	50.00	0.00	26.98	19.97
5	63.45	20.88	25.62	6.35
6	100.00	0.00	66.86	22.61
7	66.67	28.87	23.33	2.89

containing one seizure and a fixed number of normal movements. The validation is performed using a k-fold cross validation where k is the number of seizures. This means that one by one every fold is used as training set and the other folds are used a test set. Notice that normally k-1 folds are used for training and one fold as test set, but here we train the algorithm only on one fold as we only want to make use of one seizure in the training phase.

5.2.3 Results

The result of this cross validation is shown in Table 5.2. This table shows the average sensitivity and PPV over all the tests (e.g. 9 tests for patient 2) with the according standard deviation.

For patient 3 the sensitivity is 50% with a lot of false positives resulting in a PPV of 1.72%. For patients 4, 5 and 7 we get a sensitivity between 50% and 67% with a PPV of about 25%. For patients 1, 2 and 6 we get good results: a sensitivity between 98% and 100% with a PPV of about 66%.

5.2.4 Discussion

Although this approach is an ad hoc baseline method, still, for three patients, this method works well. The sensitivity is almost 100% in patients 1, 2 and 6, with

only one false detection in three on average. But notice that in patient 1 we could only validate on two seizures. For patient 2 and 6 we had more reliable results.

The results on the other patients were less convincing because there, the seizures and the normal movements overlap in the feature space, and a good separation between the two classes was not possible. So far we only looked into the energy of the signal, but more features can be extracted, e.g. frequency information. This can increase the performance on the data where the algorithm performs poorly. The extraction of different features in the time en frequency domain is elaborated in the next sections.

The training on one seizure is sensitive to changes in the seizure's feature values. Group specific information can be incorporated as prior knowledge combined with the patient-specific information. This can make the algorithm more robust and increase the sensitivity and PPV.

The algorithm is now trained patient-specific, but it would be more convenient to train it on a group specific base. The found thresholds over the patients with a good detection (1, 2 and 6) were located more or less in the same region, but of course depending on the used training set. With the margin between the two classes being large enough, a shared optimal point can be found with still a high sensitivity and PPV.

5.2.5 Conclusion

This section proposed a method to detect nocturnal hypermotor seizures in pediatric patients. To determine the patient specific thresholds we train the algorithm on only one seizure and at least 20 normal movement events. In three out of seven patients this gives good results with an average sensitivity of more than 98% and PPV of 66%. But a more profound extraction and selection of features and an more robust classifier should improve the performance as will be investigated in Section 5.3 and Section 5.4.

5.3 Support Vector Machines

The use of Support Vector Machines [183] has become very popular since the late 1990's [17]. SVMs construct an optimal n-dimensional hyperplane for the classification of two classes by maximizing the margin. This margin is defined by the largest distance to the nearest training example of both classes. By transforming

Table 5.3: *Overview of data from the patients with a least 4 hypermotor seizures. This is a subset of the dataset given in Table 2.3.*

Patient number	Nights of monitoring	Hypermotor seizures	Normal movements
2	2	9	287
5	5	26	784
6	2	7	381
Total	*9*	*42*	*1452*

the data to another space, by making use of kernel methods, the linear classification can be used to obtain a non-linear classification in the original input space.

As mentioned in Section 5.1, SVMs are sensitive to unbalanced data. Therefore, to overcome the issue of imbalance, we evaluated the following approaches:

- use a weighting factor to train the SVM model
- resample the training data by estimating the probability density function of the seizure data
- resample the seizure data based on the SMOTE technique [32]

Note that undersampling of the majority class was not considered due to the very limited number of minority class samples that were available. Furthermore we evaluated the linear and the radial basis function (RBF) kernel and the influence of the balance and the number of training points on the performance of our classifier.

Similar to the other approaches, we use a dataset of patients suffering from hypermotor seizures. However, in this study we only selected patients with at least 4 hypermotor seizures in our dataset, because for training and testing of the classification model, we need at least 3 and 1 seizures respectively. An overview of the dataset is shown in Table 5.3.

5.3.1 Feature extraction and selection

From the epochs we got from the preprocessing, we derived several features in the time and frequency domain. Table 5.4 shows an overview of these features, which

Table 5.4: *Overview of features in the time and frequency domain. These features are investigated on their ability to discriminate between normal and seizure events.*

Group	Features
Time domain	• mean and standard deviation over 12 mean amplitudes in time of each channel • mean and standard deviation over the 12 standard deviations (std) in time of amplitude of each channel • maximal amplitude of resultant in time per limb • correlation between the resultants of limbs • kurtosis of individual channels • root mean square (rms) of amplitude of individual channels • skewness of individual channels
Frequency domain	• peak frequency of spectrum of individual channels • energy in signal of individual channels • spectral edge frequency (80%, 90%, 95%) of individual channels • power in sub-bands (1-3Hz, 4-8Hz, 9-13Hz, 14-20Hz) of individual channels • normalized power in sub-bands of individual channels

are a subset of the features that are found in the literature. These are commonly used features as described in the work of Dargie [58] and Temko *et al.* [169].

The features in [169] are derived from EEG-signals for epileptic seizure detection. [58] derives features from ACM signals, in multiple cases, including human movement.

Features that are derived from individual channels are not likely to be descriptive over time or patient, as if the patient rotates a limb, the same discriminating feature can be found in another channel. Indeed, the channel which is discriminative can change depending on the orientation of the patient. Therefore we chose to only take into account features that contained more general information including information of all the channels of a limb or from all the ACMs combined. Therefore, to get more general features, for the kurtosis, rms, skewness and frequency features, only the mean, the minimal and maximal values over the calculated values over the 12 channels were considered. The other features are used as illustrated in Table 5.4.

Kruskal-Wallis test

In the first feature selection approach, we use a filtering method. We rank the different features based on the result of the Kruskal-Wallis test [23]. Using this method, we find the features with the highest univariate discriminative power, patient specifically, taking into account all the epileptic and non-epileptic epochs. The features are ranked in order of decreasing discriminative power and then combined to determine the best features over all patients. To combine the features, each feature got a score for each patient. The scores over all the patients were added, and the features with the best scores were used in the remainder of our research. Notice that this feature selection was executed on all 7 patients and not only on the 3 that are used in this SVM approach. Furthermore, because the feature scores were combined after they were calculated patient specifically, all patients get the same weighting, regardless of whether one patient has more seizures than the other.

The five features using this approach are listed below (in order of decreasing score):

- the max of resultant over both arms (Equation 5.1)
- the mean std over all channels (Equation 5.2)
- the length of the epoch, n
- the mean rms over all the channels (Equation 5.3)
- the lowest power in frequency band 1-3 Hz over all the channels (Equation 5.4)

In a more formal way, these features can be expressed using the following equations:

$$\max_t \left(\|x_{1-3}(t)\|, \|x_{4-6}(t)\| \right), \tag{5.1}$$

$$\frac{1}{12} \sum_{i=1}^{12} \sqrt{\frac{1}{n-1} \sum_{t=1}^{n} (x_i(t) - \bar{x_i})^2}, \tag{5.2}$$

$$\frac{1}{12} \sum_{i=1}^{12} \sqrt{\frac{1}{n} \sum_{t=1}^{n} x_i(t)^2}, \tag{5.3}$$

$$\min_i (f_i^{1-3Hz}), \tag{5.4}$$

where f_i^{1-3Hz} is the mean power spectral density (PSD) of one channel $x_i(t)$ between 1 and 3 Hz. $x_{a-b}(t)$ represents the channels with index a until b as a function of time. n is the number of samples (thus the length) of a certain movement epoch.

All the selected features had a p-value lower than 0.05 for each patient except for patient 3, where the mean standard deviation over all channels and the mean rms over all the channels had a p-value of 0.186 and 0.066, respectively. A lower p-value, means that this feature has a higher discriminative power.

Integrating SVM classifier

A second method we used for finding the best features is to directly use the results from a classifier trained with the according feature or a set of features as a measure for their discriminative power, a so-called wrapper method [102]. More features are added iteratively in a forward selection, so the selection procedure starts by using a univariate selection, feeding only one feature to the classifier. After testing all the different features using the cost function discussed in Section 5.3.2 applied to the results of the classifier, the best one is selected and used for further processing. Similar, more features are added until there is no further increase of the performance. After that, a backward elimination is applied to remove the features with the least improvement of the performance.

For training and testing an SVM classifier, sufficient examples of both classes are needed as already mentioned in Section 5.1. Because of this, we calculated the optimal combination of features only on patient 2, 5 and 6. However, to make sure that all of the three patients have the same influence on the feature selection, the

same number of seizures and normal movements were taken from every patient. The seizures and normal movement that were used were randomly chosen for every randomization[2].

Applying this method, the following set of features gives the best performance:

- the mean over all 12 standard deviations calculated over each individual channel (Equation 5.5)
- the highest normalized power over channels in subband 1-3 Hz (Equation 5.6)
- the lowest peak frequency over all channels (Equation 5.7)

Similarly to the previous feature selection method we can also write down these features in a more formal way, as given by the following equations:

$$\frac{1}{12}\sum_{i=1}^{12}\sqrt{\frac{1}{n-1}\sum_{t=1}^{n}(x_i(t)-\bar{x_i})^2}, \tag{5.5}$$

$$\max_i\left[mean_\omega\left[\frac{f_i^{1-3Hz}(\omega)-\bar{f}_i^{1-3Hz}}{std_\omega(f_i^{1-3Hz}(\omega))}\right]\right], \tag{5.6}$$

$$\min_i(\arg\max_\omega(f_i(\omega))), \tag{5.7}$$

using the same notation as the previous part.

In the first iteration of the feature selection where we only looked at the individual features we observed that this method produces similar results compared to the Kruskal-Wallis method. The five features from the latter method were ranked at the $1^{st}, 2^{nd}, 8^{th}, 17^{th}$ and 18^{th} place, respectively, from a total of 61 features that were investigated. In the second iteration we observe that the second best feature from the first iteration drops to the 45^{th} place (of the 60 remaining features), in the third and fourth iteration even to the 55^{th} and the 56^{th} place, respectively (of 59 and 58 features) which probably means that it is highly correlated with the best performing feature of the first iteration, as this approach also avoids to keep correlated features in the pool of best performing features.

[2]Similar to the SVM classification elaborated in Section 5.3.2, the results for using the SVM classifier for the feature selection were also averaged over 10 randomizations to get less biased results.

5.3.2 Test outline

Different classification models are estimated using state-of-the-art SVM inference. For this purpose the freely available libsvm [29] implementation was adopted[3]. Two groups of related application specific tests were set up in order to investigate: a) which alternative to adjust model inference for the high class imbalance performed best, b) the influence of the degree of class imbalance on the classification results c) the influence of the training set size on the results.

All the tests we perform, are conducted in a 10 fold randomization with different combinations of seizures and normal movements. The folds are stratified, which means that every fold contains the same number of seizure and non-seizure events. The tests are evaluated using sensitivity, the PPV and specificity. The sensitivity, PPV and specificity values are averaged over the 10 randomizations and the standard deviation is calculated. Thanks to this randomization the obtained results are less biased.

Methods for coping with class imbalance

To overcome the issue of imbalance, we explored: a) different weight factors: the first applied to the error term in the SVM learning objective corresponding to the positive examples and the second to that corresponding to the examples of the negative class. For this purpose a weight factor T is introduced which balances the weights of the error terms for both classes in the SVM objective. b) Resampling of the minority class, using 1) sampling from a probability density function estimated using the available minority examples 2) generating new minority class examples using the SMOTE technique [32].

Resampling based on the *estimation of the probability density function* estimates the distribution of the seizure data points along every feature using a non-parametric kernel density estimation, with a Gaussian function as kernel and a bandwidth equal to the standard deviation of the seizure data. Based on this estimated distribution, new points are randomly sampled. The *SMOTE technique* generates new data points by introducing synthetic examples along the line segments that connect minority class examples with their k minority class nearest neighbors.

In a first series of tests, we used the same number of seizures s and normal movements n for each patient in the training ($s = 4, n = 190$) and testing phase ($s = 2, n = 95$), to investigate the performance of the different approaches, and to

[3] http://www.csie.ntu.edu.tw/~cjlin/libsvm/

be able to compare them over all the patients. These movements were randomly chosen from each patient but with a strict separation between training and test set. We used this specific number of normal movement and seizures based on the smallest numbers available for each patient, i.e. 7 seizures for patient 6 and 287 normal movements for patient 2. For each patient, we used two thirds of the data for training and one third for testing.

In the training phase, we used a three-fold cross-validation to find the optimal hyper-parameters for the SVM classifier. These parameters include the regularization parameter C, kernel parameter σ and weight factor T. In a formal way, the support vector classification can be written according to the optimization problem of Equation 5.8.

$$\begin{aligned} \min_{w,b,\xi} \quad & \frac{1}{2} w^T w + C^+ \sum_{y_i=1} \xi_i + C^- \sum_{y_i=-1} \xi_i \\ \text{subject to} \quad & y_i(w^T \varphi(x_i) + b) \geq 1 - \xi_i, \\ & \xi_i \geq 0, \quad i = 1, 2, ..., n, \end{aligned} \tag{5.8}$$

with n being the number of training examples. x_i are the training vectors in the input space, y is the class indication vector with $y_i \in \{1, -1\}$, w denotes the normal vector in the decision hyperplane. $\frac{b}{\|w\|}$ determines the offset of the hyperplane from the origin along w. ξ_i are the slack variables that are introduced to measure the degree of misclassification (vectors at the *'wrong'* side of the hyperplane, the variable indicate the distance to this plane). C^+ and C^- are the regularization parameters for the positive and negative class, respectively. They give the degree of penalization for misclassification for both classes. $\varphi(x_i)$ is a mapping from the input space to a higher dimensional space to allow for a non-linear separation in the input space, with $K(x_i, x_j) = \varphi(x_i)^T \varphi(x_j)$ being the kernel function. We use the radial basis kernel function that is given by:

$$K(x_i, x_j) = e^{-\frac{\|x_i - x_j\|^2}{2\sigma^2}}. \tag{5.9}$$

We can interpret the hyperparameters C, σ and T in the following way: C affects the trade-off between complexity and proportion of nonseparable samples [36]. If it is too large, we may overfit and consequently store many support vectors, if it is too small, too much smoothing may be applied giving an underfit [4]. The kernel parameter σ denotes the kernel width parameter. A larger σ leads to a smoother

fit. The weight factor T gives a lower or higher weight to the error cost of the negative class, in such a way that $C^- = C \cdot T$ and $C^+ = C \cdot (1 - T)$, as used in Equation 5.8.

The different values we tested for the hyper-parameters are:

- $C : e^{-3,-2,\ldots,6}$
- $\sigma : \frac{1}{\sqrt{2}} e^{-2.5,-2,\ldots,2}$
- $T \in \{0,\ 0.01,\ 0.02,\ 0.05,\ 0.1,\ 0.2,\ 0.5,\ 0.7,\ 1\}$

The cost function gives a higher weight to the sensitivity compared to the PPV more specifically:

$$cost = -(2 \cdot sens + PPV), \tag{5.10}$$

Influence of class imbalance

In the second series of tests, we use the optimal approach found in the first tests, and investigate the influence of the number of training examples and the balance on the performance of the classification model. To investigate the influence of the balance, we keep the number of seizures in the training set fixed (4), and increase the number of normal movement from 4 (equally balanced) to 190. The same ratio is used in the test set. The different numbers we used can be found in Table 5.7 containing the results of this investigation.

Influence of training set size

For the investigation on the influence of the number of training samples, we keep the balance the same (1 seizure for 33 normal movements), and we test this for 2, 3 and 4 seizures in the training set, and 2 seizures and 66 normal movements in the test set.

5.3.3 Results

Table 5.5 shows the results of the SVM classifier using the different methods compensating for the imbalance. To be able to compare the different methods, we calculated a cost value for each test, using the same cost function for model selection

Table 5.5: *Results of performance of SVM classifier using different methods to overcome the problem of imbalance. The performance results in this table are all expressed in %. The cost in the last column corresponds to Equation 5.10.*

Balance factor								
	patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
linear	2	**90.00**	21.08	**95.00**	15.81	99.89	0.33	275.00
kernel	5	**0.00**	0.00	**0.00**	0.00	100.00	0.00	0.00
	6	**95.00**	15.81	**100.00**	0.00	100.00	0.00	290.00
	all	**15.00**	24.15	**50.00**	50.00	99.68	0.51	80.00
								188.33
	patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
RBF	2	**90.00**	21.08	**95.00**	15.81	99.89	0.33	275.00
kernel	5	**5.00**	15.81	**3.57**	9.45	98.63	1.32	13.57
	6	**95.00**	15.81	**100.00**	0.00	100.00	0.00	290.00
	all	**10.00**	21.08	**25.00**	41.83	99.26	0.87	45.00
								192.86
Probability density estimation oversampling								
	patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
linear	2	**90.00**	21.08	**53.65**	34.70	96.53	3.55	233.65
kernel	5	**95.00**	15.81	**12.10**	3.30	84.42	5.36	202.10
	6	**100.00**	0.00	**59.00**	30.15	97.68	1.91	259.00
	all	**80.00**	34.96	**10.19**	14.66	68.84	21.58	170.19
								231.58
	patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
RBF	2	**90.00**	21.08	**53.33**	32.04	96.11	4.52	233.33
kernel	5	**70.00**	25.82	**8.10**	2.05	82.84	7.00	148.10
	6	**100.00**	0.00	**56.38**	32.01	97.26	2.23	256.38
	all	**75.00**	35.36	**12.67**	11.65	82.00	16.11	162.67
								212.60
SMOTE oversampling								
	patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
linear	2	**90.00**	21.08	**77.33**	25.13	99.16	1.09	257.33
kernel	5	**75.00**	35.36	**12.24**	5.59	88.11	4.63	162.24
	6	**100.00**	0.00	**83.33**	22.22	99.37	0.89	283.33
	all	**70.00**	34.96	**16.40**	16.60	87.26	7.45	156.40
								234.30
	n°	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
RBF	2	**70.00**	25.82	**86.67**	21.94	99.68	0.51	226.67
kernel	5	**40.00**	39.44	**10.88**	11.20	93.89	3.28	90.88
	6	**75.00**	26.35	**96.67**	10.54	99.89	0.33	246.67
	all	**45.00**	43.78	**18.65**	23.46	92.53	4.92	108.65
								188.07

according to Equation 5.10. This cost value is shown in the last column of Table 5.5. Notice that the cost values have a negative value, as a lower cost implies a better performance. The higher the absolute value, however, the better the classifier performs. The cost value in bold below the performance values is the averaged cost over patients 2, 5 and 6 to get a general overview of the cost of the different techniques. In general, we can observe that the oversampling techniques work better than only using a different weighting for both classes, although this is mainly due to the low performance on patient 5, as for patients 2 and 6 the performance of using a different weighting seems to give better results. Furthermore, the SMOTE technique gives better results than the density estimation oversampling, especially the PPV values of using SMOTE are higher compared to the density estimation oversampling. For the SMOTE technique, the linear kernel performs better than the RBF kernel, as the absolute cost value of the former is higher.

To investigate the influence of the size of the training set and the balance in the data, we used the oversampling based on the SMOTE technique with the linear kernel, as this technique seemed to give the best results of the compared methods. In the first test, we evaluated three different sizes of training set with the same balance, and a test set with a fixed amount of data. Table 5.6 shows the results of this test. We observe a little increase in performance when using more data, but also here, this is mainly thanks to the performance increase of patient 5. For the other two patients no clear trend is visible.

Table 5.7 shows the results of the second test, where we tested the influence of the balance of the dataset on the performance. These results show that the balance of the dataset has a clearer influence on the performance compared to the dataset size. When the number of seizures and normal movements is equal, the performance is best. This is especially visible in patients 2 and 5. But when the imbalance increases, we see a gradual reduction of the performance, which is also reflected in the absolute cost value for each test, which decreases from 279.44 when the dataset is balanced, to 238.09 when the imbalance is largest (4 seizures compared to 190 normal movements). For patient 6 this trend is not clear, but in patients 2 and 5 the absolute cost values decrease with each test.

5.3.4 Discussion

The oversampling methods give the best results. However, we still observe that with an increasing imbalance, the performance gets worse. It should be noted that thanks to the oversampling, the SVM model is still trained with an equally balanced training set, and that only the test set remains unbalanced. Making

Table 5.6: *Results of performance using resampling based on SMOTE and a linear kernel in the SVM classifier. This table shows the influence of different sizes of the training set. The performance results in this table are all expressed in %.*

Training: 2 seizures, 66 normal Test: 2 seizures, 66 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**90.00**	21.08	**90.00**	16.10	99.55	0.73	270.00
5	**40.00**	39.44	**12.74**	16.61	91.06	8.70	92.74
6	**100.00**	0.00	**85.00**	19.95	99.24	1.07	285.00
all	**65.00**	24.15	**26.24**	12.45	93.18	4.06	156.24
							215.91
Training: 3 seizures, 99 normal Test: 2 seizures, 66 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**90.00**	21.08	**83.33**	17.57	99.24	0.80	263.33
5	**65.00**	33.75	**12.72**	6.13	86.97	6.32	142.72
6	**100.00**	0.00	**88.33**	19.33	99.39	1.06	288.33
all	**55.00**	36.89	**16.11**	13.88	90.76	5.73	126.11
							231.46
Training: 4 seizures, 132 normal Test: 2 seizures, 66 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**90.00**	21.08	**86.67**	21.94	99.39	1.06	266.67
5	**85.00**	24.15	**17.81**	6.19	86.67	5.84	187.81
6	**100.00**	0.00	**83.33**	17.57	99.24	0.80	283.33
all	**70.00**	34.96	**20.26**	16.69	89.09	6.26	160.26
							245.94

use of only the weighting factor gives bad results for patient 5. The classification problem is easier for patients 2 and 6. Therefore we believe that the oversampling methods are more robust even for more difficult datasets or problems.

In most of the trials, the performance on patient 5 is lower. This is due to the different clinical manifestation of the seizures of this patient, which are clearly shorter in time and lower in intensity.

Note that for the calculation of the performance on all patients, the varying number of seizures per patient is not taken into account. This means that the influence of patient 5 is higher than those of patient 2 en 6 as 62% of the total seizures come from patient 5 (26 on a total of 42). This explains why the performance of all patients combined is most of the time very similar to the performance on patient 5. However, we did not take the cost value into account when calculating the average cost over the three patients (in bold) for each test.

In the real setting the balance between epileptic and normal movements can differ from patient to patient, and therefore influence the patient-specific performance.

