TÉCNICO LISBOA

LAMPSY - A Novel Epilepsy Video Monitoring and Seizure Detection Device

Vicente Mayer Mendonça de Lima Garção

Thesis to obtain the Master of Science Degree in

Biomedical Engineering

Supervisor(s): Prof. Hugo Humberto Plácido da Silva
Prof. Ana Luísa Nobre Fred

Examination Committee

Chairperson: Prof. João Miguel Raposo Sanches
Supervisor: Prof. Ana Luísa Nobre Fred
Member of the Committee: Prof. Paulo Luís Serras Lobato Correia

December 2021

Declaration

I declare that this document is an original work of my own authorship and that it fulfills all the requirements of the Code of Conduct and Good Practices of the Universidade de Lisboa.

Preface

This work was conducted at Instituto de Telecomunicações (IT) (Lisbon, Portugal) and at the Departamento de Bioengenharia at Instituto Superior Técnico (IST) (Lisbon, Portugal), in collaboration with the epilepsy management units of Hospital de Santa Maria (HSM) (Lisbon, Portugal) and Hospital de Egas Moniz (HEM) (Lisbon, Portugal), under the supervision of Professors Hugo Plácido da Silva and Ana Luísa Nobre Fred of IT and IST, and the co-supervision of Dr. Carla Bentes of HSM and Mariana Abreu. This work was part of a collaborative project between IT, HSM and HEM called ”PreEpiSeizures” concerned with the development of devices to address the challenges associated with epilepsy. This work was partially funded by a research grant ”Research Grant - Image Processing for Epileptic Seizure Analysis” opened at IT in the scope of the project TeamBIT – D-0002-LX-19, funded by by the applicable financial framework (FCT/MEC through national funds and when applicable co-funded by FEDER – PT2020 partnership agreement).

Acknowledgments

Firstly, I would like to express my gratitude to Professors Hugo Plácido da Silva and Ana Luísa Nobre Fred for giving me the opportunity to work on multiple projects over the years that have been endlessly challenging and gratifying, consistently providing positive and motivating feedback, and always offering great suggestions and ideas, without which this work would not have been possible. Moreover, I would like to thank Dr. Carla Bentes for her co-supervision, the medical teams of the epilepsy monitoring units at Hospital de Santa Maria and Hospital de Egas Moniz for their availability and expertise, and Mariana Abreu for her support, advice, and work with the Ethics Committee which assured that it would be possible to place the prototype in Hospital de Santa Maria. I would also wish to extend my appreciation to Hospital de Egas Moniz and Hospital de Santa Maria for providing me with important video data; video-EEG technicians Rosa Santos of HSM and Octávia Brás of HEM, as well as Dr. Francisca Sá of HEM for their assistance in annotating seizure videos. Additionally, I would like to thank Nuno Silva of IT for helping me set up the prototype's SFTP connection to a server at IT.

A word of appreciation also goes to Instituto de Telecomunicações (IT) for the support provided to this work.

I would like to extend my appreciation to my colleagues Ana Filipa Baltazar, António Azeitona, Guilherme Pata and Rita Silva, who began this journey with me in the Project In Biomedical Engineering course, and helped define many of the ideas behind this project at its inception.

I would also like to mention my friends from "*Trompete*" and "*The Order*", some of whom I met over 13 years ago at DSL, whose fellowship and camaraderie is unrivalled by anyone I know. Moreover, I would like to thank "*Tecas*" – the friends that accompanied me through my academic journey over the last five years, because of whom, from silly moments at Ilha de Faro to many trips to Évora and Portalegre, I had by far the best years of my life. Though we may no longer have classes together, the bond we forged will be hard to break.

Furthermore, I would like to express my deepest gratitude to Rita Silva, who over the last 4 years has made me a better, more hard-working, and more caring person, and whose kind-heartedness knows no bounds.

Lastly, I also wish to thank my family - my brother and sister, without whom I would most likely be perpetually bored, and my mother and father, who have always loved me, supported me, given me the best possible education, and done everything in their power to help me succeed. Without their efforts, none of this would have been possible.

Resumo

A epilepsia é uma patologia neurológica caracterizada por crises epiléticas que afeta 50 milhões de pessoas mundialmente, 30% das quais sofrem de epilepsia fármaco-resistente. Estes pacientes têm crises frequentes que podem diminuir a sua qualidade de vida e levar a uma morte prematura. A deteção automática de crises pode permitir um melhor ajuste da medicação ao simplificar o processo de diagnóstico e melhorar a sua precisão. Pode também contribuir para a diminuição da ansiedade dos doentes e cuidadores e permitir a prestação imediata de cuidados de saúde em crises graves.

O principal objetivo deste trabalho foi criar um dispositivo de deteção de crises por vídeo que fosse preciso, acessível e discreto e preservasse a privacidade dos pacientes de raíz. Este dispositivo foi apelidado de "**Lampsy**" e um pedido de patente foi efetuado.

Foi construído um protótipo com base num Raspberry Pi integrado num candeeiro, uma câmara e um ecrã, que foi colocado no Hospital de Santa Maria para aquisições durante sessões de vídeo-EEG. Os dados foram transmitidos para um servidor seguro por SFTP, preservando a privacidade dos pacientes.

Foi desenvolvido um algoritmo de detecção de crises baseado no fluxo óptico, isolamento de movimentos epilépticos através de PCA, ICA e novas métricas, e classificação com aprendizagem automática. Foi atingido um alto nível de precisão, com uma AUC de 96,5%, sensibilidade e especificidade de 94,8%, e latência média de 52,9 segundos. As crises foram segmentadas com um erro de $\pm$ 8,85 segundos, quando comparadas com anotações feitas por profissionais de saúde.

Palavras-chave: Epilepsia, Deteção de Crises, Visão Computacional, Processamento de Sinais, Aprendizagem Automática

Abstract

Epilepsy is a neurological condition characterized by the occurrence of unprovoked seizures. It affects 50 million people worldwide, 30% of which have pharmacoresistant epilepsy, leading to frequent seizures that may considerably worsen their quality of life and potentially lead to premature death. Seizure detection and registration could be essential in facilitating the diagnosis and management of this condition and improving diagnostic accuracy, thereby enabling better medication adjustment. It can also reduce stress levels in patients and caregivers and ensure that timely help can be delivered when dangerous seizures happen. The main objective of this work was to create an accurate, affordable, unobtrusive, and accessible video-based seizure detection device that preserved patient privacy by design. This device was named "**Lampsy**", and a patent request for it was filed.