Table 5.7: *Results of performance using resampling based on a density estimation and a linear kernel in the SVM classifier. This table shows the influence of a different balance between seizures and normal movements in the training and test set. The performance results in this table are all expressed in %.*

Training: 4 seizures, 4 normal Test: 2 seizures, 2 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**100.00**	0.00	**88.33**	19.33	80.00	34.96	288.33
5	**90.00**	21.08	**73.33**	23.83	60.00	39.44	253.33
6	**100.00**	0.00	**96.67**	10.54	95.00	15.81	296.67
all	**75.00**	35.36	**79.63**	26.06	70.00	42.16	229.63
							279.44
Training: 4 seizures, 10 normal Test: 2 seizures, 5 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**90.00**	21.08	**100.00**	0.00	100.00	0.00	280.00
5	**80.00**	34.96	**72.22**	22.05	82.00	17.51	232.22
6	**100.00**	0.00	**96.67**	10.54	98.00	6.32	296.67
all	**80.00**	25.82	**66.67**	24.85	80.00	16.33	226.67
							269.63
Training: 4 seizures, 30 normal Test: 2 seizures, 15 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**90.00**	21.08	**95.00**	15.81	99.33	2.11	275.00
5	**80.00**	34.96	**43.36**	25.48	85.33	9.32	203.36
6	**100.00**	0.00	**100.00**	0.00	100.00	0.00	300.00
all	**70.00**	34.96	**55.31**	35.30	84.00	19.17	195.31
							259.45
Training: 4 seizures, 70 normal Test: 2 seizures, 35 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**90.00**	21.08	**91.67**	18.00	99.14	1.93	271.67
5	**80.00**	34.96	**24.17**	11.63	86.00	6.24	184.17
6	**100.00**	0.00	**85.00**	19.95	98.57	2.02	285.00
all	**70.00**	34.96	**30.54**	18.56	88.57	10.43	170.54
							246.95
Training: 4 seizures, 132 normal Test: 2 seizures, 66 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**90.00**	21.08	**86.67**	21.94	99.39	1.06	266.67
5	**85.00**	24.15	**17.81**	6.19	86.67	5.84	187.81
6	**100.00**	0.00	**83.33**	17.57	99.24	0.80	283.33
all	**70.00**	34.96	**20.26**	16.69	89.09	6.26	160.26
							245.94
Training: 4 seizures, 190 normal Test: 2 seizures, 95 normal							
patient	**sens.**	std.	**PPV**	std.	spec.	std.	-cost
2	**90.00**	21.08	**83.33**	22.22	99.47	0.74	263.33
5	**75.00**	35.36	**10.95**	5.02	87.79	5.21	160.95
6	**100.00**	0.00	**90.00**	16.10	99.68	0.51	290.00
all	**70.00**	34.96	**15.92**	15.53	87.58	6.85	155.92
							238.09

The number of seizures we collected is small for some patients, which makes the training of an SVM classifier more difficult if we want to make patient-specific models. However, the results of the tests on the influence of the dataset size indicated that there is no big difference when changing the amount of data. This statement needs to be interpreted carefully since the number of seizures was too small to draw any conclusions about the relation between seizure set size and classification performance. Increasing the number of seizures is likely to have a positive impact on the classification performance.

Due to the limited number of seizure examples, the weighting method does not give good results for patient 5. This can be explained by the fact that the number of the support vector candidates is small, which reduces the flexibility of setting the decision hyperplane. We also noticed that the cost of most of the different parameter combinations (C, σ, T) was the same when using only the weighting of the classes. This means that due to the limited number of training examples, the hyper-parameters only have a small influence on the decision hyperplane.

When the classifier fails to classify any event as a seizure, the PPV value becomes *'not a number'* as both numerator and denominator become zero. For calculating the average PPV and the standard deviation in the 10 randomizations, we left these values out.

In our tests we used 2 seizures in the test set using 10 randomizations. Therefore, the resolution of the sensitivity is only 5%. This is also visible in the tables showing the results.

For using the SMOTE technique, we have only a limited number of seizures in our setup when using the 3-fold cross-validation in determining the optimal parameters for the SVM. The new data points are generated on the line segments connecting the minority class examples, but sometimes there is only one nearest neighbor (for determining the line segments) for generating new data points. Nevertheless, the SMOTE technique seems to give the best overall results.

We also evaluated a cost function taking into account the values of the SVM classification indicating the distance from the data points to the decision hyperplane. This function was similar to Equation 5.10, but replacing the binary values for a TP, FP and FN for calculating the sensitivity and PPV with these distance values. However, this did not give any better results compared to our original cost function, in most cases the performance was even lower.

Table A.1 in the Appendix shows the results that were obtained when using the five best features according to the Kruskal-Wallis test. Similar to Table 5.5 the

results for the different approaches are shown. From this table we can draw similar conclusions than we previously did on Table 5.5, except that here the density estimation oversampling gives better results than the SMOTE technique. Another striking conclusion that we can draw is that the overall performance on patient 5 increases in every setting[4]. This can be explained by the fact that for the feature selection using the Kruskal-Wallis test, the relative influence of patient 5 was lower, as a scoring of all patients was taken into account. Using the SVM classifier, only patient 2, 5 and 6 were included for determining the best features. And as the seizures manifest themselves differently in patient 5, the features that are most discriminative are different. This makes us suspect that selecting the right features is crucial in obtaining a good performance.

5.3.5 Conclusions

This section gave an overview of applying SVMs to our application of detecting nocturnal epileptic seizures using ACMs. We have tested different approaches to overcome the imbalance problem. Oversampling of the minority class seems to give the best results, especially the SMOTE technique. On all three patients, this technique gives a sensitivity of 75% or more and for two patients a PPV more than 75%. The results are comparable to the former method of thresholding explained in Section 5.2, and even a bit better. A full comparison between the different methods is given in Section 5.5.

5.4 Outlier detection by estimating the Probability Density Function

In this section, we investigate if it is possible to detect seizures in a robust way, by only making use of normal movement data. This would facilitate employing the detection system in a home situation as no labeling of the data is needed beforehand. This approach is based on novelty detection. We model normal movement, and try to detect abnormal (novel, not yet or rarely seen) events, i.e. epileptic events.

Novelty detection is typically used in settings where the abnormalities occur relatively seldom, and when insufficient data is available to model the abnormal activity. This can also be seen as one-class classification, in contrast to the

[4]Except for the first two tests where we evaluated the performance of changing the balance factor where the performance is very low anyway

traditional two class (or multi-class) classifiers. We believe this approach has a high chance of success since a large number of normal activities and very few epileptic events are observed.

For modeling the normal movement, we estimate the probability density function (PDF) using non-parametric Parzen windows [15]. A non-parametric estimator was chosen since such approach does not make any assumption on the form of the data distribution and therefore is more flexible. Parzen density estimation is commonly used in the context of classification and novelty detection [127]. Rathi *et al.* [150] and Rangayyan and Wu [147] modeled the distributions of two classes based on the Parzen window approach and classify instances based on the highest probability. Yeung and Chow [197] and Tarassenko *et al.* [167,168] use the one-class classification which is the novelty or outlier detection. A PDF for normal behaviour is estimated and new instances with a probability, computed using this PDF, below a certain threshold level are considered abnormal.

The proposed approach must be contrasted to binary classification methods which use information from both classes (seizures and normal movements) to estimate a model. Requiring information of both classes for estimation implies a costly annotation process. Seizure segments must be isolated from the normal movements by humans. Our one-class method does not require knowledge about the seizure class. Furthermore, since the seizure class is populated with much less examples than the normal class it can typically be assumed that the very small portion of seizure examples has a negligible effect on the estimation of the PDF of the normal movement. In a practical system, when some patient-group specific tuning parameters are fixed, this means that a patient-specific model can be obtained with much less human interaction compared to two-class models.

5.4.1 Model estimation and hyperparameter tuning

As already described in Section 5.1 and illustrated in Figure 5.1 we use the same features for the PDF approach as already used in the SVM approach in Section 5.3 to model the normal activity (example is given in Figure 5.3). Movement events are classified in the test phase. When they are an outlier of the normal activity distribution, they are classified as seizure.

To model the normal activity, we estimated its PDF with the Parzen window method [15]. In this work a radially symmetrical Gaussian kernel function is chosen which is smooth and can be completely specified by a bandwidth (variance) parameter only.

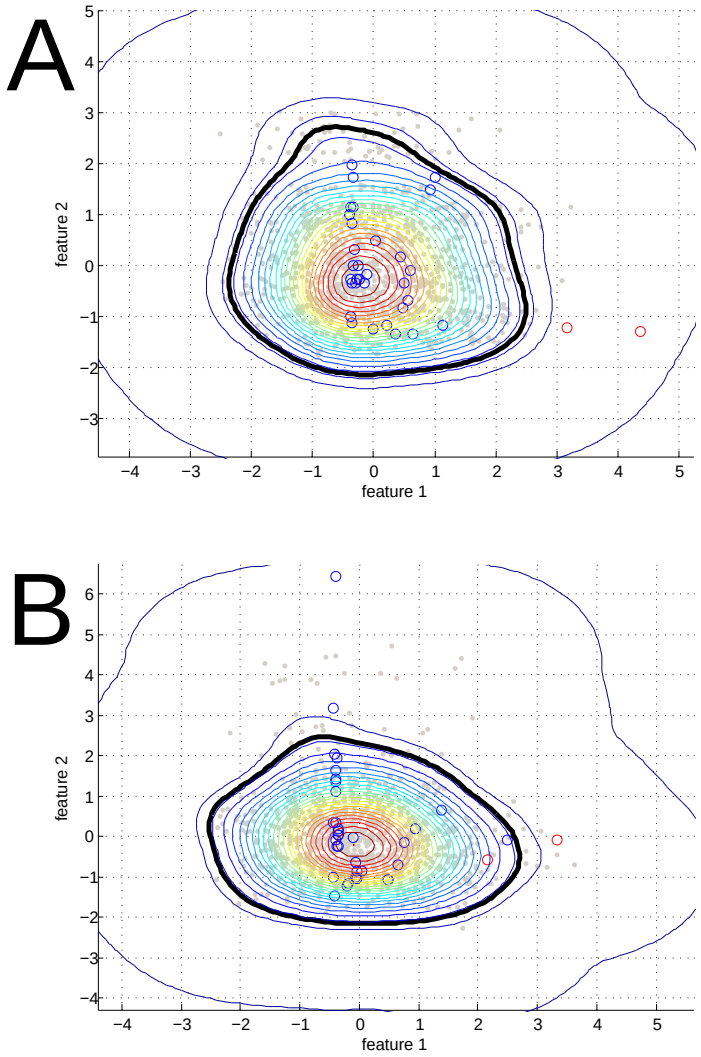


Figure 5.3: *Probability density functions of patient 2 and patient 5 for one randomization based on the first two features. The gray dots are the training examples, after the resampling. These examples were used to construct the PDF. The black border is the 95% decision line, based on the training data. The red dots are the seizures in the test phase. The normal activity from the test phase is indicated in blue. The normal activity and epileptic seizures are well separable in patient 2 (a) in patient 5, two false positives and one false negative occur (b).*

This bandwidth is a hyperparameter β that should be determined. It has an influence on the smoothness of the estimated distribution. A larger bandwidth results in a smoother density estimation.

A second hyperparameter that needs to be determined is the threshold on the PDF τ at which we separate normal from epileptic movement. We determine this threshold based on a fixed probability of 95%. This means that the 5% movements with the lowest probability to be normal are classified as epileptic. This 95% border is a conservative estimation of the number of seizures, by which we want to say that it is lower than the typical number of normal movements in the datasets. By setting this threshold too low, we ensure ourselves that this method won't miss any seizures (or at least in most cases). A disadvantage is that in this way, there are too many detections, hence false alarms, which reduces the PPV. But as said before, in this application it is more important to detect (almost) all seizures than having a low false alarm rate.

To determine β, we used the approach that is shown in Figure 5.4. Different values of β are evaluated on the performance using a cost function, more specific 2^n where $n = -2, -1, ..., 4$.

For determining the threshold τ, we first sample a fixed number of points from our training set, and calculate the probability of each point based on the PDF. From these probabilities, a histogram and its cumulative distribution function (CDF) is calculated. Using this CDF we can determine the threshold (probability value) that is associated with the 95% border.

To ensure we have enough and a fixed number of data points in the training set, we oversample the data by generating new data points according to the PDF of the original data.

We fixed this number of data points to 1000 to have a sufficient number to estimate the PDF. In case the training data does not contain such a high number of points, the resampling ensures to generate enough points.

To make the approach as general as possible we want the hyperparameters to be group specific (i.e. for all patients with hypermotor seizures) and not patient specific. In this way, they do not have to be adjusted per patient. Therefore, to determine β we made use of a leave one out approach on a patient basis. Since we have 7 patients, we use 6 of them to estimate the optimal bandwidth (training phase), and the remaining patient in a patient-specific test phase to estimate the PDF and classify movement events.

In this experiment we used two third of the normal movements that was available

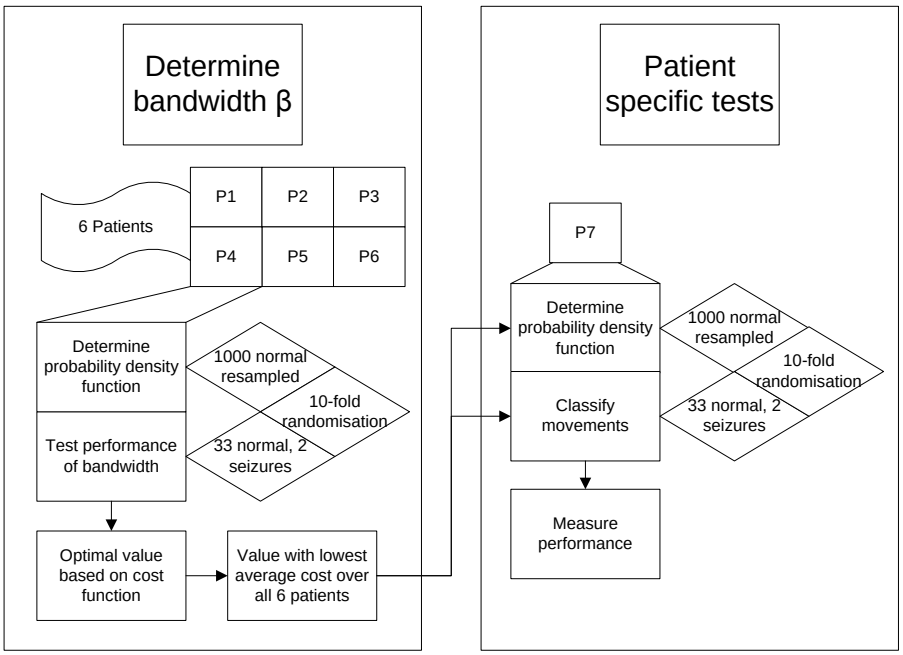


Figure 5.4: *Schematic overview of tuning β and the patients specific testing. Notice that the 1000 normal movement used for the determination of the PDF are resampled from two thirds of the normal movements of one night.*

for a patient as training, and 33 complementary normal movements and 2 seizures for testing the performance of the chosen bandwidth β.

These numbers for the test set were chosen because we wanted to have the same number of examples in each patient, and for some patients we only had 2 seizures and not much more than 100 normal movements (so 33 for one third of the night).

This training and testing is done in a 10-fold randomization, so for each value of $p1$ and for each patient, 10 PDFs are built and tested. We calculated the performance of the model as the average of 10 costs calculated using Equation 5.10, which was also used in the previous section.

Here, the cost of a false negative is higher than the cost of a false positive, because missing a seizure is worse than generating a false positive for this type of seizure.

Table 5.8: *Results of validation using the PDF approach with the Kruskal-Wallis features. For each patient the sensitivity (sens), the PPV and the specificity (spec) are shown, together with their respective standard deviation over 10 randomizations*

patient	sens. (%)	std. (%)	PPV (%)	std. (%)	spec. (%)	std. (%)
1	100.00	0.00	58.00	19.58	94.55	3.73
2	100.00	0.00	76.33	26.50	96.97	3.78
3	100.00	0.00	86.67	17.21	98.79	1.56
4	100.00	0.00	80.00	21.94	97.88	2.49
5	50.00	35.36	50.00	35.36	96.36	3.44
6	100.00	0.00	70.00	21.94	96.67	2.65
7	100.00	0.00	69.67	23.44	96.36	3.44
Total	*92.86*	*5.05*	*70.10*	*23.71*	*96.80*	*3.01*

Validation

As already mentioned, a leave one out approach is used to validate the models. Meaning that the model hyperparameter β is estimated using 6 patients and then kept fixed on the remaining patient.

For this 7th patient a performance score was computed using the same setup. For 10 randomizations a PDF of normal nocturnal movement based on two thirds of the normal movements of all nights of this patient is estimated (with fixed β). Next the sensitivity and PPV were computed on a set of 33 normal and 2 epileptic movements. Finally, the mean of the sensitivity, PPV and specificity over all the randomizations were used as validation measure.

5.4.2 Results

The results shown in Table 5.8 are obtained by using the PDF approach on the five features selected with the Kruskal-Wallis method. The results obtained using the SVM features, are to be found in the appendix in Table A.2. For all the patients, the sensitivity, the PPV and the specificity are given with the according standard deviation over all the randomizations.

The sensitivity is 100.00% in all patients except patient 5. The PPV is always 50% or more.

The optimal value for β is determined on a group of 6 patients excluding the patient in whom the performance measures are calculated. The cost over these 6 patients is averaged out, and the β related to the lowest cost is used for estimating the PDF. For all patients except patient 5, a bandwidth β of 8 was obtained in this way, while the tested values being 2^n where $n = -2, -1, ..., 4$. But when testing with another threshold τ, the optimal bandwidth β can change.

Table A.3 in the appendix illustrates the influence of β on the obtained results. As visible, this influence is limited. A smaller β gives a lower PPV (58.77% compared to 69.11% when β is 1) and an increase of the sensitivity of almost 3%. When β is 1 or higher, the difference becomes smaller, and the difference between a β of 4 and 16 is even negligible.

Figure 5.3 shows the probability density functions from patient 3 (A) and patient 5 (B), together with the epileptic activity from one randomization in red. The black decision line separates normal and epileptic movement (as is the shown PDF). For visualization purposes, only two dimensions are shown, feature 1 and feature 2. From this figure it is clearly visible that the normal and epileptic activity is well separable for patient 3 but not for patient 5 as the normal and epileptic activity overlap for this patient.

Investigating the results based on the features extracted from the SVM, which can be found in Table A.2 in the appendix, we observe that these are inferior compared to the features from the Kruskal-Wallis method. We can easily explain this by the fact that these features were chosen on information of only three patients and in a combined way. We can observe that for patient 2 and 6 the performance is comparable to the results in Table 5.8. Also in this experiment the results for patient 5 are lower. From this observation we can conclude that some features differ from patient to patient and it may be better to combine the feature that perform best on each patient individually rather than look for features than work best over all patients. Determining the optimal features based on the PDF classification, so that it is possible to extract them for every patient, would therefore be a valuable extension to the research. If we, for example, change the '*lowest peak frequency over all channels*' feature with the *length of the epoch* we already get a substantial improvement of the results as visible in Table A.4 in the appendix. Note that this last test is also performed with 3 features, so the number of features itself is not the issue.

5.4.3 Discussion

In our results section, we evaluated the method on the different patients, using the same number of seizures (2) and normal movements (33) in the test set. This makes a comparison between patients fairer. It also allows us to calculate and average sensitivity and PPV over all patients. However, for some patients, the balance between normal and epileptic movements is different when we would consider all the patient's data, which has an influence on the performance. We observed that the PPV value decreases when a relative larger number of normal movements occur with respect to the number of seizures. The influence of the balance on the PPV value is also explained in a next paragraph.

A more optimal result of the PDF classification can be achieved when choosing τ in a better way, for example when making use of prior knowledge of the relative number of seizures that occur. We see in Figure 5.5, which shows the performance for different values of τ, that for patient 2 and 6 a sensitivity and PPV value of more than 90% can be achieved when choosing the best value for τ. However, in our approach we want to introduce as little as prior knowledge as possible, as this makes a setup in a real setting more feasible. We want to be able to directly place the system in a home situation without acquiring too much data beforehand.

When we look at the manifestation of the seizures in patient 5, we see that they are short in duration and subtle in intensity and are clearly different from the seizures in the other patients. This makes it hard for the algorithm to detect them correctly, which explains the lower performance in this patient.

Novelty detection works well in situations where the occurrence of abnormal events is very low. In this case the distribution of normal events can be constructed from all the data available. If the occurrence of abnormalities is low enough, these may not influence the distribution too much.

Another advantage of this approach is that in the learning phase, the events should not be labeled, as all the events are used for estimating the PDF. In our case this would mean that, given a group-specific determined β and τ, we can immediately install the detector in the room of the child. Movement data can continuously be collected for updating the PDF. This does not require any other human action.

Initially a subset of normal movement from other patients can be used, which is gradually replaced by patient specific movements as more data is collected from that patient.

By doing the resampling of the training data and making the threshold fixed to a

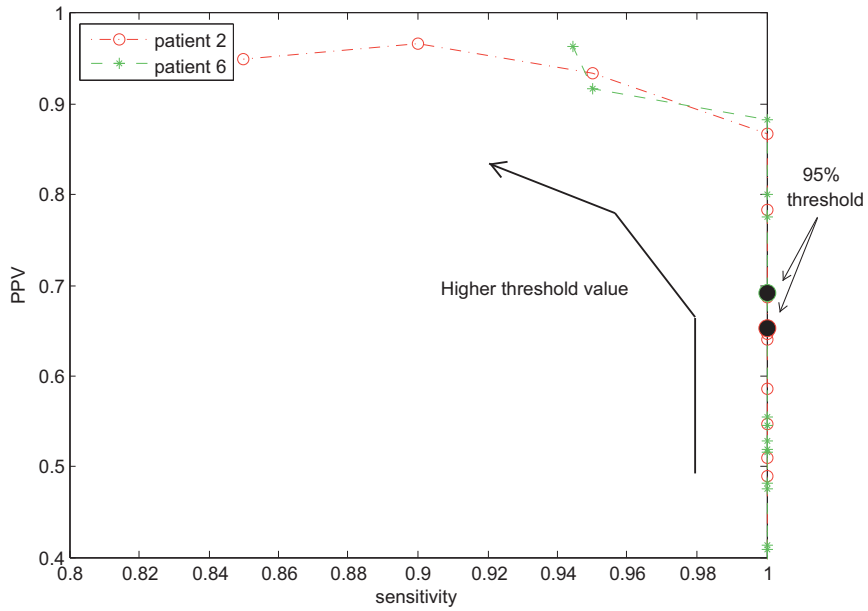


Figure 5.5: *Performance of PDF approach when using different values for the threshold τ (between 25% and 99.9%). The full black dots indicate the 95% threshold. We clearly see that an improvement is possible compared to this 95% threshold.*

probability percentage, we make the algorithm less dependent on the number of available examples and the number of features that are used. Indeed, the optimal τ value is dependent on the number of examples used, as the estimated PDF improves with the number of data points.

A possible downside of this method is that also other abnormal activity is detected. However, this could be interesting, because if an abnormal event occurs, even if it is not epileptic, the parents or caregivers can be alarmed.

Another extension of the method could be integrating an adaptation of the probability density function over time. We can make the algorithm robust against a change of normal movements of the patients over time (a change in behavior). This could be done by discarding the older movement epochs that were used for building the PDF and replacing them by new movements that are observed.

The obtained value for the PPV is dependent on the threshold, and the balance between both classes. This can easily be explained using the following reasoning: consider the case that both classes are separable, so there is no overlap between

classes, and that τ is set in a way that all seizures are detected, with a certain number of false positives. Let B be the balance, the positive examples (seizures) divided by the total number of movements (seizures + normal movements). Let n be the total number of data points, and τ the threshold, which in our case would be 95%. We can state that:

$$PPV = \frac{TP}{TP + FP} = \frac{B.n}{(1-\tau).n} = \frac{B}{1-\tau}. \tag{5.11}$$

As we assume that all seizures are detected, we can say that $TP = B.n$. To get a high PPV, $1 - \tau$ should be a good estimation of B. Note that $B <= (1 - \tau)$, otherwise seizures are missed, we get false negatives, and the ideal case we assumed would not hold anymore. So the threshold could be optimized in function of the balance.

We have now tested the algorithm on one type of seizure. However, if we want to apply it to other types of seizures we have to make sure that the seizures do not occur too often throughout the night. Indeed, if the patients suffer from too many seizures, the probability density function will be too much influenced, and seizures will not be regarded as an abnormal event anymore. So in this case many seizures will be missed, unless the threshold is set lower, which leads to more false positives.

5.4.4 Conclusion

We build a non-parametric patient specific probability density function based on normal movement to distinguish epileptic activity from normal activity. This resulted in an average sensitivity of 92.86% and a PPV of 70.10% for 7 patients.

5.5 Comparison between methods and conclusion

In this chapter three different methods for detecting hypermotor seizures using ACMs are tested, starting with a simple thresholding method, as a baseline. Afterwards a state-of-the-art SVM classifier is tested in different settings to overcome the imbalance in the data. In a last section, an outlier detection based on the probability density function estimation is applied to the data. Each of these methods cover a different application area. When little seizure information is available, the outlier detection method is most suited, as no seizure information is

needed for constructing the classification model. When more seizure information is available, the SVM approach is better suited as on the three tested patients this method gave somewhat better results using the SMOTE technique, especially with respect to the PPV. The thresholding method gives clearly worse results compared to the other techniques. Especially if we compare the results of patients 3, 4 and 7 to the PDF approach. For patients 2 and 6 the results are in the vicinity of the SVM and PDF results.

Thanks to the limited need of seizure data, the PDF approach is easiest to apply in a home situation. However, the movement patterns that differ most from normal behavior will be detected as seizures, even in healthy subjects. So, using this technique, epilepsy can not be diagnosed. It can only detect seizures in epileptic patients.

As already mentioned in the introduction, the issue of imbalance had different implications on the three methods. For the SVM classification, unbalanced data is an issue that should be addressed, as we did in Section 5.3. In contrast, the PDF approach only works in situations where there is an imbalance in the data.

In general we can say that the sensitivity of the SVM approach is lower and the PPV value is higher compared to the PDF approach. The difference between these values is due to the different manner of optimizing the classification model. In the PDF approach we use a rather conservative threshold τ in a way that in all patients but one, all the seizures are detected. In the SVM approach, we use a cost function to optimize the model where the sensitivity has a larger weight than the PPV.

The advantages of the PDF classification compared to the SVM classification are the possibility to train the model without positive training examples (no labeling needed) and the flexibility to adjust the model with an increasing number of data. In this way, when we install the system in a real home situation, a better model of normal movement can be learned over time, when more and more data becomes available.

The results we have are initial results on a group of seven patients. To make more reliable conclusions we will have to test the system and algorithm on a larger group. Furthermore, the system should also be tested in a real home situation. For now it is only tested in a hospital setting. However, when we test it at home, we will lose the EEG information, and a labeling based on the golden standard is not possible for validation.

The results in Table 5.8 and Table A.2 reveal the importance of choosing the right

features. As the manifestation of the seizures can be patient-specific, the best discriminating features can indeed change from patient to patient. It is therefore advisable to determine the features on a patient basis that is large enough, and preferably perform it patient-specifically and adding the best features, instead of determining the features group-specifically.

To conclude we can say that both the PDF and the SVM method perform well, with sensitivities over 90% and PPVs between 50% and 90%. The SVM approach is more suited when enough seizure data is available, where the PDF method also works when little information about the seizures of the patient is known.

Chapter 6

Unimodal Seizure detection - Video

In Chapter 4 the preprocessing of the video data was discussed, whereas the previous chapter elaborated on the detection of seizures with ACMs. This chapter continues the logical structure of this thesis, by introducing the detection of seizures using video. One approach that is followed is the usage of the optical flow, already explained in Chapter 4, to collect movement information that can discriminate between normal and seizure-like movements. By clustering segments with similar motion, different moving (body)parts can be described in time. This approach is introduced in Section 6.2. A second approach makes use of spatio-temporal interest points (STIPs). By extracting spatio-temporal features around these interest points on different time and space scales, no tracking of the interest points is needed. Section 6.3 goes more into detail in this approach. Finally, Section 6.4 compares both methods and concludes this chapter.

6.1 Introduction

In this chapter, we investigate the detection of seizures using video data. More specifically, we focus on myoclonic seizures. As already introduced in Chapter 1, these are small jerk-like movements that can occur as a single seizure or precede other seizures like tonic seizures. An overview of the data can be found in Table 2.4. In total, data from 3 patients is used.

As already mentioned in Chapter 1, human motion analysis has already been studied extensively in other applications such as gesture recognition or surveillance [87] [176]. The nocturnal detection of seizures is in some points more challenging. We cannot use the skin color information, which is often used in gesture analysis, due to the gray scale images during nocturnal acquisition. Also, detecting the patient's extremities is most of the time not feasible because the patients are covered with a blanket. And finally, there are different types of motor manifestations in seizures. Therefore we focus on one type of seizure only, namely the myoclonic seizure.

Karayiannis [97] already studied the detection of seizures in neonates, with a focus on clonic and myoclonic seizures. The patients were monitored in the neonatal intensive care unit where all the body parts were clearly visible for the camera. The motion of the body parts was quantified by temporal motion-strength signals based on optical flow computation, and by motion trajectories. The classification based on neural networks resulted in a sensitivity and specificity above 90%.

The difference with our study is that we focus on a different patient group. Furthermore, the setting is as natural as possible, i.e. the patients can be covered with a blanket and we do not make use of markers, just as in a normal home situation. And finally, we also make use of new features that we extract from the motion tracks and from spatio-temporal interest points.

One way to describe motion patterns in video is to extract image features and track them in time [22]. Gauglitz and Höllerer [75] give an overview of different features that can be extracted from images. The extraction of corners in images as defined by Harris and Stephens [82] is commonly used. Shi and Tomasi [160] improved this method, that is also used in the KLT tracker [173]. The scale-invariant feature transform (SIFT) [121], another popular method for extracting features in images, can detect features that are scale- and rotationally invariant [22]. Viola *et al.* [187] detected specific short-term motion patterns of pedestrians in videos instead of tracking features over a long period of time. The detection of spatio-temporal interest points is a method introduced by Laptev [111] where interest points in the space-time domain are found based on a 3D Harris detector. As the features are extracted on different time scales, the higher scales provide information over a longer time course, so no tracking of the features has to be done to get this information. In this chapter, we apply this method proposed by Laptev on our data in Section 6.3.

Another way to extract motion patterns from video is by using the optical flow information, as already explained in Section 4.2.2, and cluster the flow vectors in homogeneous blobs with similar motion directions. Previous research showed

us that clustering algorithms can be used to cluster optical flow results for image segmentation [74]. In this paper, optical flow and spectral clustering are applied on video recordings. Clustering of flow vectors is also used by Min *et al.* [133] for the classification of movements in ballet sequences. Hu *et al.* [88] used a hierarchical clustering but their goal was to cluster different global motion patterns of moving crowds, and not to cluster the vectors in each time instance.

Clustering can be considered as an unsupervised classification problem where different patterns (feature vectors) are classified in groups (clusters) [91]. Unsupervised means that the given data is not labeled i.e. no prior knowledge is available to which class an instance belongs. Clustering can also be described as grouping instances with similar characteristics, and keeping instances with dissimilar characteristics in different groups. Defining if instances are similar can be obtained by introducing certain (dis)similarity measures.

Determining if a clustering satisfies the user's needs is not trivial. The user has to supply some criterion to evaluate the result of the clustering. Jain states in his overview paper [92] that the determination of the number of clusters is one of the most difficult problems in data clustering. Most of the time the algorithm is run with different values for the number of clusters and based on another criterion function the optimal number of clusters is defined. Some researchers introduced different examples of such criteria. In the *x-means clustering* method [138] the *Akaike Information Criterion* (AIC) and *Bayesian Information Criterion* (BIC) were used. Figueiredo and Jain [66] employed the *minimum message length* (MML) in combination with a *Gaussian Mixture Model* (GMM) for estimating the optimal number of clusters. Tibshirani *et al.* [170] made use of gap statistics, here it is assumed that if the number of clusters is optimal, the influence on the clustering of random noise on the instances is minimal. Rasmussen [149] used the *Dirichlet Process* [65] to make an estimation of the number of clusters. Although the efforts of introducing these objective criteria, determining which number of cluster leads to the most meaningful result is still not solved [92].

Probably one of the most popular methods for flat clustering[1] is the *k-means clustering* algorithm [17]. Another method, the *spectral clustering* has become one of the most popular modern clustering algorithms [126]. It uses a similarity graph indicating the similarity between all the data points in the dataset that is used for clustering the data. Both methods still need to have the number of clusters as an input. We evaluate these techniques in this chapter in Section 6.2. A method that does not need to know the number of clusters beforehand is the mean-shift

[1] In contrast to hierarchical clustering, flat clustering does not output a hierarchical structure of the instances.

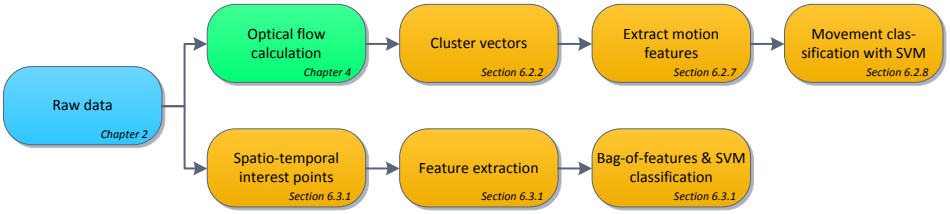


Figure 6.1: *Schematic overview of the different approaches for seizure detection based on video addressed in this chapter.*

clustering algorithm [35, 73]. We compare this method with the former methods in an evaluative study.

Figure 6.1 gives an overview of the different aspects that are discussed in this chapter. Two main approaches are discussed, which can be seen in the two different branches in the structure of Figure 6.1. The first one combines the optical flow information, already discussed in Section 4.2.2 with clustering of the vectors. The clustered blobs are tracked in time and different features are extracted, after which the motion tracks are classified using an SVM approach similar as already used in Section 5.3.

The second approach uses the STIP features for classifying motion epochs. In the neighborhood of the detected interest points, flow features are derived at different scales in space and time. From a training set, 50 prototype motions of the interest points are defined using an unsupervised k-means clustering. Using a bag-of-features approach, the interest points in the test set are assigned to a prototype cluster and using the distribution of the interest points over the prototype clusters, an SVM is trained for classification.

Finally Section 6.4 concludes with a comparison between both approaches and a general conclusion.

6.2 Clustering flow vectors

As introduced in Section 4.2.2, the extraction of motion vectors from the video sequences is carried out by the Horn-Schunck optical flow calculation [86]. In this section we want to segment the flow vectors into homogeneous blobs containing vectors with similar location and direction, to consider them as moving (body)parts in time. To this end, we first threshold the optical flow vector field, to remove low amplitude noisy motion vectors and to get a first segmentation, as shown in

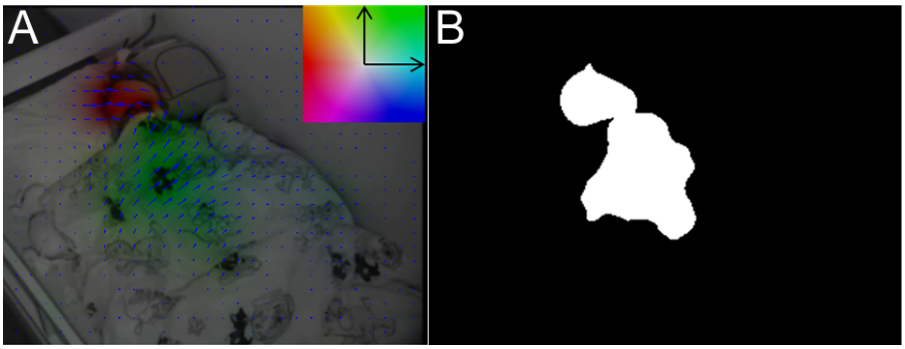


Figure 6.2: *Motion vector field in one frame from epileptic patient, the colors represent the direction of motion vectors (a) the frame after thresholding (b).*

Figure 6.2. This threshold is determined based on the maximal flow vector value of a video sequence without movement, and is done for every clustering method that we investigated. In the left part of Figure 6.2 the motion vectors in one video frame are drawn on top of the gray scale image. To enhance the visibility of the amplitude and direction of the motion, we added color information, where the hue[2] represents the direction of motion (as visible in the top-right corner of panel A), and the saturation denotes the motion amplitude. On the right side the thresholded image is shown.

After tracking the clusters in time, we extract seven features, related to amplitude, duration and direction of the motion. Then, a validation with a linear support vector machine classifier is carried out. But first we will discuss the steps that are involved in the clustering of the flow vectors.

6.2.1 Features used for clustering

Before we apply the clustering, we identify features from the motion field upon which we want to distinguish the different clusters. These features include the magnitude, direction and location of the flow vector. Features representing the appearance of the moving parts such as the gray value of the moving pixels, were not taken into account as these are not homogeneous in one segment.

The extraction of these features is trivial, except the direction of motion. For this direction, the radian angular value is used, but motion that occurs around π

[2]According to the HSI representation of color images, where H, S and I denote the *hue*, *saturation* and *intensity*, respectively.



Figure 6.3: *Result after clustering the motion vectors. This is an example of a satisfying clustering, as both moving body parts are grouped in two different clusters. The method that was used in this example is mean-shift clustering. The red circles indicate the modes of the clusters. The colors indicate the direction of the motion.*

radians can experience a phase shift of 2π which does not allow a proper clustering. Unwrapping did not always lead to satisfying results, therefore this feature was split in two, by giving the real and imaginary component of the motion vector projected to the unit circle. Note that in this way we did not combine the amplitude and angle in two features as we could by just using the real and imaginary part of the flow vectors, we wanted the angle and amplitude to be separate to be able to adjust the scaling of both features as will be explained in Section 6.2.3.

6.2.2 Clustering methods

Using the different clustering methods we want to group instances that share certain similarities or characteristics. In our case, we want to group flow vectors with similar magnitude, direction and location as discussed in the previous section. This results in different (body)parts moving independently from each other. Figure 6.3 shows a clustering that segments the parts in the video in a good way. This image shows the same video frame as was shown in Figure 6.2. The head which moves to

the left and the body (covered by a blanket) that moves to the right, are segmented in different clusters.

To this end, we tested three spectral clustering variants [62,113], k-means clustering, and mean-shift clustering on our data. Chen *et al.* [34] provides an implementation of the spectral clustering[3] that we applied in our study.

K-means clustering is an iterative clustering method where the instances closest to the cluster centers[4] are assigned to that specific cluster. The cluster centers of the next iteration are calculated by taking the centroid of the instances assigned to the corresponding cluster. *Spectral clustering*, in turn, first constructs a similarity matrix, which expresses the similarities between each instance and the other instances of the data. These similarities can be expressed using a certain similarity measure e.g. the Euclidean distance. When we would represent the data as a graph, the connecting lines between instances correspond to their similarity. What spectral clustering does, is to *cut* the connecting lines in such a graph in a way that different clusters are fully disconnected, and that the sum of the similarity values of the cut connections is as low as possible. *Mean-shift clustering* is a clustering method that searches for maxima in density functions. For each instance in the dataset, a gradient ascent is performed based on the estimated density. All the instances that converge to the same local maximum (the so called *mode*) are clustered in the same group.

In both k-means clustering and spectral clustering the number of clusters should be defined beforehand. The mean-shift clustering doesn't need this value, but the bandwidth of the density estimation window determines the size of the clusters. If a small window is used, more (smaller) local maxima are found which results in more clusters and vice versa.

Spectral clustering

As introduced in the previous paragraph, spectral clustering makes use of a similarity matrix to find the best clustering. The similarity matrix is of size $n \times n$ with n the number of instances, representing all the similarity values s_{ij} between all the different data points x_i and x_j (with i and $j \in \{1, 2, ..., n\}$). As mentioned before, we now want to partition the graph in such a way that the sum of the similarity values of the cut connections is as low as possible.

[3] http://alumni.cs.ucsb.edu/~wychen/sc.html

[4] In the first iteration these cluster centers can be chosen randomly or at certain fixed points, using prior knowledge.

In graph theory a graph can be represented by the Laplacian matrix $\mathbf{L}$, which is constructed by $\mathbf{L} = \mathbf{D} - \mathbf{W}$ where $\mathbf{W}$ is the similarity matrix and $\mathbf{D}$ is a diagonal matrix with in each element on the diagonal d_i the sum of the similarity measures s_{ij} for all values of j. Taking the eigendecomposition of the Laplacian, the optimal solution for the segmentation in two classes is found in the eigenvector with the second lowest eigenvalue[5]. When splitting the data in more than two classes, the values of the eigenvectors can be used (after normalization) to be clustered again by k-means clustering.

In small datasets, the spectral clustering can produce high quality clusterings. However on large datasets, the computational time rapidly increases, as the complexity of the algorithm is $O(n^3)$, with n the number of data points [195]. And as our dataset is rather large (each pixel representing one instance), we use three different approaches that reduce the computational complexity. The three implementations of the spectral clustering we investigated are the spectral clustering using a sparse similarity matrix and the spectral clustering using Nyström method with and without orthogonalization.

Spectral clustering with sparse similarity matrix In this approach the number of non-zero elements in the similarity matrix is reduced by only keeping the similarity values s_{ij} of the data points x_j that are the k nearest neighbors of x_i.

Spectral clustering with Nyström approximation of similarity matrix The Nyström approximation [69] randomly samples $l \ll n$ points from the dataset, and calculates the similarities between these random points, and between these points and the rest of the data [34]. From this subset of similarities the eigenvalues and eigenvector are approximated. However, using the Nyström approximation, the columns of the eigenvector matrix are not orthogonal, in contrast to the original spectral clustering. We used two implementations of the Nyström method, one with and one without an extra orthogonalization.

K-means clustering

K-means clustering is an iterative clustering method that first randomly specifies k center points. Next it calculates the Euclidean distance between these centroids and the data points. Data points are grouped according to their closest centroid.

[5] As the lowest value is zero and indicates a partition where no connections are cut, hence lowest value of the cut connections, and all data points are stored in one class.

New centroids are calculated to be used in the next iteration until the centroids converge.

Mean-shift clustering

The mean-shift algorithm is a technique for the localization of maxima in a density function. It can also be used for clustering purposes, as will be explained in this section. The mode[6] localization is an iterative procedure that starts in a pixel in the segmented area. In the next iteration the point moves to a new location according to Equation 6.1.

$$m(P) = \frac{\sum_{i=1}^{n} K(\text{dist}(P, Q_i))F(Q_i)w(Q_i)Q_i}{\sum_{i=1}^{n} K(\text{dist}(P, Q_i))F(Q_i)w(Q_i)} \tag{6.1}$$

where P is the original location, Q are the pixels in image $\mathbf{I}$ (and Q_i is one pixel in the image with a total of n pixels), K is the chosen kernel, F is the amplitude of the motion vector and w is a weight function based on vector similarity where $w = \frac{1}{2}cos(\alpha) + 1$ with α the angle between P and Q_i.

The point moves towards the location with the local maximal vector amplitude (mode), as shown in Figure 6.4(C). All the points in the segmented frame that iterate to the same mode belong to the same cluster.

$$K(x) = \begin{cases} 1 - \frac{x^2}{bandwidth^2}, & \text{if } x^2 \leq bandwidth^2 \\ 0, & \text{otherwise} \end{cases} \tag{6.2}$$

As a kernel K we use a parabolic function with a fixed bandwidth, according to Equation 6.2. The weight function w gives a measure for the similarity of the vector directions, with a value of zero if the angle between the two vectors is 180° and one if the angle is 0°. This way, vectors with a different direction are not clustered together.

Figure 6.3 in the beginning of this section showed the results of mean shift clustering. The two parts in the frame with a different movement direction are well split up.

[6] A mode is a local maximum of the density function.

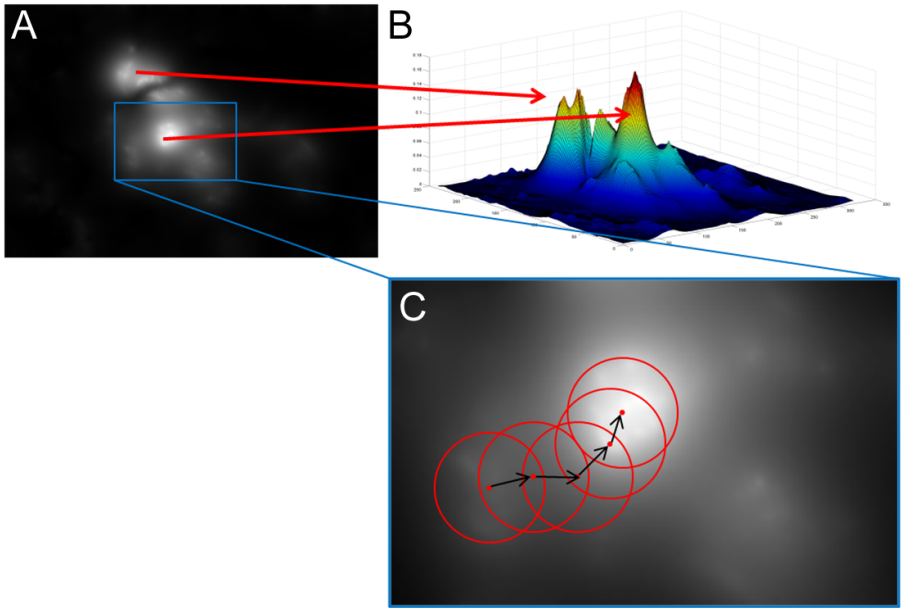


Figure 6.4: *The amplitude image of the motion vector field (a) amplitude image plotted with a clear visualization of the highest amplitudes (b) a detail of (a) with a representation of the iterative mean shift procedure (c).*

6.2.3 Determine number of clusters

Determining the number of clusters is only applied to spectral clustering and k-means clustering. To determine this optimal number, features are extracted that give a quantitative measure of the cluster quality such as the inter and intra cluster distances. The method is evaluated based on data that is manually labeled indicating the ideal number of clusters according to a human labeler.

Scaling of the features

Before the clustering, the features that are used (angle, amplitude and location) are first normalized. Different weighting factors are evaluated based on their apparent performance. That is, we want a cluster to be spatially connected (or at least be in the same region) and generally, the direction of motion is more constant in a desired cluster than the amplitude of the vectors, as the amplitude of the flow vectors decreases towards the borders of the motion.