A prototype was built using a Raspberry Pi single-board computer integrated within a light fixture, a small camera, and a display, and placed in Hospital de Santa Maria for video acquisition during video-EEG sessions. The video data was transmitted to a secure server via SFTP, preserving patient privacy. A seizure detection algorithm based on optical flow, PCA and ICA-based seizure movement isolation, novel seizure metrics, and machine learning classification was developed. High detection accuracy was achieved, with an AUC of 96.5%, sensitivity, and specificity of 94.8% at the equal error rate, and an average latency of 52.9 seconds. Seizures were segmented with an error of $\pm$ 8.85 seconds when compared to annotations by health professionals.

Keywords: Epilepsy, Seizure Detection, Computer Vision, Signal Processing, Machine Learning

Contents

List of Tables

List of Figures

List of Acronyms

ACM Accelerometry

AUC Area Under Curve

BSS Blind Source Separation

COM Center Of Mass

CV Cross-Validation

ECG Electrocardiography

EDA Electrodermal Activity

EEG Electroencephalography

FAR False Alarm Rate

FFT Fast Fourier Transform

GPC Gaussian Process Classifier/Classification

GPIO General Purpose Input-Output

GTCS Generalized Tonic-Clonic Seizures

HEM Hospital de Egas Moniz

HSM Hospital de Santa Maria

ICA Independent Component Analysis

IC Independent Component

ILAE International League Against Epilepsy

IT Instituto de Telecomunicações

L/H% Percentage of Low and High Amplitude Samples

LED Light Emitting Diode

LOSOCV Leave-One-Subject-Out Cross-Validation

MLP Multilayer Perceptron

ML Machine Learning

OF Optical Flow

P2RMS Peak-to-RMS Ratio

P2SSR Peak-to-SSR Ratio

PCA Principal Component Analysis

PC Principal Component

PPG Photoplethysmography

PPV Positive Predictive Value

RBF Radial Basis Function

ReLU Rectified Linear Unit

RMSE Root Mean Square Error

RMS Root Mean Square Level

ROC Receiver Operation Characteristic

ROI Region Of Interest

RPI Raspberry Pi

SDD Seizure Detection Device

sEMG Surface Electromiography

SFTP Secure Shell File Transfer Protocol

SNR Signal-To-Noise Ratio

SSR Squared Sum of Roots

STFT Short-time Fourier transform

SUDEP Sudden Unexpected Death In Epilepsy

SVM Support Vector Machine

TCF Temporal Consistency Factor

WHO World Health Organization

Chapter 1

Introduction

1.1 Introduction to Epilepsy and Current Challenges

Epilepsy is a neurological condition that is characterized by the repeated occurrence of unprovoked seizures. Seizures are periods of uncontrolled electrical activity in the brain that can cause involuntary movement, sensory alterations, and periods of loss of consciousness. Epilepsy is very heterogeneous and can manifest itself in many different epilepsy syndromes, some of which are more serious than others [1]. It is one of the most prevalent serious brain conditions, affecting around 50 million people worldwide, 80% of which live in low- and middle-income countries, and accounting for 0.5% of the global disease burden, according to the World Health Organization (WHO) [2]. Patients with epilepsy are at a substantially higher risk of suffering a premature death [2], especially when they do not receive immediate medical attention following severe tonic-clonic seizure episodes [3].

Tonic-clonic seizures are possibly the most worrying of the many seizure types. These seizures transition from a period of tonic movement (sustained muscle contraction, resulting in rigidity and stiffness) to a period of clonic movement (rapid jerking of extremities) and can cause loss of consciousness, incontinence, teeth-clenching, or tongue biting. Tonic-clonic seizures are arguably the most dangerous seizure type, as they can increase the risk of Sudden Unexpected Death in Epilepsy (SUDEP) (which is the most common cause of death associated exclusively with epilepsy [4]) by a factor of 15 $\times$ [3].

The prognosis for epilepsy is generally positive, with around 65% of patients going into remission and achieving some level of seizure freedom [5], partly due to antiepileptic drugs. However, these drugs can have many serious side effects which can negatively affect patients' quality of life [6], especially as physicians often underestimate the impact these side effects have on patients' lives [7]. Additionally, around 30% of patients suffer from pharmacoresistant epilepsy, also known as refractory epilepsy [8]. For some people, this can be addressed by surgically removing the area of the brain responsible for seizure activity, with a fairly high success rate. However, there are risks, and the process of preparing surgery and assessing whether a patient is eligible can be long and complicated [1, 9].

Likewise, the process of diagnosing epilepsy is also long and complex. The typical diagnosis process requires taking a detailed patient history, physical and neurological examinations, medical imaging

exams, electroencephalography (EEG), and often also long-term video-EEG - in case epileptiform activity is not detected in a regular 30-60 minute EEG or the type(s) of seizure the patient has suffered is not known [1]. Long-term video-EEG generally consists of a combination of video recording and EEG acquisition that is performed during long periods, lasting hours or several days [10]. This method is very useful for detecting seizures and diagnosing epilepsy, but its length and the arduous process of annotating seizures are challenging. Even then, diagnosis is not always accurate [11], which naturally leads patients to not have their medication properly adjusted, reducing their chance of reaching remission, and with it the chance of a normal life. One factor that affects the accuracy of the diagnosis and the adjustment of patients' medication is the fact that patients are unaware of over 50% of their seizures [12].

Some studies have estimated that half of patients with epilepsy report that they suffer from social stigma because of their condition [13]. Moreover, WHO reports that social stigma and discrimination against people with epilepsy is still common worldwide, with patients being, for example, "denied work, the right to drive or marriage" [2].

The long diagnosis process and therapy side effects can cause substantial anguish in patients and caretakers as it pertains to their quality of life and mental health [7], which is exacerbated for more serious epilepsy syndromes, such as Dravet Syndrome, Temporal Lobe Epilepsy, and Juvenile Myoclonic Epilepsy [14, 15].

1.2 Motivation

The challenges mentioned in the previous section highlight the need to provide solutions that enable patients with epilepsy and caregivers that suffer from the many physical, emotional and social consequences of epilepsy to closely monitor their and their children's condition. Patients' and caregivers' quality of life is severely affected by this condition, especially when seizure freedom is not achieved (or delayed). As such, it is essential to not only guarantee that health professionals have access to the best and most complete information about the patients' condition but also reduce patients' and caregivers' stress levels by providing them with solutions that reduce their fear of consequences like SUDEP or injury that can arise from serious seizure episodes.