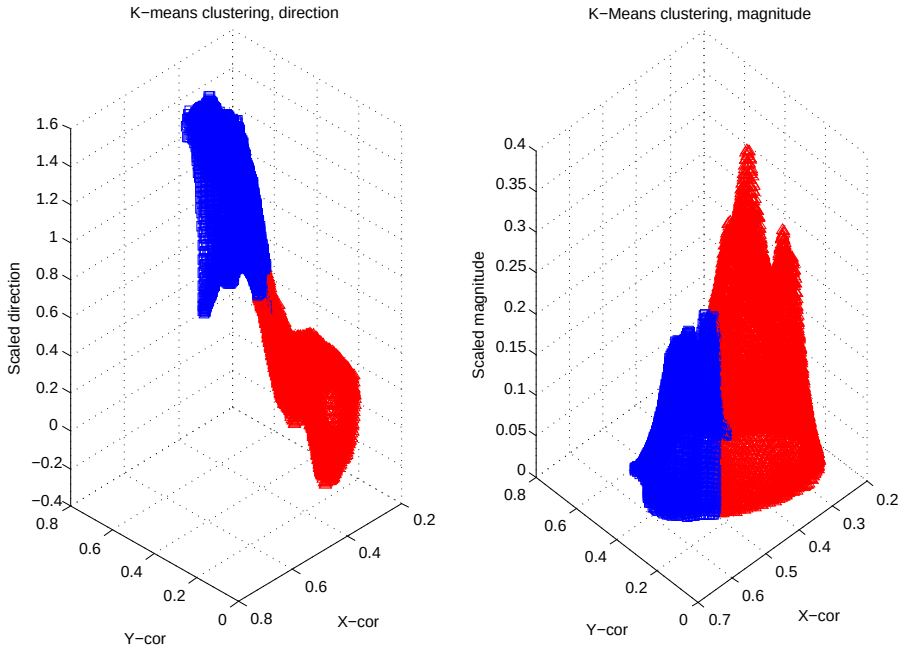


Figure 6.5: *Direction (left) and magnitude (right) plotted versus x- and y-coordinates. The direction feature has a larger weight, the algorithm clusters on this feature as the data points have a bigger variation in this feature and due to the larger weight. The right plot shows the same clustering in the X-Y-magnitude space.*

Figure 6.5 shows the data points of a video frame that is clustered. It is visible in the feature data that two blobs are identified. The axes in this figure are normalized and weighted. The x- and y-axis represent the location of the vector in the image plane, the z-axis in both plots represents the scaled direction and magnitude, respectively. We observe that the weight of the direction is larger than that of the magnitude.

Cluster analysis

To determine the number of clusters, the algorithm (one of the different clustering methods) is run different times with a varying number of clusters which is common practice in the literature [92]. Then, the different clusterings are evaluated based on objective measures that are derived from the clustered data, based on the inter-cluster distance (Equation 6.5) and the intra-cluster distance (Equation 6.4)

[148, 177]. The inter-cluster distance is the distance between the cluster centers. The larger this value, the better the clustering as there is a clear distinction between the clusters. The intra-cluster distance is the average distance between the centroid and other instances of the same cluster. The intra-cluster distance should be as low as possible to have compact clusters.

$$x_i = \begin{bmatrix} d_i^{im} \\ d_i^{re} \\ u_i \\ v_i \\ a_i \end{bmatrix} \tag{6.3}$$

$$intra = \frac{1}{n} \sum_{i=1}^{n} \|x_i - \bar{x}\| \tag{6.4}$$

$$inter = \|\bar{x} - \bar{y}\| \tag{6.5}$$

In the previous equations, x_i is a data point in the five dimensional space, d_i^{im} and d_i^{re} are the two features representing the direction of motion, u_i and v_i indicate the spatial location in the video frame, and a_i denotes the amplitude of the motion. $\bar{x}$ and $\bar{y}$ are the centroids of cluster x and y, respectively.

From both measures, minimal and maximal values and the standard deviation is derived. These features are investigated based on their ability to identify the correct number of clusters, as determined by a human labeler. The number of clusters is varied between one (all data is put in 1 cluster) and 4.

The determination of the number of clusters is done in a cascade. The algorithm is executed dividing the data into 2, 3 and 4 clusters. For every division the maximal inter-cluster distance and the standard deviation of the intra-cluster distance is calculated. Based on these results the optimal division is chosen. The cascade approach can be explained by the fact that we first decide between 1 or 2 clusters (based on the intra- and inter-cluster distances calculated from the division in two clusters), in the next step between 2 or 3 clusters and so on. The process is stopped after the result of a test in the cascade returns the lower cluster value, or when 4 clusters are reached.

The most discriminative features and the according linear decision boundary that are used for determining the number of clusters are found using a validation study of 136 manually labeled frames. 60% of the data is used as training and the remaining frames for testing. The discrimination between 1 and 2 clusters resulted

Table 6.1: *Comparison of the speed of clustering using the different methods. All clusterings are done on the same image with the same number of clusters (if relevant) and averaged out over 20 calculations.*

Time needed for clustering one frame (average of 20 calculations) Intel Core 2 Duo T9600 4Gb Ram, Windows Vista system, implemented in Matlab	
Clustering type	Processing time
Spectral Clustering with sparse similarity matrix	20.165s
Nyström with orthogonalization	0.56145s
Nyström without orthogonalization	0.3015s
K-means	0.04405s
Mean-shift	0.4192s

in a sensitivity of 83.3%[7], a PPV of 90.9%, a specificity of 76.9% and a NPV of 62.5% in the test phase. There was an unbalance in the data (more frames with one cluster than two clusters with a ratio of 2.44) which influenced the results, hence a lower NPV. For the discrimination between 2 and 3 clusters, the sensitivity, specificity, PPV and NPV are 93.30%, 38.90%, 71.80% and 77.80%, respectively. The distinction between 3 and 4 clusters could not be further elaborated due to the few number of examples of frames that occured with this number of labeled clusters.

6.2.4 Comparison of clustering methods

In this section we give a quantitative comparison of the different clustering methods based on their processing time shown in Table 6.1, furthermore we give a qualitative overview of their accuracy of clustering. It has to be noted, however, that when the number of clusters in not known, the clustering can be performed multiple times to determine the number of clusters as elaborated in the previous section.

When we visually inspect the segmentation of the clustering methods we see that the spectral clustering with the sparse similarity matrix gives good results, as visible in Figure 6.6. They are noticeable better compared to the Nyström approximation. In the latter method, the outline of the clusters seem to be noisier and more random (especially when less data is used for the similarity matrix approximation). In our study we used 150 instances for the matrix approximation. In the cluster quality

[7]In these tests, detecting one cluster is considered to be a positive detection.

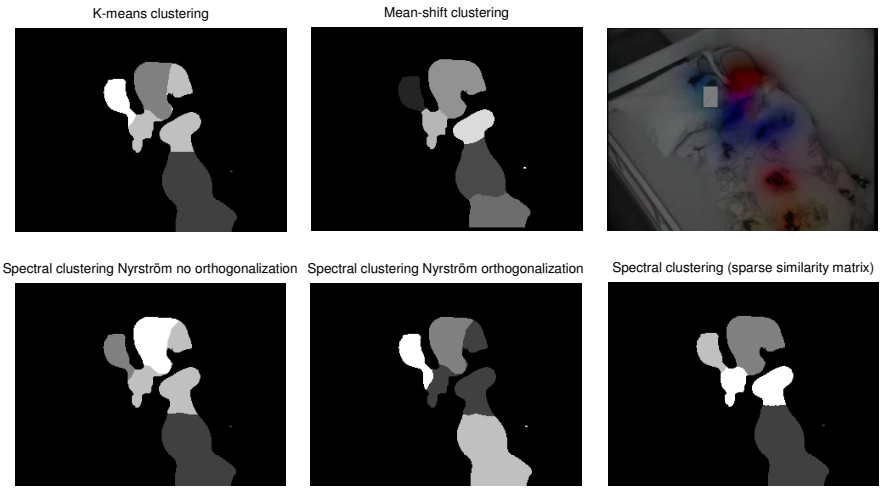


Figure 6.6: *Comparison of the different clustering methods based on the vectors shown in the upper right figure, where the colors indicate the amplitude and direction of motion.*

we do not observe any improvement of using the orthogonalization. The results of the k-means clustering is comparable to the spectral clustering with the Nyström approximation, as also here some randomness in the clustering occurs, and the outline does not seem to fit perfectly with the moving objects. The mean-shift clustering gives better results, which are more stable, comparable to the spectral clustering with sparse similarity matrix.

Although the clustering is not always perfect, the results in all methods seem to be of practical use, as a *perfect match* may not strictly be needed for further processing. In the different methods we also observe large differences in computational complexity. The spectral clustering with sparse similarity matrix takes up to 20 seconds for clustering one frame [8]. Using the Nyström approximation reduces the calculation time roughly by a factor of 50 as visible in Table 6.1. The calculation of the mean-shift clustering algorithm is in the same order of range. Only the k-means clustering is again a factor of 10 faster than the former methods. Therefore we can make the following conclusion. If the clustering needs to be fast and the clusters does not need to have a perfect fit, k-means clustering is advisable. If the speed is not so much of an issue, it is better to use the mean-shift clustering.

[8]The calculation time was averaged over 20 randomizations. The image had 11983 data points to cluster. The same image was used in all calculations over all methods.

Furthermore, the mean-shift clustering does not need to know the actual number of clusters that need to be formed, although a bandwidth has to be provided which has an influence on the size of the clusters. But this is much more intuitive to provide, as we can estimate the order of magnitude of the moving parts we want to be clustered.

6.2.5 Tracking

After the motion vectors in the frames are clustered, the clusters are tracked in time based on the method proposed by Crocker and Grier [46]. They tracked particles over time to investigate the trajectories based on the cluster frame-to-frame displacement. When multiple clusters need to be tracked, the most probable set of N identifications between N locations in two consecutive images is searched for. In our setting the modes of the different clusters are tracked. It is assumed that the location of the object would not change too much from one frame to another (proximity) and the maximum velocity defines an upper bound on the object velocity and limits the possible correspondences to a circular neighborhood around the object [198]. Clusters can be merged or split up for some frames, moving parts can go out of reach of the camera or their motion can be interrupted for a short time. To overcome these problems, it is tolerated that the followed modes are lost for maximum half a second. Otherwise the motion track is split up in two in time.

6.2.6 Feature extraction

From every cluster in every frame the average motion vector is calculated. So, every motion track results in a series of two dimensional motion vectors from which we derive different features. This is visualized in Figure 6.7.

Because the myoclonic jerks are characterized by a short and intense movement of the limbs, mostly in arms or shoulders, in one direction followed by a movement in the opposite direction, e.g. the relaxation after the contraction of the arms muscles, we incorporate features that quantify the change in motion direction.

Therefore we use two features covering the change in direction, namely slow and fast phase changes. We defined a fast phase change as a change in motion direction of more than 2 radians between two subsequent samples. For a slow phase change we look for changes of 2 radians over a longer period of time. Other features we use are the duration of the movement, the maximal, median and average vector amplitude

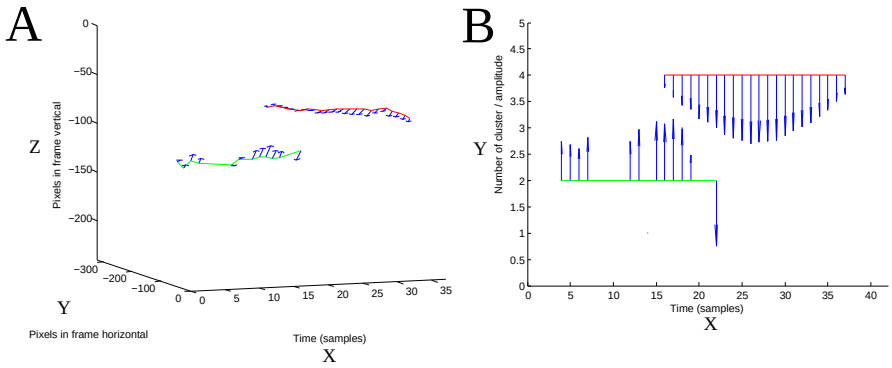


Figure 6.7: *Motion tracks (green and red) with the corresponding motion vectors in blue, the YZ-plane represent the locations in the image frame, the X-direction corresponds to the time (a) motion tracks represented with the amplitude Y-axis) and time (X-axis) (b).*

and the number of vectors in the track. The number of vectors is sometimes smaller than the actual length because the body part can stop moving for up to half a second and still be in the same track.

6.2.7 Classification

The classification of the motion tracks was done similar to the SVM classification approach elaborated in Section 5.3. In a 10-fold randomization the data was trained and tested. 113 motion tracks were extracted from 15 myoclonic jerks from one patient. Three sets of normal movement were used. The first one of the patient getting up in bed, moving around but also making some smaller movement with the head and hand, this set generated 375 tracks. In a second set, the patient lies still in the bed making eye blinks after which she starts moving her head to the left and right and making small movements under the blanket, generating 64 tracks. The third set is similar to the second one, but the patient also scratches her nose, resulting in 29 tracks. In the first set, the normal movements are larger than the myoclonic jerks so we believe that it is relatively easy for the algorithm to distinguish between the majority of these movement tracks. The second and third set contain much smaller movements, which we believe are harder to classify.

As in Section 5.3, we find the optimal values for C and T for the SVM model using a three-fold cross-validation in the training set. Because we use a linear kernel, no value for σ had to be found. As a cost function we stated that the sum

Table 6.2: *Example values of the seven features of a set of 5 normal and 5 epileptic tracks.*

	#vectors	Mean amplitude	Median amplitude	Max amplitude	Quick changes	Slow changes	Length (frames)
Normal movement							
#1	3	0.67	0.71	0.73	0	0	3
#2	48	0.72	0.75	0.95	0	2	49
#3	15	0.43	0.47	0.52	2	0	33
#4	6	0.71	0.71	0.84	0	0	7
#5	6	0.41	0.39	0.48	0	0	6
Epileptic movement							
#1	6	0.26	0.28	0.37	1	0	8
#2	5	0.09	0.10	0.11	1	0	7
#3	33	0.20	0.13	0.59	0	1	44
#4	4	0.38	0.42	0.51	2	0	7
#5	12	0.11	0.12	0.22	0	0	24

of the sensitivity and PPV had to be maximized, $cost = -(sens + PPV)$, as in contrast to the previous setting, missing seizures is not that critical for myoclonic seizures[9]. Notice that here the cost is always negative (or zero) because a lower cost implicates a better performance (hence higher sensitivity and PPV). All seven features, as shown in Table 6.2, are used in for training and testing. To have balanced datasets, the same number of normal and epileptic tracks was taken for each test, as the number of positive examples was not so small in this study. Notice that the classification was performed on the tracks, not on the seizures as a whole.

6.2.8 Results

Some typical values for the track features can be found in Table 6.2. This table shows the values for 5 epileptic and 5 non-epileptic motion tracks.

Table 6.3 shows the results of the classification when using the three different sets for the normal movement. As we expected the first test gives the best results, obtaining

[9]Depending on the decision layer on top of the detection of motion tracks related to seizures, detecting only one epileptic track in a seizure could already be enough. Determining the relation between classifying motion tracks and classifying seizures is a new classification problem.

Table 6.3: *Results of the classification of the motion tracks using different sets of normal movement. The results were obtained over a 10 fold randomization, this way the average sensitivity, PPV and specificity were calculated with their according standard deviation.*

patient	sens.(%)	std.(%)	PPV(%)	std.(%)	spec.(%)	std.(%)
1	87.67	4.17	75.58	3.42	71.50	5.12
2	68.00	14.18	93.66	6.59	95.00	6.24
3	63.00	14.18	91.39	9.64	93.00	8.23

a sensitivity of 87.67% and a PPV of 75.58%. As expected, both classifications on the datasets with smaller movement produced poorer results, especially with respect to the sensitivity.

We also observed the weights of the different features that are used by the SVM classifier for determining the decision hyperplane. A higher weight means that the corresponding feature is more important. For the first dataset we observed that the most important features are clearly the mean, the maximal amplitude and the number of vectors. The other features have a much smaller weight. When we investigate the weights using the third dataset, we observe that the best feature is still the mean amplitude, but closely followed by three features with almost the same weight, namely the length, the maximal amplitude and the number of quick changes. This confirms our suspicion that for the distinction between smaller movements features such as the change in movement direction become more important. It has to be noted that the results are averaged over 10 randomizations.

The classification is carried out on individual motion tracks. The motion tracks extracted during an epileptic jerk are possibly not all from the epileptic movement itself, but also from e.g. movements of the bed because of the seizure. So the classification can be improved on this point by looking only to the tracks related to the seizure. Moreover, the features from different tracks originating from one seizure can be combined, to further improve the detection.

The detection of seizures based on motion tracks extracted from the optical flow calculation and the mean shift clustering algorithm seems possible. In the first test on 15 myoclonic jerks, a sensitivity between 63.00% and 87.67% and a PPV between 75.58% and 93.66% is obtained.

6.3 Spatio-Temporal Interest Points (STIP)

The second approach we will discuss is based on the state-of-the-art technique proposed by Laptev [112]. In a first step, interest points in the video are found. Afterwards, spatio-temporal features from these interest points are extracted. Using a bag-of-features approach, features are fed into an SVM to generate a classification model. Each part of this procedure is explained in the following sections.

For this research we used the data from 3 patients suffering from myoclonic jerks, as given in Table 2.4. Notice that patient 3 also suffers from hypermotor seizures and was also incorporated in the tests in Chapter 5. We did not include these hypermotor seizures in this test. In total we have 27 seizures and 82 normal movements in this dataset.

6.3.1 Spatio-temporal interest points (STIP) and feature extraction

The Spatio-Temporal Interest Points detection, proposed by Laptev [111] searches for corners in the 3-dimensional video space (two spatial and one temporal dimension) based on the Harris corner detection. We apply this method on our data to find interest points. Figure 6.8 shows one frame of a video with the detected interest points marked by circles. The size of the circles represents the scale at which they are found.

The interest points are found using multiple spatio-temporal scales. Three levels of a spatial frame pyramid are applied (using a factor two subsampling). At each level, four combinations of spatial and temporal scales are obtained by using a Gaussian smoothing along the spatial and temporal dimension. The appearance and the motion around each interest point is then represented by the histograms of oriented gradients (HoG) and the histograms of optical flow (HoF). These are calculated from the neighborhood around the interest points. This is done by dividing this neighborhood in multiple cuboids in the spatial and temporal space. From every cuboid a four bin HoG and five bin HoF is calculated. However, in our study we only use the HoF as these represent the motion. Indeed, the appearance of the body part which is involved in the epileptic motion can differ and it does not contribute to the distinction between normal and epileptic movement. Furthermore, the appearance is also influenced by the blanket, whereas the motion is not as affected. In an initial test we combined the HoG and the HoF features, but this did



Figure 6.8: *STIPs detected within one frame, represented by the circles. The size of the circles indicate the scale at which they are found.*

not clearly lead to better results. As the neighborhood is divided into 18 cuboids, this results in 90 HoF features.

The motions that occur during the myoclonic seizures are often small. In the algorithm we can change the value of some parameters to increase the number of detected STIPs. A threshold T is used to eliminate interest points with a low probability. When this threshold is lowered, also noisier interest points are detected. In our tests we examine 4 different values of T being 10^n with $n \in \{-9, -10, -11, -12\}$. A value k is used in the Harris corner detection to specify the maximum sharpness of the detected corners. If this value is lowered, sharper corners (but also more line-shaped rather than corner-shaped points) will be detected. We test 2 values of k, namely $5 \cdot 10^{-4}$ and $5 \cdot 10^{-5}$. A third value we altered is the window length (25, 50, 125 or 250 frames), as the movement data is segmented in non-overlapping windows. Each window is considered as an example to be used in the training or test phase.

6.3.2 Classification

The HoF features derived from each interest point are used to classify the windows into myoclonic or normal movement. To do this, a bag-of-features approach is applied, as in Laptev *et al.* [112], on the features derived from the STIPs.

A subset of the data is used for the creation of the bag-of-features vocabulary. We use 50 clusters in this study; the cluster centers are determined by a k-means clustering on a training set. All the STIPs within one window will be assigned to their according cluster center, i.e. to the cluster center at the smallest distance in the HoF feature space. After this step, a histogram is built with each bin representing one cluster center. This histogram is fed to an SVM for classification. The use of this approach reduces the n 90-dimensional data points (STIPs) in a window, to one 50-dimensional point for each window, with n the number of STIPs in this window.

In the classification, two third of the data are used as training set and one third as test set. We use an RBF kernel for the SVM classification. We optimized the SVM model using the total sum of the misclassifications as cost function. All the tests are performed in a 15 fold randomization, randomizing the samples used in the training and test set.

6.3.3 Results

Results on global dataset

The first test is done on all the patients combined, and on every patient individually. In this test, we used one night of data for all patients.

When we look at the results, we observe that the patient-specific approach outperforms the approach combining the data of all patients. The results over all patients seem to improve with an increasing window length. The mean sensitivity over all settings of T and k increases from 11% to 37% when increasing the window length from 1 to 10 seconds. The best result is obtained using a 10 second window with $T = 10^{-10}$ and $k = 5 \cdot 10^{-5}$, reaching a sensitivity of 56% and a PPV of 72%. The performance improves when we consider the patients individually. For patient 8, an increased window length from 1 to 10 seconds raises the average sensitivity from 22% to 72% and the PPV from 45% to 77%. The parameter combination $T = 10^{-12}$ and $k = 5 \cdot 10^{-4}$ gives the best performance for most window lengths. For the ten second window a best sensitivity of 84% and a PPV of 93% is obtained.

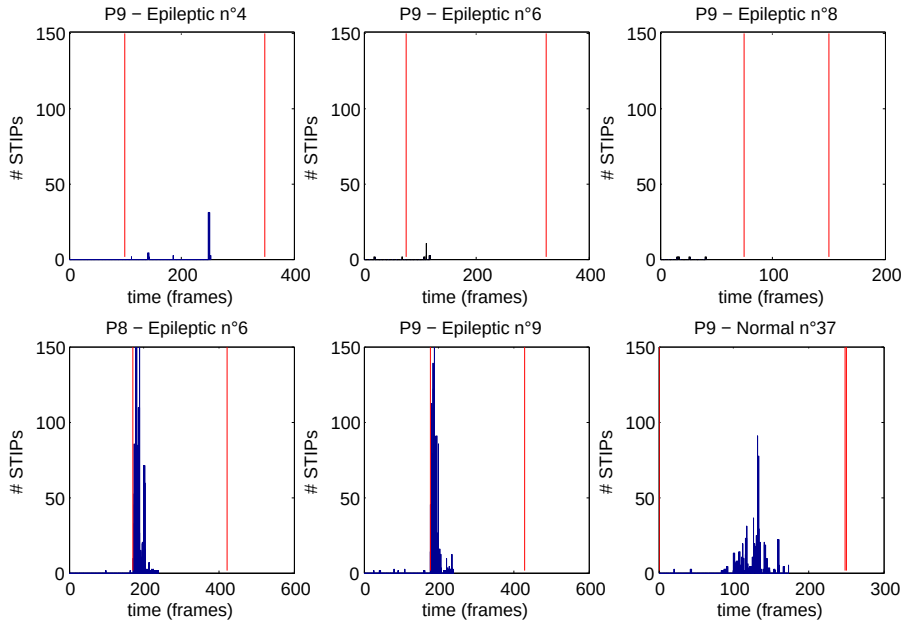


Figure 6.9: *STIPs of myoclonic and normal movements. The x-axis indicates the frame number of the video sequence. The y-axis represents the number of STIPs found in each frame. The boundaries (shown in red) indicate the beginning and ending of each seizure (or normal movement). The parameter combination generating the most STIPs is used here ($T = 10^{-12}$, $k = 5.10^{-5}$).*

However, for patients 3 and 9 no usable results could be obtained for half of the parameter combinations, as there were not enough seizures containing STIPs to train and test the classification model and obtain the performance results. Even when using the combination with the lowest threshold, some seizures had only a few or even no STIPs, which is also visible in Figure 6.9. For example, seizure 8 of patient 9 has no STIPs during the seizure. Of course it is hard or even impossible for the algorithm to correctly classify such small motions. Inspecting the video data, we could see that there were several seizures with a low intensity. Also, for some patients, myoclonic seizures manifested themselves during other, non-epileptic movement.

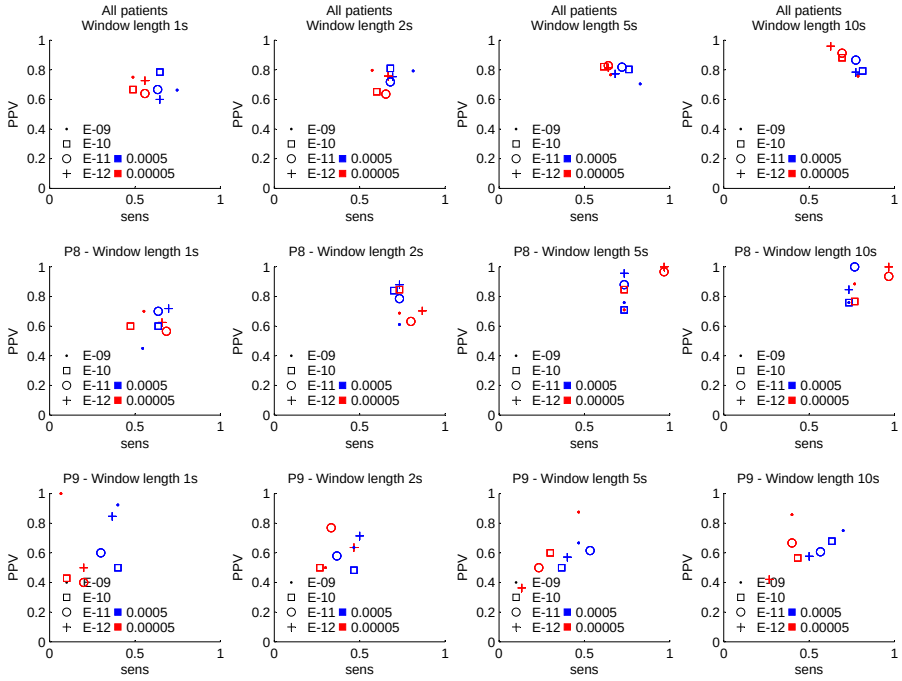


Figure 6.10: *Results for the patients with myoclonic seizures, on the subset of the data. The first row are the results for all patients combined. The second row for patient 8, and the third row for patient 9. From the left to the right the chosen window length increases from 1 second to 10 seconds. The markers in the figures indicate the performance on the sensitivity-PPV plot, and their shape and color represent their parameter combination. T ranges from 10^{-9} to 10^{-12} and the k-parameter was tested for a value of* 0.0005 *and* 0.00005.

Results on subset

To overcome the problems stated in the previous section, we validated the algorithm again on a subset of the data. Indeed, for the approach we use it is necessary that there is enough similarity within the seizure class. Furthermore some normal movement events contain jerk-like movement. To have an idea if this algorithm could work, we selected a subset of epileptic and non-epileptic data according to the following inclusion criteria:

- the jerks should be well visible in the video (thus resulting in a sufficient amount of STIPs);

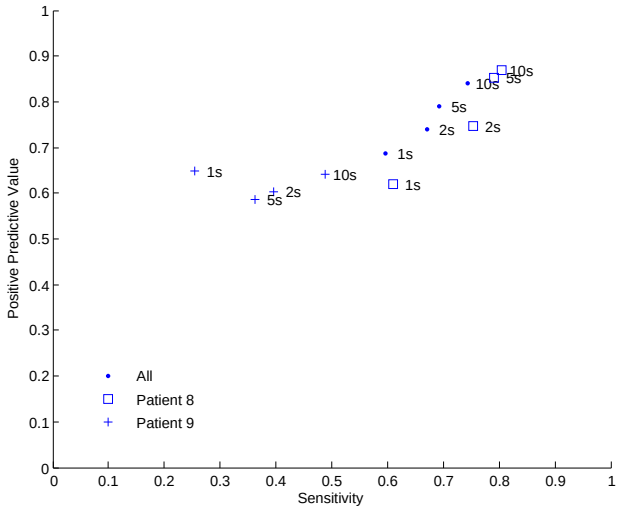


Figure 6.11: *The performance expressed in sensitivity and Positive Predictive Value on the subset of the data. The given performances are averaged out over the values of T and k, and show the influence of the window length. The window length is shown next to each data point, the marker shape indicates the corresponding dataset.*

- the jerks should be isolated in time (no influence of normal movement or other seizures);
- the normal sequences should not contain jerks.

With these criteria we composed a subset of data from 3 patients, containing in total 14 myoclonic jerks and 26 normal movements. For one patient, patient 8, we obtained 7 seizures, and for another patient, patient 9, we obtained 6 seizures. This permits us to do training and testing across both patients, by training the model on one patient and evaluating it on the other, to verify the models genericity. For patient 3, we obtained one seizure.

Figure 6.10 shows the results for patients 8 and 9 and the patients combined using different values for T and k and different window lengths. When taking only into account the influence of the window length we obtain Figure 6.11. The results shown are averaged out over the different values for T and k, to investigate the influence of the window length. For all patients individually and for the patients combined, the ten second window gives the best results. For the patients combined, the performance increases from an average sensitivity of 60% to 74% and a PPV

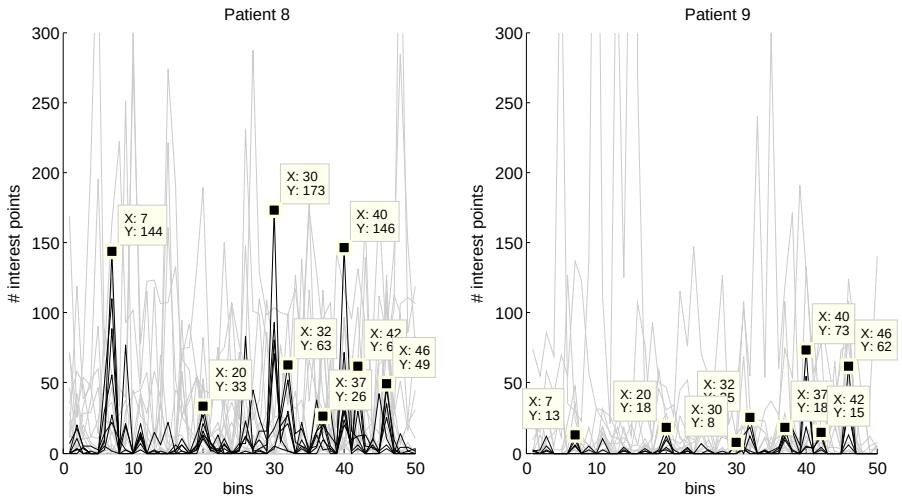


Figure 6.12: *The histograms for patient 8 (left) and patient 9 (right). The black and the gray curves are the histograms of the epileptic and normal movements, respectively. The bins that have an increased value for both patients are highlighted and labeled. It is also visible that in general the number of STIPs for the movements of patient 9 is lower than for patient 8.*

of 69% to 84%. The best performance reaches a sensitivity of 77% and a PPV of 87%, as visible in Figure 6.10. For patient 8 this is an increase from 61% to 80% and from 62% to 87%, respectively. For patient 9 the sensitivity increases from 25% to 49% but the PPV stays more or less the same (a little decrease) from 65% to 64%. But the results are more spread out with respect to the different threshold combinations. The most optimal combination for patient 8 is a window length of 5 or 10 seconds with $T = 10^{-12}$ and $k = 5 \cdot 10^{-5}$, so the combination with the highest number of STIPs as these are the lowest values for T and k we tested, and low values lead to more detected STIPs. This resulted in a sensitivity of 97% and a PPV of 100%. The best result for patient 9 is a sensitivity of 70% and a PPV of 75% obtained for a window length of 10 seconds with $T = 10^{-9}$ and $k = 5 \cdot 10^{-4}$. We see in this patient that the performance increases for this 10 second window with a decreasing value for T.

Figure 6.12 shows the 50 bins for the movements in patient 8 and patient 9, as explained in section 6.3.2. The histograms shown in black represent the myoclonic jerks, the ones in gray are normal movement. The histograms here are not normalized, so the y-axis shows the absolute number of STIPs that are assigned to the different clusters (bins). We can observe that there is a similarity in the

bins (e.g. bin 20, 40 and 46) with the highest number of STIPs assigned to, for the epileptic movement in both patients. The histograms shown here are generated using a window length of 10 seconds, a threshold of 10^{-10} and a k-parameter of $5 \cdot 10^{-4}$. Because there is similarity in the histograms, it means that it should be possible to train the classification model on one patient, and test it on the other one. We have tested this for all the different parameter combinations, using all the data from both patients (and thus without a randomization).

When inspecting the performance we notice that there is a large spread in the results regarding the parameter combination. When averaging out the results for each window length, we also can see here that a larger window length increases the performance. When we train the classification model on patient 8 and test on patient 9, the average sensitivity increases from 31% to 56% from a 1 second window to a 10 second window, whereas the PPV stays the same at 78%. The other way around we see an increase of the sensitivity and PPV from 50% to 64% and from 58% to 80%, respectively. In the second test (training on patient 9 and testing on patient 8) we observe that using a threshold of 10^{-12} and a k-parameter of $5 \cdot 10^{-4}$ gives the best performance, and for a 10 second window it even reaches a sensitivity and PPV of 100%.

6.3.4 Discussion

This study shows that we are able to detect the typical jerky movement of myoclonic seizures. The performance of the method increases with the window length. Also, the patient-specific approach gives better results compared to the group specific approach.

The influence of increasing or decreasing T or k is not always consistent. The results often differ when using a different parameter combination. Adding more STIPs (lowering T or k) can increase the number of characteristic features. But lowering the thresholds too much increases also the noisy features. Depending on how intense the motion is, the ideal value differs.

The results we get from the cross training and testing on two patients are encouraging, although the result depends strongly on which combination of parameters is used. By making use of the sub-selection of the data, we simplified the problem to some extend. A lower performance is expected when the test is executed on full nights of data. A way to improve the classifier is by retrain it using hard negatives, the seizures that were missed by the classifier in an earlier iteration.

The calculation of the STIPs and the derived features is computationally intensive. Real-time processing is not feasible for the moment. However, if we could downscale the algorithm and tune it for our specific application, an increase in speed may be obtained. Furthermore, the use of the windowing limits the real-time detection to the length of the chosen window, if we e.g. use a window of 5 seconds, since all the features can only be calculated after completion of the recording of the corresponding epoch.

The application of the STIP method on the myoclonic seizure data gives promising initial results. The best obtained result over all the patients combined reaches a sensitivity of 77% and a PPV of 87%. We can conclude that a longer window length gives better results in detecting seizures. However if the seizures are too subtle, the method seems not able to detect them.

6.4 Comparison between both methods and conclusion

In this chapter we compared two methods for detecting myoclonic jerks in video. The method that makes use of the clustering of the flow vectors is strongly dependent on a good clustering of the vectors and the clusters can be merged and split up from frame to frame. This can make the tracking more difficult.

Furthermore, this approach combines different methods, of which some are computationally intensive. An accurate calculation of the flow vectors for example, requires a high number of iterations, to this end, 1000 iterations are used in this chapter. Also the clustering step can require a significant calculation power as shown in Table 6.1. All these steps have to be performed on each video frame which makes real-time processing not feasible.

Also the calculation of the STIP features is computationally intensive. For the classification in this approach, the calculation of the bag-of-features vocabulary is hard. However, this only has to be done in the training phase. Once the words (the cluster prototypes) are learned, the classification is fast.

The features that are used in the first method are rotational invariant. It is obvious that duration or amplitude related features are rotational invariant, but this counts also for the phase changes, as the relative shift in motion direction is used. Of course the features are in some way dependent on the orientation of the patient, as body parts may be occluded (for example the torso can occlude arm movements) or movements may occur perpendicular to the image plane.

The features that are used in the STIP method, in contrast, are dependent on the orientation, as the flow vectors are directly used in the features (or more specific the histogram of flow vectors). This issue can be overcome when there are enough training samples available, to cover all possible orientations of the patient. Or the issue lapses when the patient always appears in the video images in the same orientation.

An advantage of the tracking in time, which is done after clustering the flow vectors, is that patterns that occur in specific body parts are easier to detect. Especially when they are spread over a longer period of time. It has to be noted, however, that this tracking can fail and the original part that was followed can be lost (e.g. by accidentally tracking another part after occlusion). A way to improve this tracking is by taking into account the size and shape of the cluster, presuming that these features do not change too much from frame to frame.

In both methods, a number of parameters has to be chosen by the user. In the first method, these are for example the band width for the mean-shift clustering or the number of iterations in the calculation of the optical flow. In the STIP method not only the values for k, T and the window length have to be chosen, but also, for example, the number of scales in which the interest points are searched for. This relatively large number of tunable parameters makes it hard to find the optimal setting.

An advantage of the STIP method is that thanks to the bag-of-features approach, missing some STIPs or including some noisy points has not so much influence on the classification. The method is more robust to noise compared to the clustering of the flow vectors, as a small change in the vector field can lead to a different clustering, and thus resulting in different tracks or features. However, the features derived from the STIP method are more difficult to interpret compared to clustering the flow vectors. In the latter method the motion of the body parts can be seen in the motion tracks and the derived features are more intuitive. But in general, it is advisable to use the STIP method because of the higher robustness against noise.

A downside of video detection in general is that the video cannot record subtle movement that occur under the blankets. For these type of movements, other modalities have to be used which may include ACMs, thermal infrared cameras (if the blankets don't mask the heat too much), or depth cameras.

An extension to this work could be to reconstruct movement under the blankets by modeling the patient's posture by making use of prior knowledge (e.g. if we know the orientation of the head, we can estimate how the patient is positioned under the blankets). This could help defining the exact movement when the patient is

occluded.

Chapter 7

Multimodal seizure detection

Chapter 5 and Chapter 6 discussed the detection of epileptic seizures based on one modality, ACM data and video, respectively. This chapter investigates if a combination of both modalities improves the performance. Section 7.2 elaborates on the methodology of selecting the (synchronized) epochs for both modalities and how they are combined. Section 7.3 compares the results of the uni-modal and multi-modal approach, and to conclude, in Section 7.4 a discussion and conclusion about this multi-modal integration is given.

7.1 Introduction

In the previous chapters, we have focused on the detection of seizures based on the usage of only video recordings and only ACMs. In this chapter we want to integrate both modalities to investigate if a better performance is feasible when using a multi-modal approach. For the ACM based approach, we used the SVM classifier with the Kruskal-Wallis features, discussed in Section 5.3. The detection based on video recordings [53] is based on the spatio-temporal interest point (STIP) method proposed by Laptev [111, 112], which was discussed in Section 6.3.

Conradsen *et al.* [41] apply a multi-modal approach for seizure detection, using EMG, acceleration and angular velocity. Their multi-modal approach (based on simulated data) performs better than the uni-modal approach (based on real seizures). Other studies already investigated for example the integration of EEG and ECG in epileptic seizure detection [76]. In that work, an early integration

(combining features from both modalities or *feature fusion*) and late integration (combining output from classifiers on both modalities or *decision fusion*) are considered. In our study we also investigate an early and a late integration of the data. Although fusion at a late stage in the processing stream naturally incurs performance loss, it is a feasible solution when fusion at a lower level is not practical [184]. Integration of video data with other modalities is for example already used in speech recognition [11, 64, 78].

We have tested this multi-modal integration on the hypermotor seizures and normal nocturnal movements of patient 6. A total of 7 seizures and 381 normal movements were used in this study.

7.2 Multi-modal integration

To be able to objectively compare both approaches, based on ACMs and on video, we use the same segmentation of the data for both modalities. This means that the segmentation of the data, both ACM and video, is based on the movements in the ACM signals. The training and testing is done in a 10-fold randomization, just like in our prior studies. This means that we randomly select a number of normal and epileptic movements for the training and test set and perform the modeling and validation 10 times using a different combination of movements, and average out the obtained results. This makes the results less dependent on the division of the data in a training and test set. Also in this randomization, the folds are stratified, meaning that the number of seizures and normal movements is the same for each randomization. We use two thirds of the data as the training set, and one third as test set. The same randomizations are used in both approaches and for the ACM/video integration.

The ACM and video detection methods are the ones that are already described in Chapter 5 and Chapter 6, respectively. For the ACM detection we used the SVM approach as in Section 5.3, for the video detection we also applied the SVM using the STIP features of Section 6.3.

In a first test we combined the normalized features from both approaches, the three features from the ACM data, and the bag-of-features output from the video approach, in an early integration. In a second test, we combined the outputs of both individual classifiers from the video and ACM detection, in a late integration. The output values of both classifiers give a probability of the sequence belonging to the epileptic seizure class, as we use the libsvm implementation for extending SVM to give probability estimates based on Wu *et al.* [193] and Lin and Weng [116].

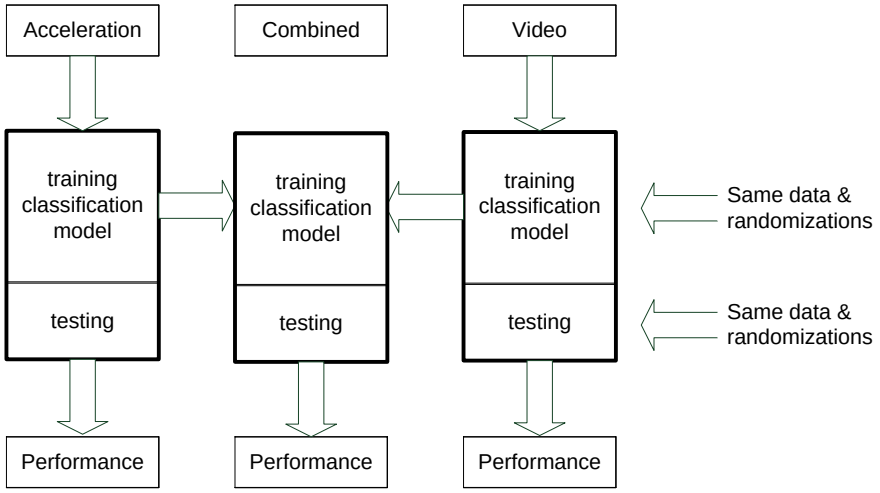


Figure 7.1: *Overview of the late integration of both modalities. The high level classifier makes use of the outputs of both the video detection and the ACM detection algorithm.*

To integrate the ACM and the video detection, first we made sure that the data to process is exactly the same. In the training and test sets, the same randomization of both modalities is used and the video data is not cut into smaller windows, the entire movements are classified as one event. We used all the available data for training and testing, so in contrast to Section 5.3 and Section 5.4 where we use 2 seizures and 33 normal movements for testing to be able to objectively compare different patients, we have a larger imbalance in the data as we have 3 seizures and 127 normal movements (one third of the data) in the test set.

In the early integration, the features are used to train an SVM model, similarly as the model for each individual modality.

In the late integration, the outputs of the two detection algorithms are used as an input in a new Linear Discriminant Analysis (LDA) classifier. The LDA classifier [131] finds a linear transformation that maximizes the inter-class variance and minimizes intra-class variance to get the best separation between the classes. The probability that a movement belongs to the seizure class in both approaches is used as input features in the integrated classification. The LDA model is trained based on these new features making use of the training data. In the test phase, the movement events are first classified by the ACM and video detection algorithms, and their outputs are again used for the LDA classifier, which produces the final classification. An overview of the working method for the late integration is shown

Table 7.1: *Results of the single and multi-modal approach, in terms of sensitivity, positive predictive value and specificity. The last column gives the cost value of each approach, using the cost function of Equation 5.10. These results were obtained on the data from patient 6.*

	sens	std	PPV	std	spec	std	-cost
ACM	83.33	23.57	100.00	0.00	100.00	0.00	266.67
video	70.00	33.15	97.22	8.33	99.92	0.25	237.22
combi (features)	**83.33**	23.57	**96.00**	12.65	**99.84**	0.50	270.83
combi (decisions)	**86.67**	23.31	**97.50**	7.91	**99.92**	0.25	262.67

in Figure 7.1.

7.3 Results

Table 7.1 shows an overview of the results for the uni-modal and multi-modal detection. We see that the PPV is typically larger in all results compared to the sensitivity. Of the uni-modal detection, the ACM detection seems to give the best results. This is also visible in the cost value in the last column, which is calculated similarly to the cost in Section 5.3.

The combined detection using the early integration seems to give a lower performance than the ACM detection alone. This means that the video features does not have any added value in this integration. It has to be said that the influence of the video features is larger (50 features) compared to the ACM features which can partially explain this lower performance. The late integration of both modalities does have a po sitive influence on the performance, as visible the overall cost values. Although there is a decline in the PPV from 100.00% to 97.50% compared with the ACM detection, the sensitivity increases to 86.67%.

7.4 Discussion

The test is now performed on an initial dataset to see whether an integrated approach would give better results. We used the data from one patient, as we have good ACM and video data in these datasets, in some other datasets of patients we used in Chapter 5 the video data had artifacts what makes it more difficult to use

for video detection. Although the first indications are promising, this approach should be tested on a larger group of patients and seizures.

For the moment we have integrated ACM and video signals, but in the future, it would be interesting to also integrate other modalities. The datasets we have collected also contain ECG, EMG, EEG, electrooculogram (EOG) and audio data. Combining multiple modalities can improve the performance. However, we have to keep in mind that adding more sensors would make the recording of the data more cumbersome, which is contrary to the research goals of making the monitoring easier and more comfortable for the patient.

There are also other ways to combine the input from different classifiers. For example voting methods, Bayesian fusion methods on the soft outputs of the classifiers (as available in our approach) or using Fuzzy integrals or Fuzzy templates [156]. If we have a dataset with more positive examples, we can integrate both modalities based on the stacked generalization [192]. This would be a logical extension of our current approach. Up till now, we used only one training set for learning the high level model, instead of the separate training sets for the low level and high level modeling in the stacked generalization.

Another interesting possibility for the fusion on the feature level would be applying the bag-of-features approach on the 12 channel ACM data. This is feasible by detecting spikes in the signals and derive features around them or simply by taking the raw input data around each spike. This would make it easier to integrate both modalities. Other ways that the multi-modality can be used is for example performing the training on the multi-modal data (in our case e.g. video and ACMs) and test it on unimodal data (e.g. only video). In this way, the different modalities can provide extra information to train a better model, but once the detection algorithm is employed in a real setting, fewer modalities have to be registered which makes the detection system easier and more comfortable to use.

Chapter 8

Conclusion

Don't cry because it's over,
smile because it happened
anonymous

This chapter summarizes the most important findings of this dissertation, listed chapter per chapter in Section 8.1. In a last part in Section 8.2, an outlook is given on interesting topics for future research.

8.1 Concluding remarks

The goal of this thesis was to investigate if it is possible to detect nocturnal seizures in epileptic children without interfering too much in the environment of the patient. Furthermore, in several types of seizures, and certainly the most violent ones, a motor component is present. For these reasons we applied minimally intrusive sensors that are able to register the motion of the patient, ACMs attached to the limbs, and a video camera that is placed in the corner of the room.

Next the main findings will be discussed chapter by chapter.

- **Data preprocessing:** First, the data was preprocessed and segmented in movement epochs, both on the ACM and the video signals. In both modalities the movements were well detectable, although subtle movements that are occluded in the video are easier to detect in the ACM data. Furthermore,

the computational load of the video processing is higher compared to the ACM data processing. The reduction in ACM data after the preprocessing step varies between 97% and 98.5%, which helps to speed up the further processing of the data.

The movement detection is already a first indication of the behavior (or a change in behavior) of the patient during the night, as for example abnormally long movements can already denote seizures or other relevant events.

The recommendations for the video acquisition setting are that the spatial resolution can be downsampled to 320×240 pixels and the temporal resolution to 10 frames per second. This reduces the processing time, but the downsampling in time can have a negative influence on the detection of specific patterns such as jerk-like movement. Therefore, it is only recommended for the movement detection step.