Seizure detection devices (SDDs) and methods can have a very important role to play in addressing these challenges. In a clinical setting, seizure detection, registration, segmentation or classification methods could facilitate and validate seizure annotation for video-EEG. In an ambulatory setting, these methods could be used to relay information to doctors which could make medication adjustment easier by helping classify seizures and their localization in the brain, thereby covering for unreliable patient histories by registering all seizures. Furthermore, these methods could also help reduce the need for lengthy video-EEG sessions that would have this purpose and tranquilize worried caregivers and patients and improve their quality of life by alerting them when seizures may be occurring, especially at night.

The patient groups who would benefit most from seizure monitoring and detection devices or methods would be: (a) Patients who suffer from refractory (pharmacoresistant) forms of epilepsy; (b) patients

whose medication is not yet fully adjusted, such as those who have been recently diagnosed or still are undergoing the diagnosis process; (c) infants who suffer from epilepsy, especially those with serious forms such as Dravet Syndrome and Juvenile Myoclonic Epilepsy; (d) Patients who suffer from nighttime seizures; (e) patients who are not aware or do not recall seizure events; and (f) and patients who suffer from tonic-clonic seizures, and as such are at higher risk of SUDEP or injury [3].

Multiple inquiries of patients with epilepsy and caregivers have been performed and published recently. These have concluded that despite overwhelming interest in SDDs, and an expectation that these could help increase autonomy and freedom, decrease anxiety, alert caregivers or family members, and help patients more objectively log seizures and share this information with health professionals, most patients and caregivers have not acquired them, due to very high prices and concerns about effectiveness and accuracy [16–19]. Although most patients would not be ready to pay more than $200, most devices on the market exceed this price point, with some even costing over $1000 [1]. Patients expected a sensitivity above 90% and a low false alarm rate, preferably lower than one per week. This standard has not been met by the devices currently on the market [20]. Besides these expectations, other factors such as comfort and usability, privacy/confidentiality of data, low intrusiveness, low visibility, low obtrusiveness, attractive appearance, and good technical support were also considered important [16, 18, 20–22].

Many scientific papers have been published detailing novel methods for seizure detection. These methods differ in their modality, the types of seizures they detect, their accuracy, level of intrusiveness or obtrusiveness, computational complexity, and many other relevant parameters. A lot of them report relatively high accuracy and detect a wide array of seizure types, with some methods being more specific than others as to what seizure types they attempt to detect [23, 24].

There are also various SDDs currently commercially available. Among these are several wearable devices such as smartwatches or bracelets; under-mattress sensors; surface Electromyography (sEMG) sensors; and many others (see Chapter 3). Interestingly, despite the existence of video-based motion detectors and monitors, and the fact that multiple published papers detail video-based seizure detection algorithms, **there does not seem to be any video-based seizure detector** on the market, with medical approval or otherwise [20]. While some patients shunned video-based devices due to privacy concerns, caregivers supported and emphasized the need for video monitoring, feeling that the benefits of this modality outweighed these concerns [21, 25]. Existing products suffer from various problems, such as very high prices, unsatisfactory sensitivity and specificity, high obtrusiveness or intrusiveness, among others [20, 21].

Developing a device that satisfies these requirements is an important challenge. Affordability requires deft component choice. Achieving the accuracy that patients need requires developing an algorithm that surpasses the current state of the art. And developing a truly unobtrusive device requires a novel outlook; a new way of envisioning wearable devices. One of the main ideas that preceded and motivated this work was the recent emergence of a novel concept, namely that of *"invisibles"* [26], which are unnoticeable devices that can be seamlessly integrated into a user's life. An example of such a device is "Emerald", a discreet device for monitoring patients with Parkinson's Disease that monitors sleep stages

[1] List of popular seizure alert systems on the *Epilepsy Action* UK-based epilepsy charity website: https://www.epilepsy.org.uk/info/daily-life/safety-aids-equipment/alarms-monitors, last accessed on 22/10/2021.

and gait patterns using radio waves [27]. *"Invisibles"* are the future of *wearables* in the ongoing drive to further integrate modern life with technology, particularly in the context of healthcare, where *"invisibles"* can facilitate physiological monitoring and enable improvements in quality of life and diagnostic accuracy for patients with conditions like epilepsy, without negatively affecting their day-to-day life or exposing their medical condition to others, which could lead to stigma [2].

These factors motivated the development of an accurate, affordable, accessible, unobtrusive, and privacy-preserving video monitoring and **automatic seizure detection device** encased within a light fixture called **Lampsy**.

1.3 Objectives

The main objective of this work was to develop **an "invisible" device capable of detecting major motor seizures in epilepsy patients** (especially *generalized tonic-clonic seizures*, but also *focal to bilateral tonic-clonic* and *clonic* seizures (see Section 2.1), in order to provide a solution to the challenges outlined in the previous sections. It is essential to deliver peace of mind to patients with epilepsy and caregivers and to simplify, shorten and improve the accuracy of the diagnosis process by providing accurate seizure registration and assisting in seizure classification, thereby ultimately reducing the need for long-term video-EEG sessions, assisting the adjustment of medication, and helping the process of preparing surgery for refractory epilepsy.

Within this scope, one of the main components of this project was the development of a **seizure detection algorithm** to be coupled with the device. This algorithm would need to automatically record, denoise, process, segment, and classify seizures in video footage, as well as deliver a timely alarm.

An additional target was to match the needs of patients and caregivers, by developing a device that is affordable, easy to use, unobtrusive and matched the accuracy requirements of patients and caregivers.

From the outset, the author defined four main tenets that would help guide and define the ideation, development, and operation of this device:

- *Affordability*: The design and component choice of the device should take into account its price, to make it accessible to everyone who needs it.
- *Unobtrusiveness*: The device should not interfere with the day-to-day life of the user, and should be as unnoticeable as possible to facilitate its use and integration with the user's home, without compromising its functionality.
- *Accessibility*: The device should require a minimum amount of input or setup from the user, in order to enable people with disabilities, older people, or less technologically adept users to use the device with ease and without technical difficulties.
- *Privacy*: The device and algorithm must preserve the privacy of the user by design, by creating representations of the video footage that do not allow identification of the patient, as well as ensuring the data is encrypted and its transmission is secure.

This work proposes an innovative solution called Lampsy. Lampsy will be a video monitoring and automatic seizure detection device discretely encased within a light fixture. It must record and process video, detect and isolate possible seizure episodes, issue an alarm, and securely transmit isolated video segments of possible seizure episodes to a secure location or database. The light fixture casing ensures that Lampsy is integrated within the user's home, thereby following the principles of *"invisibles"* and ensuring *unobtrusiveness*. Very importantly, Lampsy must generate unidentifiable representations of the video data, encrypt, and securely transmit it only with explicit consent, such that it respects the principles of *privacy by design*. It must also include a cover for the camera, an ON/OFF button, and a RECORD button, to ensure *accessibility* and practicality. Lampsy must be made of cheap and widely available components, ensuring *affordability*.