For this movement detection, we found out that a variable threshold based on the noise level gives better results compared to a fixed threshold. The sensitivity and PPV reaches 100.00% and 95.40%, respectively. Also, the parameter values for the variable threshold are robust over different datasets.

Based on the video data, the detection of respiration is also possible if the chest is visible in the camera. This can be an added value in the preprocessing of the video data, for the detection of seizures.

- **ACM detection:** In the detection of seizures using ACMs, we applied three different methods. From our results we could conclude that the choice of features is crucial for a good performance. The best features can differ from patient to patient because of a different manifestation of the seizures. By choosing the optimal features, a higher performance is possible for detecting hypermotor seizures, regardless of the classifier. However, the different classifiers we tested had different properties. The outlier detection method is, for example, well suited to be implemented in a home situation without too much prior knowledge of the seizure manifestations. If more data is available, the SVM approach is preferable, as it gives better results compared to the other techniques when using SMOTE. The simple thresholding gives the least good performance.

  The issue of imbalance of the data has different implications on the proposed methods. For the outlier detection, the imbalance is necessary but for the SVM approach it complicates generating a model. For this latter approach, the resampling of the minority class gives the best results to overcome the issue of imbalance compared to adjusting the weighting (of the penalization of misclassification) for both classes. In contrast to binary classification methods,

which use information of both classes (seizures and normal movements) to estimate a classification model, the outlier detection does not need labeled data for constructing the model. Also because the fact that the seizure data does not have a significant influence on the probability density function of the movements.

The SVM and outlier detection approach both perform well with a sensitivity above 90% and a PPV between 50% and 90%. The performance of the outlier detection can even be improved, as the used threshold value of 95% is a rather conservative estimation.

- **Video detection:** In the video detection we tested the detection of myoclonic seizures. Here we noticed that it is hard to detect subtle movements in the video. Furthermore, the occurence of myoclonics during other (non-)epileptic movements complicates the detection. On a subset of seizures in the STIP detection, the detection results were satisfying, and a promising outcome is the similarity in pattern of myoclonics in the considered patients. This may indicate that a common model for this type of seizures is able to detect seizures over different patients.

  Both the clustering of flow vectors and the STIP method are computationally intensive. The advantage of using the flow vector approach is that tracking over a longer period of time allows to detect specific motion patterns in body parts, especially when they last longer. Furthermore, the motion tracks are easy to interpret and the derived features, based on duration, amplitude and change in motion direction, are intuitive. However, the STIP method is more robust against noise, which makes this method preferable over the former.

- **Multi-modal:** From the multi-modal approach we can conclude that the video detection (sensitivity of 70.00% and PPV of 97.22%) in general performs poorer compared to the ACM detection (sensitivity of 83.33% and PPV of 100.00%). An objective comparison could be made as the same type of seizures, the same segments and the same classification method was used. Of course the ACMs are directly attached to the body, and can therefore record the motion more accurately. Also the sampling frequency of this modality is higher and nonetheless, the processing of the data is less computational intensive compared to video. On the other hand, the video detection is a contactless method, and its main disadvantage is also an advantage as contactless registration of the movement is even less intrusive for the patient which really breaks down barriers for introducing the monitoring in a home situation. Furthermore, we can conclude that a combined approach can yield slightly better results with respect to the sensitivity compared to the

uni-modal approach, which was illustrated in the late integration of both modalities, which resulted in a sensitivity of 86.67% and a PPV of 97.50%. The early integration showed somewhat poorer results with a sensitivity of 83.33% and a PPV of 96.00%.

8.2 Future Research

The future research part focuses on two domains. First, we want to increase the performance of the methods we used on patients where the performance was lower. Furthermore the proposed methods can also be applied in a larger scope, in other domains than epilepsy. This will be discussed in Section 8.2.1. Second, these methods should be applied in a real home setting to evaluate their clinical impact. Section 8.2.2 elaborates on the first results of such a home monitoring system.

8.2.1 Research directions

Integrate multiple modalities

Monitoring epileptic seizures only based on modalities that record the motion of the patients means that only a subgroup of the patients (subjects that suffer from clearly distinguishable motor seizures) is eligible to be monitored by such a system. Moreover, the seizures may resemble non-epileptic activity such as the hypnic jerk that is experienced before falling asleep, but that does not denote abnormal neurological activity. Also, patients that suffer from epilepsy sometimes show damage in the motor cortex of the brain, which can lead to the occurrence of disturbed motion patterns that are not related with epilepsy. This makes it hard to classify the movements as epileptic or non-epileptic without any knowledge of other physiological parameters. One way to try to overcome this problem is by integrating multiple modalities.

Therefore, it is advisable to integrate the ACM data with data from other sensors, such as ECG, audio, video or skin conductivity. But there is a trade-off between the number of sensors and the comfort for the patient during sleeping, unless contactless sensors are used. To keep the number of sensor nodes attached to the patient as low as possible, multiple sensors can be integrated in one node, e.g. the sensors for measuring the skin conductivity and the heart rate can be integrated in the wrist band that already measures the acceleration of the arm.

Clustering of specific movement patterns

Another interesting way to proceed is to cluster nocturnal motion patterns in an unsupervised way, apart from if it is epileptic activity or not. In this manner, patterns that occur often during the night can be identified, for example jerk-like movements that have an epileptic or non-epileptic origin, automatisms such as lip smacking or specific head motions. In this way the movements are not categorized based on their electrical origin, but only on their clinical manifestation. This generates an overview of the nocturnal activities of the subjects during the night and aids in constructing a global picture of the possible conditions of the subject, such as nocturnal epilepsy but also sleep related problems as sleep apnoea or restless leg syndrome.

Patient specific approach versus group specific approach

In our research, we mainly investigated the patient specific approach, where we built the classification models based on patient specific information. Although we built group specific models, for example for the myoclonic seizure detection using the STIP method, but on a relatively small group of patients. However, if we could use group specific models, that give (almost) equal performance to patient specific models, this would reduce the need for a lot of patient specific information. This means, in turn, that the developed algorithms can be employed much easier for more patients. To increase the performance of the classification model, a combination between group specific and patient specific information can be used.

8.2.2 Integrated detection in real setting

To study the clinical impact of a system that can monitor the movements of a subject during the night, we developed a mobile stand-alone system that integrates video, ACMs and a radar sensor which is easy to install in a subject's room in a home situation. The system is able to detect and record nocturnal movements based on these modalities, and gives an overview of the collected information over the night. The ACM detection explained in Chapter 3 and the video detection from Chapter 4 are integrated in this system and can process the data in real-time.

So far we have collected data from 3 patients in a semi-home situation, that is, they were staying at the rehabilitation center or institution for a longer period of time in their own room which mimics the situation at home. Figure 8.1 gives a

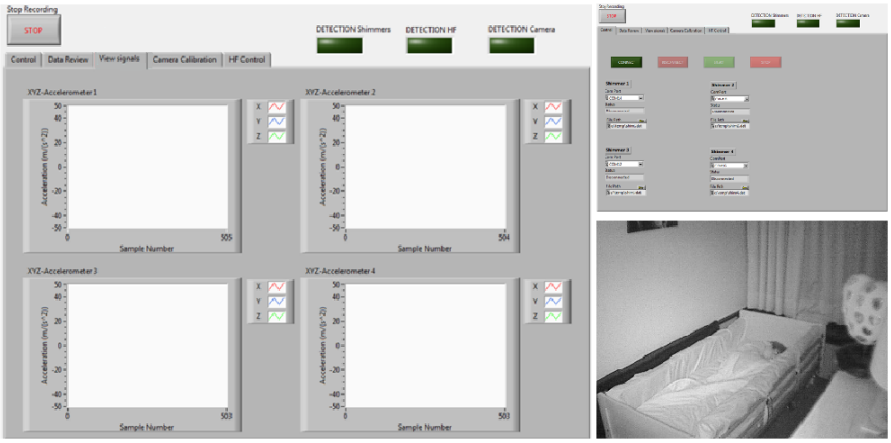


Figure 8.1: *Snapshot of the system interface for data acquisition. The left figure shows the screen that displays the ACM signals. The upper right figure shows the Shimmer status screen, which indicates if the wireless nodes are connected and streaming data. The lower right figure is a snapshot of the video that is captured during the night with this system.*

view of the interface of the system, together with a camera view of the infrared camera that monitors the movements.

A first investigation of the acquired data learns us that the three modalities separately are able to detect most of the movement, respectively 80.34%, 84.91% and 86.46% of the movements are detected for the ACMs, camera and radar sensor assuming that the combination can detect all movements.

This platform lends itself to expanding the functionalities with developments in other parts of this thesis. A valuable extension would be integrating the detection of respiration that is explained in Section 4.3, and also the seizure detection methods itself. The only obstacle for the video detection is that these algorithms are not yet feasible in real-time.

The number of seizures a patient has during the night offers useful information for the neurologist. This permits a better observation of the patient, which can increase the quality of life, thanks to the better insight in the effect of medication. It will also benefit the parents because they can be warned if their child has a seizure, and can offer help to their child during or after the seizure.

Appendix A

Additional results

This appendix gives an overview of the results that were obtained for the accelerometer detection for example with respect to a different feature set, or the influence of the bandwidth in the outlier detection based on the PDF estimation.

Table A.1: *Results of performance of SVM classifier using different methods to overcome the problem of imbalance, using the five best features according to the Kruskal-Wallis test.*

Balance factor								
linear	patient	**sens.**	std.	**PPV.**	std.	spec.	std.	-cost
kernel	2	**80.00**	25.82	**93.33**	14.05	99.79	0.44	253.33
	5	**0.00**	0.00	**0.00**	0.00	99.37	1.66	0.00
	6	**80.00**	42.16	**89.58**	19.80	99.68	0.71	249.58
	all	**10.00**	21.08	**16.67**	25.82	98.95	1.57	36.67
								167.64
RBF	patient	**sens.**	std.	**PPV.**	std.	spec.	std.	-cost
kernel	2	**90.00**	21.08	**90.00**	16.10	99.68	0.51	270.00
	5	**5.00**	15.81	**25.00**	50.00	99.68	0.51	35.00
	6	**50.00**	52.70	**100.00**	0.00	100.00	0.00	200.00
	all	**15.00**	24.15	**14.29**	20.25	98.21	1.79	44.29
								168.33
Probability density estimation oversampling								
linear	patient	**sens.**	std.	**PPV.**	std.	spec.	std.	-cost
kernel	2	**90.00**	21.08	**68.67**	29.15	98.63	1.41	248.67
	5	**70.00**	34.96	**12.44**	6.50	88.95	3.92	152.44
	6	**100.00**	0.00	**83.33**	22.22	99.37	0.89	283.33
	all	**65.00**	24.15	**12.90**	5.84	89.47	4.30	142.90
								228.15
RBF	patient	**sens.**	std.	**PPV.**	std.	spec.	std.	-cost
kernel	2	**95.00**	15.81	**74.00**	23.98	98.95	1.11	264.00
	5	**60.00**	39.44	**12.99**	10.08	91.05	4.25	132.99
	6	**100.00**	0.00	**80.00**	26.99	99.05	1.44	280.00
	all	**45.00**	36.89	**19.41**	13.27	95.79	3.33	109.41
								225.66
SMOTE oversampling								
linear	patient	**sens.**	std.	**PPV.**	std.	spec.	std.	-cost
kernel	2	**90.00**	21.08	**83.33**	22.22	99.47	0.74	263.33
	5	**45.00**	36.89	**9.84**	8.18	91.68	3.27	99.84
	6	**85.00**	33.75	**87.04**	20.03	99.58	0.74	257.04
	all	**55.00**	28.38	**17.22**	10.35	92.21	5.39	127.22
								206.74
RBF	patient	**sens.**	std.	**PPV.**	std.	spec.	std.	-cost
kernel	2	**90.00**	21.08	**75.67**	26.44	99.05	1.16	255.67
	5	**30.00**	25.82	**15.59**	19.47	95.68	3.42	75.59
	6	**80.00**	34.96	**87.04**	20.03	99.58	0.74	247.04
	all	**45.00**	28.38	**16.58**	14.53	94.00	3.75	106.58
								192.76

Table A.2: *Results of the accelerometer classification using the feature set selected with the SVM approach. For each patient the sensitivity (sens), the PPV and the specificity (spec) are shown, together with their respective standard deviation over 10 randomizations*

patient	sens. (%)	std. (%)	PPV (%)	std. (%)	spec. (%)	std. (%)
1	0.00	0.00	0.00	0.00	99.70	0.96
2	85.00	24.15	95.00	15.81	99.39	1.92
3	0.00	0.00	0.00	0.00	99.39	1.28
4	75.00	26.35	72.50	30.94	97.27	3.33
5	27.78	36.32	35.19	44.44	96.97	2.47
6	100.00	0.00	62.00	22.84	95.15	3.56
7	94.44	16.67	50.74	10.77	94.85	2.49
Total	*54.60*	*14.79*	*45.06*	*17.83*	*97.53*	*2.29*

Table A.3: *Influence of the bandwidth on results of the accelerometer classification using the feature set selected with the Kruskal-Wallis method.*

	$\beta = 0.25$		$\beta = 1$		$\beta = 4$		$\beta = 16$	
patient	sens. (%)	PPV (%)	sens. (%)	PPV (%)	sens. (%)	PPV (%)	sens. (%)	PPV (%)
1	100.00	40.33	100.00	57.33	100.00	58.00	100.00	58.00
2	100.00	73.00	100.00	80.00	100.00	76.33	100.00	76.33
3	100.00	70.67	100.00	85.00	100.00	86.67	100.00	86.67
4	100.00	56.33	95.00	76.67	100.00	80.00	100.00	80.00
5	60.00	47.83	44.44	44.44	50.00	50.00	50.00	50.00
6	100.00	65.22	100.00	69.00	100.00	70.00	100.00	69.00
7	100.00	58.00	100.00	71.33	100.00	69.67	100.00	69.67
Total	94.29	58.77	91.35	69.11	92.86	70.10	92.86	69.95

Table A.4: *Results obtained by using the SVM-features for the accelerometer based detection on the PDF but replacing the* 'lowest peak frequency over all channels' *feature by the* length of the epoch.

patient	sens. (%)	std. (%)	PPV (%)	std. (%)	spec. (%)	std. (%)
1	100.00	0.00	62.33	16.56	95.76	2.56
2	100.00	0.00	91.67	18.00	99.09	2.05
3	100.00	0.00	90.00	16.10	99.09	1.46
4	55.00	15.81	68.33	27.72	97.58	2.39
5	22.22	36.32	22.22	37.27	95.45	3.27
6	100.00	0.00	71.67	15.81	97.27	1.72
7	100.00	0.00	67.33	19.68	96.36	2.78
Total	82.46	7.45	67.65	21.59	97.23	2.32

Bibliography

[1] E.H. Adelson, J.R. Bergen, "The extraction of spatiotemporal energy in human and machine vision," *In Proc. of the IEEE Workshop on Visual Motion*, pp. 151-156, 1986.

[2] R.A.P.P. Adriaans, C. Ungureanu, M.J.P. van Bussel, M.J. Zwarts, J.B.A.M. Arends, R.M. Aarts, "Strategies for real-time detection of nocturnal epileptic seizures," *In Proc. of the Annual Symposium IEEE EMBS Benelux Chapter*, 2011.

[3] F.Q. AL-Khalidi, R. Saatchi, D. Burke, H. Elphick, "Tracking human face features in thermal images for respiration monitoring," *In Proc. of the IEEE/ACS International Conference on Computer Systems and Applications (AICCSA)*, pp. 1-6, 2010.

[4] E. Alpaydin, *Introduction to Machine Learning*, MIT Press, 2004.

[5] P. Anandan, A computational framework and an algorithm for the measurement of visual motion, *International Journal of Computer Vision*, Vol. 2(3), pp. 283-310, 1989.

[6] N.J. Azar, T.F. Tayah, L. Wang, Y. Song, B.W. Abou-Khalil, Postictal breathing pattern distinguishes epileptic from nonepileptic convulsive seizures, *Epilepsia*, Vol. 49(1), pp. 132-7, 2008.

[7] L. Bao, S. Intille, Activity Recognition from User- Annotated Acceleration Data, *Lecture Notes Computer Science*, Vol. 3001, pp. 1-17, 2004.

[8] H. Barman, L. Haglund, H. Knutsson, G. Granlund, "Estimation of velocity, acceleration and disparity in time sequences," *In Proc. of the IEEE Workshop on Visual Motion*, pp. 44-51, 1991.

[9] J.L. Barron, D.J. Fleet, S.S. Beauchemin, T.A. Burkitt, "Performance of Optical Flow Techniques," *In Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '92)*, pp. 236-242, 1992.

[10] J.L. Barron, D.J. Fleet, S.S. Beauchemin, Performance of optical flow techniques, *International Journal of Computer Vision*, Vol. 12(1), pp. 43-77, 1994.

[11] S. Basu, C. Neti, A. Senior, N. Rajput, L. Subramanium, A. Verma, "Audio-Visual Large Vocabulary Continuous Speech Recognition in the Broadcast Domain," *In Proc. of the IEEE Workshop on Multimedia Signal Processing*, 1999.

[12] S.S. Beauchemin, J.L. Barron, The Computation of Optical Flow, *ACM Computing Surveys*, Vol. 27(3), pp. 433-467, 1995.

[13] G. Becq, S. Bonnet, L. Minotti, M. Antonakios, R. Guillemaud, P. Kahane, Classification of epileptic motor manifestations using inertial and magnetic sensors, *Comput Biol Med*, Vol. 41(1), pp. 46-55, 2011.

[14] A.T. Berg, S.F. Berkovic, M.J. Brodie, J. Buchhalter, J.H. Cross, W. van Emde Boas, J. Engel, J. French, T.A. Glauser, G.W. Mathern, S.L. Moshé, D. Nordli, P. Plouin, I.E. Scheffer, Revised terminology and concepts for organization of seizures and epilepsies: report of the ILAE Commission on Classification and Terminology, 2005-2009, *Epilepsia*, Vol. 51(4), pp. 676-85, 2006.

[15] C. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.

[16] M. Black, P. Anandan, "A framework for the robust estimation of optical flow," *In Proc. of the International Conference on Computer Vision (ICCV93)*, pp. 231-236, 1993.

[17] H. Blockeel, *Machine Learning and Inductive Inference*, Acco, 2009.

[18] A.S. Blum, Respiratory physiology of seizures, *J Clin Neurophysiol.*, Vol. 26(5), pp. 309-15, 2009.

[19] W.T. Blume, H.O. Lüders, E. Mizrahi, C. Tassinari, W. van Emde Boas, J. Jr. Engel, Glossary of descriptive terminology for ictal semiology: Report of the ILAE task force on classification and terminology, *Epilepsia*, Vol. 42, pp. 1212-1218, 2001.

[20] S. Bonnet, R. Héliot, A magnetometer-based approach for studying human movements, *IEEE Trans Biomed Eng*, Vol. 54(7), pp. 1353-5, 2007.

[21] B. Bonroy, P. Schiepers, G. Leysens, D. Miljkovic, W. Wils, L. De Maesschalck, S. Quanten, E. Triau, V. Exadaktylos, D. Berckmans, B. Vanrumste, "Acquiring a dataset of labeled video images showing discomfort in demented elderly," *Telemed J E-Health*, Vol. 15(4), pp. 370-378, 2009.

[22] G. Bradski, A. Kaehler, *Learning OpenCV*, O'Reilly Media, 2008.

[23] N. Breslow, "A generalized Kruskal-Wallis test for comparing K samples subject to unequal patterns of censorship," *Biometrika*, Vol. 57, pp. 579-594, 1970.

[24] T. Brezmes, J.L. Gorricho, J. Cotrina, "Activity Recognition from accelerometer data on mobile phones," *In Proc. of the 10th International Work- Conference on Artificial Neural Networks (IWANN '09)*, pp. 796-799, 2009.

[25] T. Brox, A. Bruhn, N. Papenberg, J. Weickert, "High accuracy optical flow estimation based on a theory for warping," *In Proc. of the European Conference on Computer Vision (ECCV)*, Vol. 3024, pp. 25-36, 2004.

[26] E.J. Burt, C. Yen, X. Xu, "Multiresolution flow-through motion analysis," *In Proc. of the Conf. Comput. Vis. Part. Recog.*, pp. 246-252, 1983.

[27] C. Carlson, V. Arnedo, M. Cahill, O. Devinsky, "Detecting nocturnal convulsions: efficacy of the MP5 monitor," *Seizure*, Vol. 18(3), pp. 225-7, 2009.

[28] A. Casson, D. Yates, S. Smith, J. Duncan, "Wearable Electroencephalography," *Engineering in Medicine and Biology Magazine*, Vol. 29, pp. 44-56, 2010.

[29] C.C. Chang, C.J. Lin, LIBSVM: A library for support vector machines, *ACM Transactions on Intelligent Systems and Technology*, Vol. 2(3), pp. 27:1-27:27, 2011.

[30] D. Chapman, R. Panelli, J. Hanna, T. Jeffs, *Sudden unexpected death in epilepsy: continuing the global conversation*, Camberwell, Australia, 2011.

[31] J. Chase, B. Nelson, J. Bodily, Z. Wei, D.J. Lee, "Real-time optical flow calculations on FPGA and GPU architectures: a comparison study," *In Proc. of the 16th international symposium on field-programmable custom computing machines (FCCM '08)*, pp. 173-182, 2008.

[32] N.V. Chawla, K.W. Bowyer, L.O. Hall, W.P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research*, Vol. 16, pp. 321-357, 2002.

[33] L. Chen , X. Yang , Y. Liu , D. Zeng , Y. Tang , B. Yan , X. Lin , L. Liu, H. Xu, D. Zhou, "Quantitative and trajectory analysis of movement trajectories in mentary motor area seizures of frontal lobe epilepsy," *Epilepsy & Behavior*, Vol. 14(2), pp. 344-353, 2009.

[34] W.-Y. Chen, Y. Song, H. Bai, C.-J. Lin, E.Y. Chang, "Parallel Spectral Clustering in Distributed Systems," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 33(3), pp. 568-586, 2011.

[35] Y.Z. Cheng, Mean Shift, Mode Seeking, And Clustering, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 17, pp. 790-799, 1995.

[36] V. Cherkassky, F. Mulier, *Learning from Data: Concepts, Theory, and Methods*, John Wiley & Sons, 1998.

[37] Y. Cho, Y. Nam, Y.-J. Choi, W.-D. Cho, "Smart-Buckle: human activity recognition using a 3-axis accelerometer and a wearable camera," *In Proc. of the 2nd International Workshop on Systems and Networking Support for Health Care and Assisted Living Environments (HealthNet '08)*, 2008.

[38] T. Choudhury, S. Consolvo, B. Harrison, A. LaMarca, L. LeGrand, A. Rahimi, A. Rea, G. Borriello, B. Hemingway, P. Klasnja, K. Koscher, J. Landay, J. Lester, D. Wyatt, D. Haehnel, The mobile sensing platform: An embedded activity recognition system, *IEEE Pervasive Computing*, Vol. 7(2), pp. 32-41, 2008.

[39] J.G. Cleland, A.A. Louis, A.S. Rigby, U. Janssens, A.H. Balk, Non-Invasive Home Telemonitoring for Patients with Heart Failure at High Risk of Recurrent Admission and Death: The Trans-European Network Ű Home-Care Management Systems study (TEN-HMS), *Journal of the American College of Cardiology*, Vol. 45(10), pp. 1654-64, 2012.

[40] I. Conradsen, S. Beniczky, P. Wolf, D. Terney, T. Sams, H.B.D. Sorensen, "Multi-modal intelligent seizure acquisition (MISA) - A new approach towards seizure detection based on full body motion measures," *In Proc. of the 31st annual international conference of the IEEE EMBS*, pp. 2591-2595, 2009.