1.4 Original Contributions

This project generated original scientific, methodological, and engineering contributions:

- A patent request was filed for "LAMPSY - Dispositivo e Método para Monitorização de Epilepsia Preservando a Privacidade do Paciente e Encapsulado em Luminária". (Privacy Preserving Epilepsy Monitoring Device and Method Encased in Light Fixture) after official approval was issued by the Administrative Council (Conselho de Gestão) of Instituto Superior Técnico, following a request to the IST patent office;
- A device based on a Raspberry Pi 4 with an IR LED camera and encased within a light fixture was developed as an architectural basis for Lampsy and placed in Hospital de Santa Maria for raw video data acquisition of patients with epilepsy;
- This prototype preserved patient privacy by design, as it contained a wireless 4G router to securely and automatically transmit the recorded data via SFTP (Secure File Transfer Protocol) to a Network Attached Storage (NAS) unit in the Image Processing and Analysis lab in IT (Instituto de Telecomunicações);
- Novel noise reduction methods for Optical Flow data were developed;
- A novel seizure detection and segmentation algorithm was developed, based on Optical Flow (OF), Independent Component Analysis (ICA), and Machine Learning (ML) classification;
- A novel ICA source selection metric named the *Temporal Consistency Factor* was proposed, as well as novel features for classification;
- A novel methodology for unobtrusive seizure monitoring and integration of "invisible" devices in a user's home environment was envisioned.
- An extensive review of the state of the art in seizure detection methods and devices, as well as patients' and caregivers' opinions regarding them was performed.

1.5 Thesis Outline

This work is divided in 6 chapters: Introduction, Background, Prototype, Seizure Detection Algorithm, Results and Discussion, and Conclusion and Outlook.

- The current chapter, “Introduction”, is comprised of the motivation for this work, the main objectives the author set out to achieve within the project and the original contributions achieved within this work.
- Chapter 2 features an introduction to epilepsy, and an introduction to some signal processing tools used in this work. The aforementioned introduction to epilepsy contains an account of its epidemiology, a description of seizures and seizure types, a section on epilepsy diagnosis, and finally one on the prognosis and treatment for this condition, as well as an current challenges. After this, Optical Flow, Independent Component Analysis and machine learning methods for classification are discussed.
- Chapter 3 includes an outline of the views and opinions of patients with epilepsy and caregivers on seizure detection devices, followed by a thorough account of the current state of the art in epilepsy monitoring and seizure detection methods and devices, with a specific focus on video-based seizure detection devices.
- Chapter 4 concerns the development of a video monitoring device that served as the architectural basis for Lampsy. This chapter details the choice of components and the construction of this device and the adopted data transmission and recording methodology.
- Chapter 5 describes the development of the seizure detection algorithm elaborated within this project. This chapter entails a short account of the software and datasets used in this process, the pre-processing and noise reduction performed on this data, and a detailed description of how the algorithm itself works. This description includes the ICA-aided artifact removal and segmentation, as well as the isolation of seizure episodes using novel signal processing seizure metrics, and the classification of possible episodes using different machine learning methods.
- Chapter 6 includes results obtained from the novel methods implemented in this work, such as the noise reduction, the source selection and artifact removal using ICA and seizure metrics, and the algorithm’s performance in detecting and segmenting seizure episodes.
- Finally, chapter 7 contains concluding thoughts, limitations, and future work.

Chapter 2

Background

2.1 Epilepsy

Epilepsy is a chronic neurological condition characterized by the **repeated occurrence of unprovoked (or reflex) seizures**, which are short periods of abnormal, uncontrolled, and excessive electrical brain activity that disrupt normal brain function [1, 2].

Seizures can be said to be provoked or unprovoked, in which case they may be triggered, for example, by alcohol withdrawal, fever, or stroke [28]. According to the World Health Organization (WHO), "Epilepsy is defined as having two or more unprovoked seizures" [1]. This has been the general definition, because while after one unprovoked seizure there is around a 40-50% chance that a second seizure will occur [29], after two the chance of recurrence in the following 5 years rises to roughly three in four [30]. However, the International League Against Epilepsy (ILAE) published a report in 2014 [31] that argued that some seizures caused by stroke, head trauma or infection lead to a high risk of recurrence, similar to that of two unprovoked seizures, and therefore defined epilepsy as "*a disease of the brain defined by any of the following conditions:*

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

In a sense, characterizing epilepsy as one condition or disease can be regarded as somewhat incomplete, as there are various different epilepsy syndromes that are characterized by, for example, the occurrence of different seizure types, at different ages, with different prognoses and treatment options [1].

[1] WHO Epilepsy Fact Sheet, available online at https://www.who.int/news-room/fact-sheets/detail/epilepsy

2.1.1 Epidemiology

The prevalence of a disease is defined as the number of people per 1000 that suffer from it at a given point in time within a specific population [11]. It is estimated by WHO that epilepsy affects roughly 50 million people worldwide [2], and according to a study in *The Lancet Neurology* had an age-standardized prevalence in the general population of between 0.5% and 0.7% [32]. In poorly developed countries this value is seemingly higher according to various studies, reaching 6-10/1000 [11]. Additionally, 80% of people with epilepsy live in developing countries [2]. The chance of someone developing epilepsy over their lifetime, i.e. the lifetime prevalence, is also much higher, reaching 5% [11].

The incidence of a disease is defined as the number of people that have been diagnosed with it over the past year, per 100 000 people at risk for said condition [11]. According to WHO, the incidence of epilepsy is much higher in poorly developed countries, with around 139 people in 100 000 being diagnosed each year in these countries, as opposed to 49 in 100 000 in developed countries. This is mainly due to factors such as the prevalence of diseases like malaria, the higher incidence of traffic and birth-related trauma, and differences in medical infrastructure and access to care. Incidence also varies with age, being higher for children and older people, and lowest for adults [1].