[41] I. Conradsen, S. Beniczky, P. Wolf, J. Henriksen, T. Sams, H.B.D. Sorensen, "Seizure onset detection based on a uni- or multi-modal intelligent seizure acquisition (UISA/MISA) system," *In Proc. of the 32nd Annual International*

Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2010), pp. 3269-72, 2010.

[42] I. Conradsen, S. Beniczky, K. Hoppe, P. Wolf, H.B.D. Sorensen, Automated algorithm for generalised tonic-clonic epileptic seizure onset detection based on sEMG zero-crossing rate, *IEEE Transactions on Biomedical Engineering*, Vol. 59(2), pp. 579-585, 2012.

[43] I. Conradsen, M. Moldovan, P. Jennum, P. Wolf, D. Farina, S. Beniczky, Dynamics of muscle activation during tonic-clonic seizures, *Epilepsy Research*, in press, 2012.

[44] P. Corbishley, E. Rodríguez-Villegas, Breathing detection: towards a miniaturized, wearable, battery-operated monitoring system, *IEEE Trans Biomed Eng*, Vol. 55(1), pp. 196-204, 2008.

[45] C. Cortes, V. Vapnik, "Support-vector networks," *Mach Learn*, Vol. 20(3), pp. 273-297, 1995.

[46] J.C. Crocker, D.G. Grier, Methods of digital video microscopy for colloidal studies, *Journal of Colloid and Interface Science*, Vol. 179, pp. 298-310, 1996.

[47] K. Cuppens, L. Lagae, B. Vanrumste, "Towards automatic detection of movement during sleep in pediatric patients with epilepsy by means of video recordings and the optical flow algorithm," *In Proc. (IFMBE Proceedings) of the 4th European congress of the international federation for medical and biological engineering (MBEC 2008)*, vol. 22, pp. 784-789, 2008.

[48] K. Cuppens, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, "Detection of nocturnal frontal lobe seizures in pediatric patients by means of accelerometers: a first study," *In Proc. of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2009)*, Vol. 1-20, pp. 6608-6611, 2009.

[49] K. Cuppens, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, "Automatic video detection of body movement during sleep based on optical flow in pediatric patients with epilepsy," *Medical & Biological Engineering & Computing*, Vol. 48(9), pp. 923-931, 2010.

[50] K. Cuppens, B. Bonroy, A. Van de Vel, B. Ceulemans, L. Lagae, T. Tuytelaars, S. Van Huffel, B. Vanrumste, "Automatic video detection of nocturnal epileptic movement based on motion tracks," *In Proc. of the International Conference on bio-inspired systems and signal processing (Biosignals)*, pp. 342-345, 2012.

[51] K. Cuppens, C.-W. Chen, K. Wong, A. Van de Vel, L. Lagae, B. Ceulemans, T. Tuytelaars, S. Van Huffel, B. Vanrumste, H. Aghajan, Using spatio-temporal interest points (STIP) for myoclonic jerk detection in nocturnal video, *In Proc. of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2012)*, 2012.

[52] K. Cuppens, C.-W. Chen, K. Wong, A. Van de Vel, L. Lagae, B. Ceulemans, T. Tuytelaars, S. Van Huffel, B. Vanrumste, H. Aghajan, "Integrating Video and Accelerometer Signals for Nocturnal Epileptic Seizure Detection," *In Proc. of the 14th ACM International Conference on Multimodal Interaction (ICMI)*, 2012.

[53] K. Cuppens, P. Karsmakers, A. Van de Vel, S. Luca, B. Bonroy, M. Milosevic, T. Croonenborghs, B. Ceulemans, L. Lagae, S. Van Huffel, B. Vanrumste, Handling unbalanced data in nocturnal epileptic seizure detection using accelerometers, *Internal report*, 2012.

[54] K. Cuppens, P. Karsmakers, A. Van de Vel, B. Bonroy, M. Milosevic, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, Accelerometer based home monitoring for detection of nocturnal hypermotor seizures based on novelty detection, *International Conference on Pattern Recognition Applications and Methods (ICPRAM)*, 2013.

[55] Q. Dai, B. Leng, "Video object segmentation based on accumulative frame difference," *In Proc. of the 26th Picture Coding Symposium (PSC 2007)*, 2007.

[56] A. Dalton, S. Patel, A.R. Chowdhury, M. Welsh, T. Pang, S. Schachter, G. Ólaighin, P. Bonato, "Detecting epileptic seizures using wearable sensor technologies," *In Proc. of the First AMA-IEEE Medical Technology Conference on Individualized Healtcare*, 2010.

[57] A. Dalton, S. Patel, A. Chowdhury, M. Welsh, T. Pang, S. Schachter, G. OLaighin, P. Bonato, Development of a body sensor network to detect motor patterns of epileptic seizures, *IEEE Transactions on Biomedical Engineering*, Vol. PP(99), 2012.

[58] W. Dargie, "Analysis of time and frequency domain features of accelerometer measurements," *In Proc. of the 18th International Conference on Computer Communications and Networks (ICCCN '09)*, pp. 1-6, 2009.

[59] G. Debard, P. Karsmakers, M. Deschodt, E. Vlaeyen, E. Dejaeger, K. Milisen, T. Goedeme, B. Vanrumste, T. Tuytelaars, Camera-Based Fall Detection

on Real World Data, *Outdoor and Large-Scale Real-World Scene Analysis, in book series: Lecture Notes in Computer Science*, Vol. 7474, pp. 356-375, 2012.

[60] C.M. DeGiorgio, S.C. Schachter, A. Handforth, M. Salinsky, J. Thompson, B. Uthman, R. Reed, S. Collins, E. Tecoma, G.L. Morris, B. Vaughn, D.K. Naritoku, T. Henry, D. Labar, R. Gilmartin, D. Labiner, I. Osorio, R. Ristanovic, J. Jones, J. Murphy, G. Ney, J. Wheless, P. Lewis, C. Heck, Prospective long-term study of vagus nerve stimulation for the treatment of refractory seizures, *Epilepsia*, Vol. 41, pp. 1195-1200, 2000.

[61] M.O. Derawi, "Accelerometer-Based Gait Analysis, A survey," *In Proc. of the NISK conference 2010*, 2010.

[62] J. De Tollenaere, Spectrale clustering in: Zelflerende Spraakherkenning, *Master thesis*, pp. 5-18, Faculty of Engineering, KU Leuven, Belgium, 2008.

[63] O. Devinsky, D. Gazzola, W.C.Jr. LaFrance, Differentiating between nonepileptic and epileptic seizures, *Nature Reviews Neurology*, Vol. 7, pp. 210-220, 2011.

[64] G. Fanelli, J. Gall, L. Van Gool, "Hough Transform-based Mouth Localization for Audio-Visual Speech Recognition," *In Proc. of the British Machine Vision Conference*, 2009.

[65] T.S. Ferguson, A bayesian analysis of some nonparametric problems, *Annals of statistics*, Volume 1, pp. 209-230, 1973.

[66] M. Figueiredo, A.K. Jain, Unsupervised learning of finite mixture models, *IEEE transactions on pattern analysis and machine intelligence*, Vol. 24(3), pp. 381-396, 2002.

[67] D.J. Fleet, A.D. Jepson, Computation of component image velocity from local phase information, *International Journal of Computer Vision*, Vol. 5, pp. 77-104, 1990.

[68] D.J. Fleet, Y. Weiss, Optical Flow Estimation in: Mathematical models for Computer Vision: The Handbook, Springer, 2005.

[69] C. Fowlkes, S. Belongie, F. Chung, J. Malik, Spectral Grouping Using the Nyström Method, *IEEE Trans. Pattern Anal. Mach. Intell.*, Vol. 26(2), pp. 214-225, 2004.

[70] M. Friedrich, Radar based motion detection module, *Datasheet Hygrosens Instruments GmbH*, 2008.

[71] D. Fudurić, D. Siladi, M. Žagar, Human Movement Detection Based on Acceleration Measurements and k-NN Classification, *In Proc. of the International Conference on Computer as a Tool (EUROCON 2007)*, pp. 589-594, 2007.

[72] C.-S. Fuh , P. Maragos, Region-based optical flow estimation, *In Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '89)*, pp. 130-135, 1989.

[73] K. Fukunaga, L.D. Hostetler, Estimation Of Gradient Of A Density-Function, With Applications In Pattern-Recognotion, *IEEE Transactions on Information Theory*, Vol. 21, pp. 32-40, 1975.

[74] S. Galic, S. Loncaric, Spatio-temporal image segmentation using optical flow and clustering algorithm, *In Proc. of the First International Workshop on Image and Signal Processing and Analysis (IWISPA 2000)*, pp. 63-68, 2000.

[75] S. Gauglitz, T. Höllerer, "In-Depth Evaluation of Popular Interest Point Detectors on Video Streams," *Technical Report: Department of Computer Science, University of California, Santa Barbara*, 2009.

[76] B. Greene, G. Boylan, R. Reilly, P. de Chazal, S. Connolly, Combination of EEG and ECG for improved automatic neonatal seizure detection, *Clin Neurophysiol*, Vol. 118(6), pp. 1348-59, 2007.

[77] L. Grote, D. Zou, H. Kraiczi, J. Hedner, Finger plethysmography - a method for monitoring finger blood flow during sleep disordered breathing, *Respiratory Physiology & Neurobiology*, Vol. 136(2-3), pp. 141-52, 2006.

[78] M. Gurban, J.-P. Thiran, T. Drugman, T. Dutoit, "Dynamic modality weighting for multi-stream HMMs inaudio-visual speech recognition," *In Proc. of the 10th international conference on Multimodal interfaces (ICMI '08)*, pp. 237-240, 2008.

[79] O. Gurné, V. Conraads, L. Missault, W. Mullens, J.L. Vachierys, W. Van Mieghem, W. Droogne, A.C. Pouleur, P. Troisfontaine, S. Huez, E. Nellessens, B. Peperstraete, P. Blouardo, J. Vanhaecke, D. Raes, A critical review on telemonitoring in heart failure, *Acta Cardiol.*, Vol. 67(4), pp. 439-44, 2012.

[80] N. Gyorbiro, A. Fabian, G. Homanyi, An activity recognition system for mobile phones, *Mobile Networks and Applications*, Vol. 14(1), 82-91, 2008.

[81] J. Han, B. Bhanu, "Human Activity Recognition in Thermal Infrared Imagery," *In Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR Workshops)*, 2005.

[82] C. Harris, M. Stephens, A combined corner and edge detector, *In Proc. of the 4th Alvey Vision Conference*, pp. 147-151, 1988.

[83] D.J. Heeger, Optical flow using spatiotemporal filters, *International Journal of Computer Vision*, Vol. 1, pp. 279-302, 1988.

[84] C. Hildebrand, *Introduction to Numerical Analysis: Second Edition*, Courier Dover Publications, 1987.

[85] L.J. Hirsch, E.J. Donner, E.L. So, M. Jacobs, L. Nashef, J.L. Noebels, J.R. Buchhalter, Abbreviated report of the NIH/NINDS workshop on sudden unexpected death in epilepsy, *Neurology*, Vol. 76(22), pp. 1932-8, 2011.

[86] B.K.P. Horn, B.G. Schunck, "Determining optical-flow," *Artif Intell*, Vol. 17(1-3), pp. 185-203, 1981.

[87] W.M. Hu, T.N. Tan, L. Wang, S. Maybank, "A survey on visual surveillance of object motion and behaviors," *IEEE Transactions on Systems Man and Cybernetics Part C-Applications and Reviews*, Vol. 34(3), pp. 334-352, 2004.

[88] M. Hu, S. Ali, M. Shah, Learning Motion Patterns in Crowded Scenes Using Motion Flow Field, *In Proc. International Conference on Pattern Recognition (ICPR) (2008)*, pp. 1-5, 2008.

[89] A.M. Husain, S.R. Sinha, Nocturnal epilepsy in adults, *J Clin Neurophysiol.*, Vol. 28(2), pp. 141-5, 2011.

[90] H. Inooka, Y. Ohtaki, H. Hayasaka, A. Suzuki, R. Nagatomi, "Development of advanced portable device for daily physical assessment," *In Proc. of the International Joint Conference SICE-ICASE*, pp. 5878-5881, 2010.

[91] A.K. Jain, M.N. Murty, P.J. Flynn, Data clustering: a review, *ACM Comput. Surv.*, Vol. 31(3), pp. 264-323, 1999.

[92] A.K. Jain, Data clustering: 50 years beyond K-means, *Pattern Recognition Letters*, Vol. 31(8), pp. 651-666, 2010.

[93] P. Jallon, S. Bonnet, M. Antonakios, R. Guillemaud, "Detection system of motor epileptic seizures through motion analysis with 3D accelerometers," *In Proc. of the 31st annual international conference of the IEEE EMBS*, pp. 2466-2469, 2009.

[94] P. Jallon, "A bayesian approach for epileptic seizures detection with 3D accelerometers sensors," *In Proc. of the 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2010)*, pp. 6325-8, 2010.

[95] X. Ji, H. Liu, Advances in View-Invariant Human Motion Analysis: A Review, *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, Vol. 40(1), pp. 13-24, 2010.

[96] N.B. Karayiannis, G.Z. Tao, J.D. Frost, M.S. Wise, R.A. Hrachovy, E.M. Mizrahi, "Automated detection of videotaped neonatal seizures based on motion segmentation methods," *Clinical Neurophysiology*, Vol. 117(7), pp. 1585-1594, 2006.

[97] N.B. Karayiannis, Y. Xiong, G. Tao, J.D. Jr. Frost, M.S. Wise, R.A. Hrachovy, E.M. Mizrahi, Automated detection of videotaped neonatal seizures of epileptic origin, *Epilepsia*, Vol. 47, pp. 966-980, 2006.

[98] L.K. Kearney, W.B. Thompson, D.L. Boley, Optical flow estimation: An error analysis of gradient-based methods with local optimization, *IEEE Tmns. Part. Anal. Artif. Mach. Intell.*, Vol. 9, pp. 229-244, 1987.

[99] J. Kearney, W. Thompson, "Bounding constraint propagation for optical flow estimation," *Motion understanding: robot and human vision (Aggarwal J, Martin W (eds))*, Kluwer Academic Publishers, 1988.

[100] Y.H. Kim, A.C. Kak, "Error analysis of robust optical flow estimation by least median of squares methods for the varying illumination model," *IEEE Trans Pattern Anal Mach Intell*, Vol. 28(9), pp. 1418-1435, 2006.

[101] T. Ko, "A survey on behavior analysis in video surveillance for homeland security applications," *In Proc. of the 37th IEEE Applied Imagery Pattern Recognition Workshop*, pp. 84-91, 2008.

[102] R. Kohavi, J.H. George, Wrappers for feature subset selection, *Artificial Intelligence*, Vol. 97(1-2), pp. 273-324, 1997.

[103] J. Kortelainen, J. Virkkala, "FFT averaging of multichannel BCG signals from bed mattress sensor to improve estimation of heart beat interval," *In Proc. of the IEEE Eng Med Biol Soc 1*, 2007.

[104] J.M. Kortelainen, M.O. Mendez, A.M. Bianchi, M. Matteucci, S. Cerutti, "Sleep staging based on signals acquired through bed sensor," *IEEE Transactions on Information Technology in Biomedicine*, Vol. 14(3), pp. 776-785, 2010.

[105] U. Kramer, S. Kipervasser, A. Shlitner, R. Kuzniecky, "A novel portable seizure detection alarm system: Preliminary results," *Journal of Clinical Neurophysiology*, Vol. 28(1), pp. 36-38, 2011.

[106] N. Krishnan, D. Colbry, C. Juillard, S. Panchanathan, "Real time human activity recognition using tri-Axial accelerometers," *In Proc. of the Sensors, Signals and Information Processing Workshop (SENSIP)*, 2008.

[107] N. Krishnan, S. Panchanathan, "Analysis of Low Resolution Accelerometer Data for Continuous Human Activity Recognition," *In Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2008)*, pp. 3337-3340, 2008.

[108] Y.-M. Kuo, J.-S. Lee, P.-C. Chung, A Visual Context-Awareness-Based Sleeping-Respiration Measurement System, *IEEE Transactions On Information Technology In Biomedicine*, Vol. 14(2), pp. 255-265, 2010.

[109] J.R. Kwapisz, G.M. Weiss, S.A. Moore, Activity recognition using cell phone accelerometers, *ACM SIGKDD Explorations Newsletter*, Vol. 12(2), pp. 74-82, 2010.

[110] M. Langer, R. Mann, Optical snow, *International Journal of Computer Vision*, Vol. 55(1), pp. 55-71, 2003.

[111] I. Laptev, "On space-time interest points," *International Journal of Computer Vision*, Vol. 64(2-3), pp. 107-123, 2005.

[112] I. Laptev, M. Marszalek, C. Schmid, B. Rozenfeld, "Learning realistic human actions from movies," *In Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.

[113] Y. Lee, S. Choi, Minimum entropy, k-means, spectral clustering, *In Proc. of the IEEE International Joint Conference on Neural Networks*, Vol. 1, 2004.

[114] J. Lester, T. Choudhury, G. Borriello, A practical approach to recognizing physical activities, *In Proc. of the 4th international conference on Pervasive Computing (PERVASIVE'06)*, pp. 1-16, 2006.

[115] Z.J. Li, A.M. da Silva, J.P.S. Cunha, "Movement quantification in epileptic seizures: A new approach to video-EEG analysis," *IEEE Transactions on Biomedical Engineering*, Vol. 49(6), pp. 565-573, 2002.

[116] C.-J. Lin, R.C. Weng, Simple probabilistic predictions for support vector regression, *Technical report, Department of Computer Science, National Taiwan University*, 2004.

[117] J.J. Little, A. Verri, "Analysis of differential and matching methods for optical flow," *In Proc. of the IEEE Workshop on Visual Motion*, pp. 173-180, 1989.

[118] Q. Liu, R.J. Sclabassi, M. Sun, "Change detection in epilepsy monitoring video based on markov random field theory," *In Proc. of the Intelligent signal processing and communication systems (ISPACS 2004)*, 2004.

[119] B. Lo, S. Velastin, "Automatic congestion detection system for underground platforms," *In Proc. of the Int. Symp. Intell. Multimedia, Video Speech Process.*, pp. 158-161, 2001.

[120] J. Lockman, R.S. Fisher, D.M. Olson, "Detection of seizurelike movements using a wrist accelerometer," *Epilepsy & Behavior*, Vol. 20(4), pp. 638-641, 2011.

[121] D.G. Lowe, Distinctive Image Features from Scale-Invariant Keypoints, *Int. J. Comput. Vision*, Volume 60(2), pp. 91-110, 2004.

[122] H. Lu, E. How-Lung, B. Mandal, D.W.S. Chan, N. Yen-Ling, "Markerless video analysis for movement quantification in pediatric epilepsy monitoring," *In Proc. of the 33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2011)*, 2011.

[123] B.D. Lucas, T. Kanade, "An iterative image registration technique with an application to stereo vision," *In Proc. of the DARPA imaging understanding workshop*, pp. 121-130, 1981.

[124] H.J. Luinge, P.H. Veltink, C.T.M. Baten, Ambulatory measurement of arm orientation, *J Biomech*, Vol. 40(1), pp. 78-85, 2007.

[125] D. Lulic, A. Ahmadian, A.A. Baaj, S.R. Benbadis, F.L. Vale, Vagus nerve stimulation, *Neurosurg Focus*, Vol. 27(3), 2009.

[126] U. Luxburg, A tutorial on spectral clustering, *Statistics and Computing*, Volume 17(4), pp. 395-416, 2007.

[127] M. Markou, S. Singh, "Novelty detection: a review - part 1: statistical approaches," *Signal Processing*, Vol. 83(12), pp. 2481-2497, 2003.

[128] F. Massé, J. Penders, A. Sereteyn, M. van Bussel, J. Arends, "Miniaturized wireless ECG-monitor for real-time detection of epileptic seizures," *In Proc. of the Wireless Health annual conference (WH '10)*, pp. 111-117, 2010.

[129] M. Mathie, B. Celler, N. Lovell, A. Coster, Classification of basic daily movements using a triaxial accelerometer, *Medical & Biological Engineering and Computing*, Vol. 42(5), pp. 679-687, 2004.

[130] U. Maurer, A. Smailagic, D. Siewiorek, M. Deisher, Activity recognition and monitoring using multiple sensors on different body positions, *In Proc. of the International Workshop on Wearable and Implantable Sensor Networks (BSN'06)*, pp. 113-116, 2006.

[131] G.J. McLachlan, Discriminant analysis and statistical pattern recognition, Wiley Series in Probability and Mathematical Statistics, 1992

[132] E. Miluzzo, N. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. Eisenman, X. Zheng, A. Campbell, "Sensing meets mobile social networks: The design, implementation and evaluation of the CenceMe application," *In Proc. of the 6th ACM Conference on Embedded Networked Sensor Systems (SenSys '08)*, pp. 337-350, 2008.

[133] J.H. Min, R. Kasturi, O. Camps, Extraction and temporal segmentation of multiple motion trajectories in human motion, *Image and Vision Computing*, Vol. 26, pp. 1621-1635, 2008.

[134] T.B. Moeslund, E. Granum, A Survey of Computer Vision-Based Human Motion Capture, *Computer Vision and Image Understanding*, Vol. 81, pp. 231-268, 2001.

[135] T.M.E. Nijsen, P.J.M. Cluitmans, J. Arends, P.A.M. Griep, "Detection of subtle nocturnal motor activity from 3-D accelerometry recordings in epilepsy patients," *IEEE Transactions on Biomedical Engineering*, Vol. 54(11), pp. 2073-2081, 2007.

[136] T.M.E Nijsen, R.M. Aarts, P.J.M. Cluitmans, P.A.M. Griep, "Timefrequency analysis of accelerometry data for detection of myoclonic seizures," *IEEE Transactions on Information Technology in Biomedicine*, Vol. 14(5), pp. 1197-1203, 2010.

[137] M. Pediaditis, M. Tsiknakis, P. Vorgia, D. Kafetzopoulos, V. Danilatou, D. Fotiadis, "Vision-Based Human Motion Analysis in Epilepsy - Methods and Challenges," *In Proc. of the 10th IEEE International Conference on Information Technology and Applications in Biomedicine (ITAB)*, pp. 1-5, 2010.

[138] D. Pelleg, A. Moore, X-means: Extending k-means with efficient estimation of the number of clusters, *In Proc. of the 17th international conference on machine learning*, pp. 727-734, 2000.

[139] K. Petersen, "Biomedical Applications of MEMS," *In Proc. of the International Electron Devices Meeting (IEDM '96)*, pp. 239-242, 1996.

[140] C.F. Poets, Apparent life-threatening events and sudden infant death on a monitor, *Paediatr Respir Rev*, Vol. 5 (Suppl. A), pp. 383-6, 2004.

[141] M.-Z. Poh, T. Loddenkemper, N.C. Swenson, S. Goyal, J.R. Madsen, R.W. Picard, Continuous monitoring of electrodermal activity during epileptic seizures using a wearable sensor, *In Proc. of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS 2010)*, pp. 4415-8, 2010.

[142] M.-Z. Poh, D.J. McDuff, R.W. Picard, Advancements in Noncontact, Multiparameter Physiological Measurements Using a Webcam, *IEEE Transactions On Biomedical Engineering*, Vol. 58(1), pp. 7-11, 2011.