2.1.2 Seizure Types

There are many different types of seizures. Some only cause periods where the patient loses awareness, while others can cause motor symptoms, altered sensory experiences and loss of consciousness [33]. Seizures can cause a variety of types of motor activity, which can be characterized as follows:

- Atonic movement activity is the sudden loss of muscle tone [1, 33];
- Tonic movement comprises sustained increases in muscle contraction causing stiffness, extension of extremities and a rigid posture [1, 33];
- Clonic motor activity is characterized by uncontrollable and rapid muscle contraction and relaxation, leading to jerking of extremities [1, 33];
- Myoclonic activity, or myoclonus, are very short unintentional muscle contractions. Slightly longer unintentional contractions are labeled epileptic spasms [1, 33];
- Automatisms are nonpurposeful, repetitive and semi-coordinated movements often associated with impaired awareness [34, 35].

In 2017, the ILAE defined a new operational classification of epileptic seizures, as shown in *Fig. 2.1* [36]. This new classification defined new seizure types and designated clearer names for some of the existing seizure types.

Seizures are categorized primarily by their onset as **focal, generalized or unknown onset seizures**. Focal seizures originate in a seizure focus in one hemisphere of the brain, typically in a small group of neurons in the cerebral cortex. Changes in these neurons' excitability, possibly caused by tumors, localized infections or abnormalities, trauma, or stroke, produce the abnormal and uncontrolled electrical

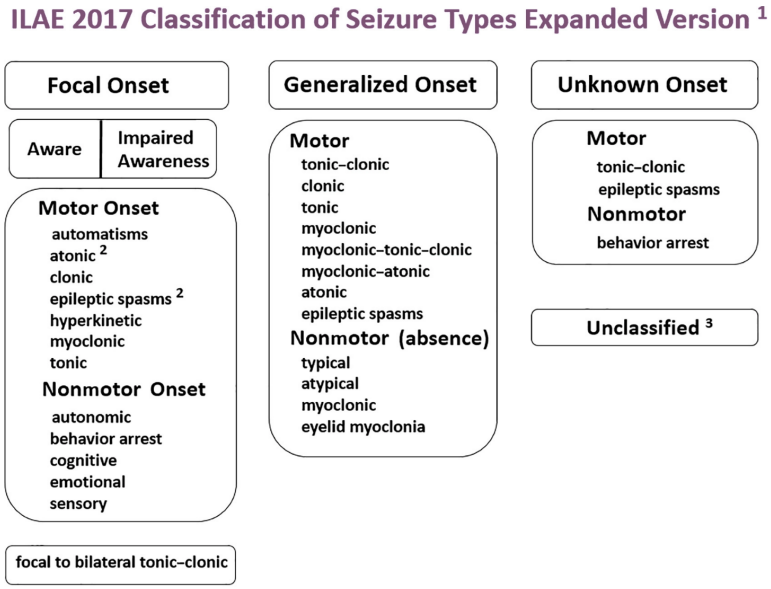

Figure 2.1: Operational classification of seizures regarding onset (focal, generalized or unknown), awareness in case of focal seizures, nature of first prominent sign or symptom (motor or nonmotor) and type of movement or nonmotor symptom, as defined by the International League of Epilepsy (ILAE) in 2017. From: "*Operational classification of seizure types by the International League Against Epilepsy: Position Paper of the ILAE Commission for Classification and Terminology*" [36]

(epileptiform) activity that characterizes a seizure. This activity can then spread to other areas of the brain. The localization of this epileptiform activity is directly related with the aforementioned motor symptoms, loss of consciousness, or altered sensory phenomena [33].

Focal seizures were previously called partial seizures and were classified primarily as simple partial seizures, and complex partial seizures. This distinction concerned whether the patient was conscious during the entirety of the seizure (in simple partial seizures) or there was a period of impaired awareness (complex partial seizures). Impaired awareness can range from lower responsiveness or confusion to unconsciousness or lack of recollection of the event [1]. The aforementioned 2017 ILAE report defined a new classification of these seizures as focal aware and focal impaired awareness seizures, as this would be clearer [36]. Often, focal seizures start with an *aura*, which is a subjective sensory phenomenon that may precede an observable seizure or constitute a nonmotor seizure if it occurs by itself [37].

These seizures can also be classified as motor or nonmotor onset focal seizures according to the nature of the first prominent sign or symptom. As the seizure progresses and other signs become prominent, these can be added as descriptors. It's important to underscore that this classification is operational, and therefore flexible to the specific circumstances of the seizure [36].

Motor or nonmotor focal seizures can be further classified according to their most prominent motor or nonmotor sign [36]. The movement characteristics of focal seizures are similarly localized in parts of the body, with tonic or clonic activity often being lateralized. This movement generally also occurs contralaterally to the affected brain hemisphere [1]. Focal seizures with motor onset can present a myriad of signs, differing in the type of movement (such as clonic, tonic, atonic), with automatisms being the most common motor manifestation [1, 36]. Nonmotor symptoms include: emotional phenomena

such as déjà-vu or experiencing emotions like fear, sadness or happiness; sensory phenomena, such as seeing or hearing things that do not correspond to reality; and a variety of other alterations in cognition and perception, for which we refer the reader to the glossary published by the ILAE [37].

Focal seizures can also spread to the entirety of the cortex and become generalized seizures. These were formerly called secondarily generalized seizures, but were later reclassified by the ILAE as focal to bilateral tonic-clonic seizures [36]. Often, these occur in a clonic-tonic-clonic pattern, with clonic focal activity preceding a generalization leading to tonic-clonic activity [1].

Generalized onset seizures, on the other hand, are seizures that start in both hemispheres, stemming from alterations to structural and functional connectivity as well as physiological alterations [38]. These seizures can also be classified according to their motor or nonmotor components. Nonmotor generalized seizures are called absence seizures (formerly known as petit mal seizures). These are characterized by a sudden impairment of awareness and responsiveness, with the patient often staring blankly ahead for a short period of time, after which the patient has no recollection of the event [1, 6].

Generalized seizures with motor symptoms include tonic, atonic, clonic, myoclonic, and generalized tonic-clonic seizures (GTCS). The latter, formerly known as grand mal seizures, are possibly the most important, both in the context of epilepsy as a whole and in the specific context of this work, because they are arguably the most dangerous seizure type, as they can increase the risk of **Sudden Unexpected Death in Epilepsy (SUDEP)** (which is the most common cause of death associated exclusively with epilepsy [4]) by a factor of 15$\times$ [3]. Additionally, a study by W. Allen Hauser published in *Epilepsia* in 1993 ascertained that GTCS are the most common type of generalized seizure and the second most common type of seizure, constituting 23% of all epileptic seizures [39]. These are also the most common nonfebrile seizure type in children [40].

GTCS start with loss of consciousness and postural control, followed by a short period of tonic motor activity (which is, as mentioned, sustained muscle contraction leading to stiffness, extension of extremities and a rigid posture), usually lasting less than 30 seconds. Patients often clench their teeth, bite their tongue, and breathing can stop (apnea) leading to cyanosis (blue coloring of extremities and lips due to lack of oxygen). Following this, as there is a transition from tonic to clonic movement (fast and uncontrollable contraction and relaxation of multiple muscles, leading to jerking motion), where a "vibratory" period starts [41], and there is very high frequency movement, which slows down as the clonic phase ends. Patients also often lose bladder and bowel control [1, 6, 33].

2.1.3 Diagnosis and Monitoring

Diagnosis of epilepsy requires acquiring a detailed patient history, gathered with the patient. This history can entail information given by the patient themselves, such as memories of seizures and events that triggered or preceded seizures, or information gathered by witnesses, such as patient behaviour during seizures. This is an important step in ascertaining information on seizure type, as well as any previous event that may have increased seizure risk, such as stroke, trauma, infection, etc. [1]. Other causes of seizures also need to be ruled out, to confirm that these were unprovoked seizures. Physical

and neurological examinations need to be performed to rule out other neurological conditions, and brain imaging exams such as magnetic resonance imaging (MRI) or computed tomography (CT) may be performed to find possible seizure causes such as the above-mentioned injury, stroke, infection and others [1, 6]. A significant challenge is the unreliability of patient histories. Studies have found that patients are unaware of (and therefore do not report) over 50% of their seizures, which often leads to misdiagnosis or poor adjustment of medication [12].

One of the most important, if not the most important technique applied to epilepsy diagnosis is Electroencephalography (EEG), which is a medical technique that detects cortical electrical activity in the brain, using electrodes placed on the surface of the scalp or intracranially [42].

Nonetheless, often, to enable better seizure classification and diagnose specific epilepsy syndromes, seizures may need to be observed and recorded using long-term Video-EEG. Video-EEG is a technique whereby EEG acquisition is combined with video footage of the patient over a period lasting between hours and weeks, during which the patient is occasionally sleep deprived in order to increase the probability of a seizure occurring [10]. This method is very useful for detecting seizures and diagnosing epilepsy, but its length and the arduous process of annotating seizures are challenges.

A significant challenge in epilepsy treatment is diagnostic accuracy. Wrong diagnoses can cause patients to not have their medication properly adjusted, which can significantly worsen prognosis. Factors that lengthen the diagnostic process and affect its accuracy are: (a) the heterogeneity of possible manifestations of epilepsy; (b) the different causes of seizures; and (c) the need to observe seizures in order to properly diagnose epilepsy due to the lack of permanent neurological symptoms outside of seizures [11].

2.1.4 Prognosis, Treatment And Challenges

Epilepsy can be treated with medication and in general the prognosis for new epilepsy cases is fairly positive. According to the British National General Practice Study of Epilepsy, conducted in 1995 and published in *The Lancet Neurology*, 65% of patients who received treatment after the first seizure reached remission (defined by being seizure-free for 5 years). While this number was 5% higher for those who did not receive treatment, the study mentions that this results from the fact that patients who did not receive treatment generally had milder forms of epilepsy and thus could achieve remission even without treatment [43].

According to a 2016 study published in *Seizure* [5], which compared data from the Epilepsy Unit in the Western Infirmary in Glasgow, Scotland from 1982 to 2012, around 64% of patients with epilepsy had gone into remission (defined by being seizure-free for at least 1 year). Nevertheless, this number had not changed significantly in the 30 years that this Epilepsy Unit had been open. This can indicate that although treatment options have improved over the years in some respects, there have not been significant advances in increasing the number of patients who reach remission. Nonetheless, there have been some recent advances in the development of antiepileptic drugs, some of which have benefits such as reducing drug interactions and side effects, although they do not seem contribute to a decrease in

the prevalence of pharmacoresistant epilepsy [44].

A study published in *Epilepsia* in 2007 concluded that approximately 20-40% of patients have drug-resistant epilepsy, also called refractory epilepsy [45], which can lead to increased mortality and comorbidities [44]; other studies placed this number at over 30% [8]. Refractory epilepsy can, however, be treated with surgery, by resecting (removing) an area of the brain responsible for the patient's symptoms, with special attention to what areas in and around the *epileptogenic zone* could be responsible for essential or important brain functions. Before surgery is chosen as a viable therapy, patients need to undergo a variety of tests to locate this area and prepare the surgical proceedings, including medical imaging methods and long-term video-EEG [1]. Surgery has generally been shown to be very effective in achieving remission in patients with refractory epilepsy [9].

A 2005 review study on the epidemiology of epilepsy in Europe found that between 8% and 25% of patients with epilepsy have, on average, over 1 seizure per month, and between 11% and 30% have over one seizure per week [46].

Despite the generally positive prognosis, it is estimated that the risk of premature death in people who suffer from epilepsy is up to three times greater than the general population [2]. The biggest cause of premature death caused by epilepsy is SUDEP [4]. The mechanisms that underlie this event are unknown, but it is generally associated with the occurrence of a generalized tonic-clonic seizure. Proposed mechanisms include cardiac causes such as ictal bradycardia or pulmonary causes such as pulmonary edema and oxygen desaturation, both because of and during GTCS [1]. There are certain additional risk factors for SUDEP, such as high seizure frequency, youth, and long duration of epilepsy [47]. A study published in *Lancet* found that there was a 42% increase in mortality in newly diagnosed epilepsy cases that was only present in patients who were not in remission (i.e. those that were not responding to treatment). Patients with chronic epilepsy had double the mortality risk and also double the risk of SUDEP, namely 2.46 per 1000 compared to 1.08 per 1000 for newly diagnosed patients. Additionally, this risk was higher for those under 30 [48]. Furthermore, patients with epilepsy are also at increased risk for comorbidities. A study in 2004 found that patients were at double the risk to develop psychiatric conditions and at increased risk to develop neurodegenerative diseases such as Alzheimer's or Parkinson's Disease [49].

WHO estimates that around 25% of epilepsy cases are preventable, by, among others, reducing the risk of head trauma, improving sanitary conditions to reduce infections, and reducing the risk of stroke by reducing diabetes, obesity, alcohol, and tobacco use [2].

There are still significant challenges in the treatment of epilepsy, besides the already mentioned lack of dramatic advances in recent years regarding refractory epilepsy and remission rates. One such challenge is the lack of high quality drug trials targeting epilepsy [50], which can stifle the development of new therapies. In the same trend, although many new drugs for epilepsy have recently entered the market, these have not been shown to be more successful than older drugs when not used in conjunction with them; moreover, there have not been enough studies to discover their complete mechanism of action, their optimal concentration, and their efficacy for all specific epilepsy syndromes and manifestations [51]. Adverse side effects of antiepileptic drugs can often be serious, which can negatively affect

patients' quality of life [6], especially as physicians often underestimate their impact on patients' lives [7].