[143] D.L. Polla, "MEMS Technology for Biomedical Applications," *In Proc. of the 6th International Conference on Solid-State and Integrated-Circuit Technology*, Vol. 1, pp. 19-22, 2001.

[144] R. Poppe, "A survey on vision-based human action recognition," *Image and Vision Computing*, Vol. 28(6), pp. 976-990, 2010.

[145] T. Pylvänäinen, "Accelerometer Based Gesture Recognition Using Continuous HMMs," *Pattern Recognition and Image Analysis, in book series: Lecture Notes in Computer Science*, pp. 413-430, 2005.

[146] P. Rajendran, A. Corcoran, B. Kinosian, M. Alwan, Falls, Fall Prevention, and Fall Detection Technologies, *Eldercare Technology for Clinical Practitioners, in book series: Aging Medicine*, pp. 187-202, 2008.

[147] R.M. Rangayyan, Y. Wu, "Modeling and classification of knee-joint vibroarthrographic signals using probability density functions estimated with Parzen windows," *In Proc. of the 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS 2008)*, pp. 2099-2102, 2008.

[148] B. Raskutti, C. Leckie, An Evaluation of Criteria for Measuring the Quality of Clusters, *In Proc. of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI '99)*, pp. 905-910, 1999.

[149] C. Rasmussen, The infinite Gaussian mixture model, *Advances in neural information processing systems (NIPS)*, Volume 12, pp. 554-560, 2000.

[150] Y. Rathi, J. Malcolm, S. Bouix, R. McCarley, L. Seidman, J. Goldstein, C.F. Westin, M. Shenton, "Disease classification: A probabilistic approach," *In Proc. of the IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 1345-1348, 2010.

[151] E. Rauhala, J. Virkkala, S.L. Himanen, "Periodic limb movement screening as an additional feature of Emfit sensor in sleep-disordered breathing studies," *Journal of Neuroscience Methods*, Vol. 178(1), pp. 157-161, 2009.

[152] N. Ravi, N. Dandekar, "Activity recognition from accelerometer data," *In Proc. of the 17th Conference on Innovative Applications of Artificial Intelligence (IAAI)*, pp. 1541-1546, 2005.

[153] J. Rémi, J.P. Silva Cunha, C. Vollmar, Ö.B. Topçuŏglu, A. Meier, S. Ulowetz, P. Beleza, S. Noachtar, Quantitative movement analysis differentiates focal seizures characterized by automatisms, *Epilepsy Behav.*, Vol. 20(4), pp. 642-7, 2011.

[154] G.B. Richerson, G.F. Buchanan, The serotonin axis: Shared mechanisms in seizures, depression, and SUDEP, *Epilepsia*, Vol. 52 (Suppl. 1), pp. 28-38, 2011.

[155] T. Rowekamp, L. Peters, "A real-time smart sensor system for visual motion estimation," *In Proc. of the 1997 European conference on design and test (EDTC '97)*, pp. 613, 1997.

[156] D. Ruta, B. Gabrys, Automated detection of videotaped neonatal seizures based on motion segmentation methods, *Computing and Information Systems*, Vol. (7), pp. 1-10, 2000.

[157] P. Sand, S. Teller, "Particle video: Longrange motion estimation using point trajectories," *In Proc. of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2195-2202, 2006.

[158] H.M. Schepers, D. Roetenberg, P.H. Veltink, "Ambulatory human motion tracking by fusion of inertial and magnetic sensing with adaptive actuation," *Med Biol Eng Comput*, Vol. 48(1), pp. 27-37, 2010.

[159] C. Schuldt, I. Laptev, B. Caputo, "Recognizing human actions: A local SVM approach," *In Proc. of the 17th International Conference on Pattern Recognition (ICPR)*, pp. 32-36, 2004.

[160] J. Shi, C. Tomasi, Good Features to Track, *In Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '94)*, pp. 593-600, 1994.

[161] E.L. So, J. Bainbridge, J.R. Buchhalter, J. Donalty, E.J. Donner, A. Finucane, N.M. Graves, L.J. Hirsch, G.D. Montouris, N.R. Temkin, S. Wiebe, T.L. Sierzant, "Report of the American Epilepsy Society and the Epilepsy Foundation Joint Task Force on Suddent Unexplained Death in Epilepsy," *Epilepsia*, Vol. 50(4), pp. 917-22, 2009.

[162] M. Sonka, V. Hlavac, R. Boyle, "Motion analysis," *Image processing, analysis, and machine vision*, Kluwer Academic Publishers, 2008.

[163] C. Stauffer, W. Grimson, "Adaptive background mixture models for real-time tracking," *In Proc. of the IEEE Conf. Comput. Vis. Pattern Recognit.*, Vol. 2, pp. 246-252, 1999.

[164] A. Steventon, M. Bardsley, J. Billings, J. Dixon, H. Doll, S. Hirani, M. Cartwright, L. Rixon, M. Knapp, C. Henderson, A. Rogers, R. Fitzpatrick, J. Hendy, S. Newman, Effect of telehealth on use of secondary care and mortality: findings from the Whole System Demonstrator cluster randomised trial, *British Medical Journal*, Vol. 344, 2012.

[165] S. Suzuki, T. Matsui, H. Kawahara, H. Ichiki, J. Shimizu, Y. Kondo, S. Gotoh, H. Yura, B. Takase, M. Ishihara, "A non-contact vital sign monitoring system for ambulances using dual-frequency microwave radars," *Med Biol Eng Comput*, Vol. 47(1), pp. 101-105, 2009.

[166] E.M. Tapia, S.S. Intille, W. Haskell, K. Larson, J. Wright, A. King, R. Friedman, "Real-Time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor," *In Proc. of the 11th IEEE International Symposium on Wearable Computers*, pp. 37-40, 2007.

[167] L. Tarassenko, P. Hayton, N. Cerneaz, M. Brady, "Novelty detection for the identification of masses in mammograms," *In Proc. of the Fourth International Conference on Artificial Neural Networks*, pp. 442-447, 1995.

[168] L. Tarassenko, A. Hann, D. Young, "Integrated monitoring and analysis for early warning of patient deterioration," *British Journal of Anaesthesia*, Vol. 97(1), pp. 64-68, 2006.

[169] A. Temko, E. Thomas, W. Marnane, G. Lightbody, G. Boylan, "EEG-based neonatal seizure detection with support vector machines," *Clinical Neurophysiology*, Vol. 122(3), pp. 464-473, 2011.

[170] R. Tibshirani, G. Walther, T. Hastie, Estimating the number of clusters in a data set via the gap statistic, *Journal of the royal statistical society: Series B*, Volume 63(2), pp. 411-423, 2001.

[171] P. Tinuper, F. Provini, F. Bisulli, E. Lugaresi, Hyperkinetic manifestations in nocturnal frontal lobe epilepsy. Semeiological features and physiopathological hypothesis, *Neurol Sci.*, Vol. 26 (Suppl. 3), pp. 210-4, 2005.

[172] P. Tinuper, F. Provini, F. Bisulli, L. Vignatelli, G. Plazzi, R. Vetrugno, P. Montagna, E. Lugaresi, Movement disorders in sleep: guidelines for differentiating epileptic from non-epileptic motor phenomena arising from sleep, *Sleep Med Rev.*, Vol. 11(4), pp. 255-67, 2007.

[173] C. Tomasi, T. Kanade, "Detection and Tracking of Point Features," *Technical Report: Carnegie Mellon University (CMU-CS-91-132)*, 1991.

[174] J. Tompkin, "Optical Flow An Introduction," *Internal report: University College London*, 2008.

[175] H. Top, Optical flow en bewegingillusies, *Master thesis*, Faculty of Mathematics & Natural Sciences, University of Groningen,The Netherlands, 2007.

[176] P. Turaga, R. Chellappa, V.S. Subrahmanian, O. Udrea, "Machine recognition of human activities: a survey," *IEEE Trans Circuits Syst Video Technol*, Vol. 18(11), pp. 1473-1488, 2008.

[177] R.H. Turi, S. Ray, Determination of Number of Clusters in K-Means Clustering and Application in Colour Image Segmentation, *In Proc. of the 4th International Conference on Advances in Pattern Recognition and Digital Techniques (ICAPRDT'99)*, pp. 137-143, 1999.

[178] S. Ullman, *The interpretation of visual motion*, MIT Press, 1979.

[179] S. Uras, E. Girosi, A. Verri, V. Torre, A computational approach to motion perception, *Biol. Cybern.*, Vol. 60, pp. 79-97, 1988.

[180] M. van Bussel, J. Arends, F. Massé, A. Serteyn, I. Tan, J. Penders, P. Griep, Detection of major epileptic seizures with heart rate changes: feasibility test of a state-of-the-art wearable sensor, *Epilepsia*, Vol. 51 (Suppl. 4), pp. 56, 2010.

[181] A. Van de Vel, K. Cuppens, B. Bonroy, M. Milosevic, L. Lagae, S. Van Huffel, B. Vanrumste, B. Ceulemans, Epilepsy alert! From seizure detection to SUDEP prevention, *Internal report*, 2012.

[182] W. Van Looy, K. Cuppens, B. Bonroy, A. Van de Vel, B. Ceulemans, L. Lagae, S. Van Huffel, B. Vanrumste, "Detection of body movement using optical flow and clustering," *In Proc. of the 2nd International Conference on Positioning and Context-Awarenss (PoCA)*, 2011.

[183] V.N. Vapnik, Statistical learning theory, *Wiley*, New York, 1998.

[184] P.K. Varshney, Multisensor data fusion, *Electronics and Communications Engineering Journal*, Vol. 9, pp. 245-253, 1997.

[185] P.H. Veltink, H.B.J. Bussman, F. Koelma, H.M. Franken, W.L.J. Martens, R.C. van Lummel, "The feasibility of posture and movement detection by accelerometry," *In Proc. of the 15th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, 1993.

[186] P.H. Veltink, H.B.J. Bussmann, W. de Vries, W.L.J. Martens, R.C. van Lummel, Detection of static and dynamic activities using uniaxial accelerometers, *IEEE Transactions on Rehabilitation Engineering*, Vol. 4(4), pp. 375-385, 1996.

[187] P. Viola, M.J. Jones, D. Snow, Detecting Pedestrians Using Patterns of Motion and Appearance, *Int. J. Comput. Vision*, Volume 63(2), pp. 153-161, 2005.

[188] U. Von Luxburg, A Tutorial on Spectral Clustering, *Statistics and Computing*, Vol. 17(4), pp. 395-416, 2007.

[189] L.A. Wang, W.M. Hu, T.N. Tan, "Recent developments in human motion analysis," *Pattern Recognit.*, Vol. 36(3), pp. 585-601, 2003.

[190] C.W. Wang, A. Ahmed, A. Hunter, "Vision Analysis in Detecting Abnormal Breathing Activity in application to Diagnosis of Obstructive Sleep Apnoea," *In Proc. of the 28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS 2006)*, pp. 4469-4473, 2006.

[191] A. Weinstock, P. Giglio, S.L. Kerr, P.K. Duffner, M.E. Cohen, Hyperkinetic Seizures in Children, *J Child Neurol*, Vol. 18(8), pp. 517-524, 2003.

[192] D. Wolpert, Stacked generalization in: Neural Networks, pp. 241-259, 1992.

[193] T.-F. Wu, C.-J. Lin, R.C. Weng, Probability estimates for multi-class classification by pairwise coupling, *Journal of Machine Learning Research*, Vol. 5, pp. 975-1005, 2004.

[194] L. Xu, J. Jia, Y. Matsushita, Motion Detail Preserving Optical Flow Estimation, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 34(9), pp. 1744-1757, 2012.

[195] D. Yan, L. Huang, M.I. Jordan, "Fast approximate spectral clustering," *In Proc. of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 907-916, 2009.

[196] J. Yang, "Toward physical activity diary: Motion recognition using simple acceleration features with mobile phones," *In Proc. of the 1st International Workshop on Interactive Multimedia for Consumer Electronics (IMCE '09)*, pp. 1-10, 2009.

[197] D.Y. Yeung, C. Chow, "Parzen-window network intrusion detectors," *In Proc. of the 16th International Conference on Pattern Recognition*, Vol. 4, pp. 385-388, 2002.

[198] A. Yilmaz, O. Javed, M. Shah, Object tracking: A survey, *ACM Computing Surveys*, Vol. 38, 2006.

Biography

Kris Cuppens was born on May 21, 1984. He received a M.Sc. degree ("industrieel ingenieur") in electronics-ICT from the Katholieke Hogeschool Limburg in 2007. After his studies he started to work at the Mobilab research group of the Katholieke Hogeschool Kempen on a research project about the detection of epileptic seizures based on video recordings. His Ph.D. research is carried out at the Mobilab research group and the KU Leuven in the faculty of Applied Sciences, department of Electrical Engineering in the SCD/SISTA research division, on the same research topic.

Publications by the author

Papers in international journals

1. **K. Cuppens**, L. Lagae, B. Ceulemans, S. Van Huffel and B. Vanrumste. Automatic video detection of body movement during sleep based on optical flow in pediatric patients with epilepsy. *Medical and Biological Engineering and Computing*, Vol. 48(9), pp. 923-931, 2010.

2. A. Van de Vel, **K. Cuppens**, B. Bonroy, M. Milosevic, S. Van Huffel, B. Vanrumste, L. Lagae, B. Ceulemans, Long term home monitoring of hypermotor seizures by patient-worn accelerometers, *Epilepsy & Behavior, accepted for publication*, 2012.

Papers submitted to international journals

3. **K. Cuppens**, P. Karsmakers, A. Van de Vel, S. Luca, B. Bonroy, M. Milosevic, T. Croonenborghs, B. Ceulemans, L. Lagae, S. Van Huffel, B. Vanrumste, Accelerometer based home monitoring for detection of nocturnal hypermotor seizures based on novelty detection, *IEEE Transactions on Information Technology in Biomedicine, submitted*, 2012.

Papers in proceedings of international conferences

4. **K. Cuppens**, P. Karsmakers, A. Van de Vel, B. Bonroy, M. Milosevic, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, Handling unbalanced data in nocturnal epileptic seizure detection using accelerometers. *In Proc.*

of the 2nd International Conference on Pattern Recognition Applications and Methods (ICPRAM), Barcelona, Spain, accepted for publication, Feb. 2013.

5. **K. Cuppens**, C.-W. Chen, K. Wong, A. Van de Vel, L. Lagae, B. Ceulemans, T. Tuytelaars, S. Van Huffel, B. Vanrumste, H. Aghajan, Integrating Video and Accelerometer Signals for Nocturnal Epileptic Seizure Detection. *In Proc. of the 14th ACM International Conference on Multimodal Interaction (ICMI), Santa Monica, CA, USA*, Oct. 2012.

6. **K. Cuppens**, C.-W. Chen, K. B.-Y. Wong, A. Van de Vel, L. Lagae, B. Ceulemans, T. Tuytelaars, S. Van Huffel, B. Vanrumste, H. Aghajan, Using Spatio-Temporal Interest Points (STIP) for myoclonic jerk detection in nocturnal video. *In Proc. of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE EMBC), San Diego, CA, USA*, Aug. 2012.

7. **K. Cuppens**, B. Bonroy, A. Van de Vel, B. Ceulemans, L. Lagae, T. Tuytelaars, S. Van Huffel, B. Vanrumste. Automatic video detection of nocturnal epileptic movement based on motion tracks. *In Proc. of the International Conference on bio-inspired systems and signal processing (Biosignals), Algarve, Portugal*, pp. 342-345, Feb. 2012.

8. W. Van Looy, **K. Cuppens**, B. Bonroy, A. Van de Vel, B. Ceulemans, L. Lagae, S. Van Huffel, B. Vanrumste, Detection of body movement using optical flow and clustering. *In Proc. of the 2nd International Conference on Positioning and Context-Awarenss (PoCA), Brussels, Belgium* Mar. 2011.

9. A. Decaigny, **K. Cuppens**, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, Accelerometers used for the detection of nocturnal frontal lobe seizures in pediatric patients. *In Proc. of the European Conference on the Use of Modern Information and Communications Technologies (ECUMICT), Ghent, Belgium*, pp. 331-342, Mar. 2010.

10. **K. Cuppens**, B. Vanrumste, B. Ceulemans, L. Lagae, S. Van Huffel, Detection of epileptic seizures using video data *In Proc. of the 6th International Conference on intelligent environments (IE10), Kuala Lumpur, Malaysia*, Jul. 2010.

11. **K. Cuppens**, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, Detection of nocturnal frontal lobe seizures in pediatric patients by means of accelerometers: A first study. *In Proc. of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE EMBC), Minneapolis, MN, USA*, pp. 6608-11, Sep. 2009.

12. **K. Cuppens**, L. Lagae, B. Vanrumste, Towards automatic detection of movement during sleep in pediatric patients with epilepsy by means of video recordings and the optical flow algorithm. *In Proc. of the 4th European Conference of the International Federation for Medical and Biological Engineering (IFBME MBEC), Antwerp, Belgium*, pp. 784-789, Nov. 2008.

13. **K. Cuppens**, S. Omloop, P. Colleman, B. Ceulemans, B. Vanrumste, Detection of Nocturnal Epileptic Seizures of Pediatric Patients Using Accelerometers: Preliminary Results. *In Proc. of the IEEE/EMBS Benelux chapter, Heeze, The Netherlands*, Dec. 2007.

14. D. Tormans, T. Swinnen, **K. Cuppens**, S. Omloop, B. Geraets, P. Colleman, V. Claes, B. Ceulemans, B. Vanrumste, Nocturnal monitoring of pediatric patients with epilepsy based on accelerometers. *In Proc. of the Program for Research on Integrated Systems and Circuits (ProRISC) conference, Veldhoven, The Netherlands*, Nov. 2007.

Abstracts in proceedings of international conferences

15. **K. Cuppens**, A. Van de Vel, P. Karsmakers, B. Bonroy, M. Milosevic, S. Luca, T. Croonenborghs, T. Tuytelaars, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, Detection Of Nocturnal Convulsive Seizures In Children By Integrating Video And Accelerometer Recordings. *in Proc. of the 4th Dutch Bio-Medical Engineering Conference, Egmond aan Zee, The Netherlands, accepted for publication*, Jan. 2013.

16. I. Van den Herrewegen, **K. Cuppens**, M. Broeckx, H. Vertommen, M. Mertens, L. Peeraer, Development of a model to analyse foot biomechanics using dynamic 3D surface scanning. *in Proc. of the 37th Congress of the Société de Biomécanique, Toulouse, France)*, Oct. 2012. *Computer Methods in Biomechanics and Biomedical Engineering* Vol. 15 (Suppl. 1), pp. 85-86, 2012.

17. I. Van den Herrewegen, **K. Cuppens**, M. Broeckx, H. Vertommen, M. Mertens, L. Peeraer, Towards analysis of foot motion using dynamic 3D surface scanning. *in Proc. of the 18th Congress of the European Society of Biomechanics, Lisbon, Portugal*, Jul. 2012. *Journal of Biomechanics* Vol. 45 (Suppl. 1), pp. S197, 2012.

18. B. Bonroy, R. Gransier, P. Dunias, **K. Cuppens**, K. Meijer, B. Vanrumste, Inductive Sensor Technology For Ambulantory Measurement Of Physical Activity. *in Proc. of the IEEE/EMBS Benelux chapter, Leuven/Brussels, Belgium*, Dec. 2011.

19. **K. Cuppens**, A. Van de Vel, B. Bonroy, B. Ceulemans, L. Lagae, T. Tuytelaars, S. Van Huffel, B. Vanrumste, Detection of myoclonic shocks in nocturnal video by means of motion tracks. *in Proc. of the IEEE/EMBS Benelux chapter, Leuven/Brussels, Belgium*, Dec. 2011.

20. A. Van de Vel, **K. Cuppens**, B. Bonroy, M. Milosevic, R. Krols, L. Gijsemans, J. Vervisch, L. Lagae, S. Van Huffel, B. Vanrumste, B. Ceulemans, Detection systems and algorithms for motor seizures in pediatric patients with epilepsy. *in Proc. of the 29th International Epilepsy Congress, Rome, Italy*, Sep. 2011.

21. A. Van de Vel, **K. Cuppens**, B. Bonroy, M. Milosevic, R. Krols, L. Gijsemans, J. Vervisch, L. Lagae, S. Van Huffel, B. Vanrumste, B. Ceulemans, Accelerometers for noctural motor seizure detection in pediatric patients with epilepsy. *FlandersBio's annual life sciences convention (Knowledge for growth), Ghent, Belgium*, May 2011.

22. A. Van de Vel, **K. Cuppens**, B. Bonroy, M. Milosevic, R. Krols, L. Gijsemans, J. Vervisch, L. Lagae, S. Van Huffel, B. Vanrumste, B. Ceulemans, Accelerometers for noctural motor seizure detection in pediatric patients with epilepsy. *in Proc. of the 9th European Paediatric neurology society congress, Dubrovnik, Croatia*, May 2011.

23. **K. Cuppens**, A. Van de Vel, B. Bonroy, M. Milosevic, R. Krols, L. Gijsemans, J. Vervisch, T. Tuytelaars, B. Ceulemans, L. Lagae, S. Van Huffel, B. Vanrumste, Extraction of features for myoclonic shock detection in video based on mean shift clustering for constructing motion tracks. *in Proc. of the 9th European Paediatric neurology society congress, Cavtat, Dubrovnik, Croatia*, May 2011.

24. **K. Cuppens**, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, Optical flow based motion tracks for the detection of nocturnal myoclonic shocks in pediatric patients. *in Proc. of the 3rd Dutch Bio-Medical Engineering Conference (BME2011), Egmond aan Zee, The Netherlands*, Jan. 2011.

25. **K. Cuppens**, L. Lagae, B. Ceulemans, S. Van Huffel, B. Vanrumste, Towards automatic detection of myoclonic shocks in pediatric patients by means of video recording. *in Proc. of the 9th International Conference and Workshop on Ambient Intelligence and Embedded Systems, Geel, Belgium*, Sep. 2010.

26. B. Ceulemans, **K. Cuppens**, L. Lagae, S. Van Huffel, B. Vanrumste, Detection of nocturnal frontal lobe seizures in pediatric patients by means of accelerometers: preliminary results. *in Proc. of the 8th European Paediatric Neurology Society congress (EPNS), Harrogate, UK*, pp. S67, Sep. 2009. *European Journal of Paediatric Neurology* Vol. 13 (Suppl. 1), Page S67, 2009.

27. **K. Cuppens**, L. Lagae, B. Vanrumste, Towards automatic detection of movement during sleep in pediatric patients with epilepsy by means of video recordings and the optical flow algorithm. *in Proc. of the International Summit of Life Sciences Biomedica (Biomedica), Liège, Belgium*, pp. D10, Apr. 2009.

28. **K. Cuppens**, L. Lagae, B. Vanrumste, Towards automatic detection of movement during sleep in pediatric patients with epilepsy by means of video recordings and the optical flow algorithm. *in Proc. of the 7th International Conference and Workshop on Ambient Intelligence and Embedded Systems (AmiEs), Kiel, Germany*, Sep. 2008.

29. **K. Cuppens**, S. Omloop, P. Colleman, B. Ceulemans, B. Vanrumste, Detection of Nocturnal Epileptic Seizures of Pediatric Patients Using Accelerometers. *in Proc. of the Belgian Day on Biomedical Engineering, Brussels, Belgium*, Nov. 2007.

Arenberg Doctoral School of Science, Engineering & Technology
Faculteit ingenieurswetenschappen
Department elektrotechniek
Onderzoeksgroep SCD/SISTA
Kasteelpark Arenberg 10, B-3001 Leuven (Belgium)