As mentioned, the lengthy diagnosis process can be a significant challenge in quickly starting and properly calibrating the medication the patient needs. Moreover, there are enormous challenges with respect to refractory epilepsy. Some of these concern the length and complexity of the process of preparing surgery and finding the epileptogenic zone, while other challenges stem from the low number of high quality studies. This situation results in a lack of knowledge about whether some forms of epilepsy are inherently refractory or could be treated with the right medications under specific conditions [1, 50, 51].

Beyond the significant burden that these challenges cause in patients with epilepsy and caregivers, particularly in those with refractory epilepsy, studies have estimated that half of patients with epilepsy also report that they suffer from social stigma because of their condition [13]. WHO reports that social stigma and discrimination against people with epilepsy is still common worldwide, with patients being, for example, "denied work, the right to drive or marriage" [2]. This stigma is often associated with lack of understanding of this condition by many in the general population [52].

Finally, some epilepsy syndromes are particularly difficult for patients and caregivers. For example, Dravet Syndrome, also known as Severe Myoclonic Epilepsy in Infancy, which is characterized by focal and generalized myoclonic, absence, tonic, and tonic-clonic seizures in young infants, with the possibility of developing cognitive and behavioral disorders [53], can lead to a very significant burden on caregivers, affecting their mental health and socioeconomic conditions [15]. Studies have also reported significant impact on caregivers of patients with other serious epilepsy syndromes such as Temporal Lobe Epilepsy and Juvenile Myoclonic Epilepsy [14].

2.2 Signal Processing And Classification Techniques

2.2.1 Optical Flow

Optical flow (OF) can be described as "**the distribution of apparent velocities of movement of brightness patterns in an image**", according to Horn and Schunck's seminal paper [54], published in 1981. As optical flow is effectively a projection of the real motion of objects in three-dimensional space into a two-dimensional plane, its calculation is useful for many computer vision applications, such as video compression, restoration or even physiological monitoring, such as what is explored in the current work [55]. Not only is OF useful for tracking the motion of objects in a video sequence, it can also be used to determine object positions, shapes, or acquire depth information [54].

OF is computed by assuming that **the brightness of one point in an object or pattern remains consistent if the object moves** (i.e. over a short time-frame the value of the brightness in a point in a physical object in space remains the same if that point moves to a different point in space). This is called the **brightness constancy constraint** (which can also be referred to as the color constancy constraint, in case OF is calculated using multiple colors instead of just brightness in black and white footage) [54, 55]. However, this constraint is generally not valid in real world conditions because lighting

conditions and object surfaces are not uniform, and as such as an object moves its apparent brightness in the video sequence changes. Nevertheless, this constraint is essential in the calculation of OF, and is a useful approximation. It can be formulated mathematically as follows [55] (Eq. 2.1):

$$I(x, y, t) = I(x + \frac{\partial x}{\partial t}, y + \frac{\partial y}{\partial t}, t) \tag{2.1}$$

In this equation, $I(x, y, t)$ represents the brightness I of a specific point in the 2-D video frame space (x and y) and in time (t). Moreover, because this constraint requires the brightness of a point to remain constant over time, it can naturally also be expressed as the derivative of I with respect to time equalling zero, as done by Horn and Schunck [54] (Eq. 2.2):

$$\frac{\partial I(x, y, t)}{\partial t} = 0 \tag{2.2}$$

This can be expressed as a linear combination of the partial derivatives of I multiplied by the partial derivatives of x, y and t with respect to t, according to the chain rule [54]. The partial derivatives of x and y with respect to time are the horizontal and vertical velocities of the point in (x, y). These can be expressed as $u(x, y, t)$ (the horizontal velocity of a point in an object in 2-D space) and $v(x, y, t)$ (the vertical velocity of such a point) (Eq. 2.3):

$$u\frac{\partial I}{\partial x} + v\frac{\partial I}{\partial y} + \frac{\partial I}{\partial t} = 0 \tag{2.3}$$

However, this is merely an approximation, as it assumes that brightness is differentiable, i.e. it must vary smoothly across space, which is not accurate for edges or occlusions [54]. Moreover, this equation is not sufficient to estimate flow locally. This is due to the fact that movement on an edge, when viewed through a small aperture (such as when computing OF locally), can only be estimated in the direction orthogonal to the edge, and not parallel to it. Similarly, when computed locally, OF can only be computed in the direction of the brightness gradient; in other directions there is not enough information to compute it. This problem was explained by Horn and Schunck in 1981, but only later did it receive its current name: the **aperture problem**, coined by David Marr in his posthumously published 1982 book "*Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*" [56]. Thus, locally, OF estimation is undetermined, and therefore other constraints need to be introduced to enable its estimation. One of these is proposed by Horn and Schunck [54]: the **smoothness constraint**. This assumes that the brightness variations in pixel *neighbourhoods* behave consistently, as motion tends to occur in objects larger than one pixel, and points in one object tend to have similar (or at least consistently varying) velocities.

OF can be calculated in evenly spaced points in the image, which can be called **dense optical flow**, or it can be calculated sparsely, by tracking relevant features and only computing OF on them [55]. An

example of a dense OF method is the aforementioned Horn and Schunck method [54], and an example of a sparse method is one developed by Peter Sand and Seth Teller in 2008 in [57].

The method used in this work was the Farnebäck method [58], an established and computationally efficient method for calculating dense Optical Flow based on polynomial expansion. This method, although somewhat noisy, was chosen because it enables accurate OF calculation with only two frames, enabling it to be calculated in real-time. Moreover, the fact that it does not employ spatiotemporal smoothing makes it adequate for accurately portraying very fast movement, such as what is present in clonic seizure movement. More details about the implementation and the choice of method are presented in Chapter 5.

2.2.2 Independent Component Analysis (ICA)

Independent Component Analysis (ICA) is a computational method used for separating and identifying hidden statistically independent and non-Gaussian components (or sources) from multidimensional data. This multivariate data is comprised of a mixture of independent signals, and ICA's goal is to separate the mixed signals into their original independent components [59].

ICA is a Blind Source Separation (BSS) method, because it aims to separate these hidden sources with minimal to no information about them [60]. The classic representation of source separation and ICA is the *Cocktail Party Problem*, which can be described as follows:

Suppose there are multiple people holding conversations in a cocktail party. Most people concentrate on a specific conversation, as the human brain is capable of separating these voices and tuning in to a specific conversation, but if there was a microphone in the room, it would record a mixture of the various voices. The aim of BSS methods is to separate the different independent signals, such as the different voices, from mixtures such as what would be captured from a microphone. Generally, multiple mixed signals (in this case, multiple microphones placed around the room to record different mixtures) are required to separate multiple sources. ICA requires at least as many observations (microphones) as independent sources, in order to properly separate them [59].

The mathematical problem behind ICA is therefore to find a transformation whereby one can obtain the hidden sources from the observed signals. If the observations are represented mathematically as an m-dimensional random vector $\boldsymbol{x} = (x_1, \ldots, x_m)^T$, and the independent sources as an n-dimensional random vector $\boldsymbol{s} = (s_1, \ldots, s_n)^T$, then $\boldsymbol{W}$ is a linear transformation, such that (2.4):

$$\boldsymbol{s} = \boldsymbol{W}\boldsymbol{x} \tag{2.4}$$

When this is stated in the context of the source separation problem, the observed signals $x_i(t)$ can be seen as samples of the random vector $\boldsymbol{x}$, and the k independent sources $s_j(t)$ as samples of $\boldsymbol{s}$. These observations are comprised of a linear combination of the hidden sources and their respective mixing weights, which can be denoted by a_{ij} (2.5):

$$x_i(t) = \sum_{j=1}^{k} a_{ij} s_j(t) \tag{2.5}$$

The weights a_{ij} can be represented as a matrix $\boldsymbol{A}$, called the **mixing matrix** (2.6):

$$\boldsymbol{x} = \boldsymbol{A}\boldsymbol{s} \tag{2.6}$$

The mixing matrix is the transformation that represents the mixing of the hidden sources into the observations, hence being the inverse of $\boldsymbol{W}$, which can be designated as the **unmixing matrix** [59].

The goal of ICA is to estimate these matrices as well as the independent components based on the assumption that these components are statistically independent and non-Gaussian. ICA algorithms differ in how they estimate independence [59].

The algorithm used in this work was FastICA, a fixed-point algorithm developed by Aapo Hyvärinen [61]. Its pre-processing steps consist of centering and whitening the data (applying a linear transformation such that the components have zero covariance and unit variance). The independent components are estimated with an iterative fixed-point algorithm that uses a maximization of non-Gaussianity as a proxy for independence [61].

Often, particularly in cases where there are more observations than hidden sources or the data has very high dimensionality, dimension reduction methods such as Principal Component Analysis (PCA) are used before whitening the data [62].

PCA is a method that reduces the dimensionality of the data by generating new variables (Principal Components) that contain as much of the variability of the original data as possible, in order to minimize loss of information [63]. PCA works by performing an eigendecomposition on the data's covariance matrix, thereby estimating the matrix's eigenvectors and their respective eigenvalues. These eigenvectors are called Principal Components (PCs), and their respective eigenvalues represent the amount of variance in the data that they account for. The first Principal Component therefore accounts for more variance in the data than all successive components, each of which is orthogonal to the previous components and accounts for more variability in the data than the following PCs. Often, to reduce dimensionality, the last eigenvectors, which explain the least amount of variance in the data, are rejected, and only the first are kept. This selection is made in accordance to the desired dimensionality and the amount of information (as a percentage of the total variance of the original signal which is explained by the first few eigenvectors) which should be maintained. This ensures that minimal information from the original data is lost, while dimensionality is reduced [63].

The details of the implementation of this method as well as the justification for its use is explained in chapter 5.

2.2.3 Machine Learning And Classification Methods

Machine Learning (ML) is a field that is concerned with the development of algorithms and models that instead of being explicitly programmed to perform a task, automatically learn from experience, gradually improving their performance in the task at hand [64]. An example of a machine learning problem is the estimation of a house price. There are various parameters that affect this price: some are categorical, such as if the house has a pool or not, and some are numerical, such as the square footage or the median income of the neighbourhood. These variables are called **features**. The algorithm must learn using a **training set**, which contains features and **labels**, which, in this case, are the house prices of houses with a specific set of features with specific values. In order to learn, the model should adjust to the data. This requires a performance measure, such as an *error or cost function*, which must be minimized. After learning, the algorithm must be able to respond to unseen data - a **test set** - with an output (or **prediction**). In this case, the test set would contain houses with a set of features but without a label (a price), and its output would be a price for this house. [65].

This is an example of a *statistical learning* problem - the model should learn from the data it receives in order to find hidden patterns and make predictions on future, unseen data [66]. It is also an example of a supervised learning method. ML methods can be divided in *supervised learning* and *unsupervised learning* methods. In supervised learning, the algorithm receives inputs and outputs, i.e. features and associated labels, while unsupervised learning methods only receive inputs, and algorithms must learn only from the datasets in question. Supervised learning problems include *regression* and *classification* problems, while unsupervised learning methods include clustering and dimensionality reduction. Regression problems aim to estimate the numerical value of an unknown variable, while classification problems aim to determine which class a variable belongs to. The aforementioned house price problem is an example of a regression problem [66, 67], while an example of a classification problem would be object recognition in images. Classifiers and other ML models have *hyperparameters*, which are parameters that control the behaviour of the model and can be tuned using the training set and methods like *cross-validation*. Cross-validation (CV) divides the training set in small batches, trains the algorithm using the remaining data points and then calculates predictions using these small batches successively until the whole training set has been used. K-fold CV divides the training set in k batches and tests on each of them, while leave-one-out CV is a case of this where the model is trained on the full dataset except for one data point, and is evaluated on this data point (and successively for all data points). A special case of this is *leave-one-subject-out* cross-validation (LOSOCV), which separates the data in batches of subjects, in case the data contains multiple data points for each different subject. This method was used in the algorithm devised in this work, because multiple data points were generated for each seizure video and every seizure video could then only be classified by training it with the remaining data points, in order to avoid overly optimistic results.

An important notion in supervised learning problems is the *bias-variance* tradeoff. Bias is the error between model predictions and the real values. In the house price example, the bias would be the difference between the price estimates outputted by the model and the real prices. Variance measures how much the model adjusts to small differences in different datasets, i.e. how variable its predictions

are depending on how it was trained. If a model aggressively minimizes errors in the training set, becoming overly complex, it will have low bias, since the prediction error in the training set is low, but a high variance, as it will produce very variable results when tested on different data. This is called *overfitting*. A model that is not complex enough will not adjusted enough to different data and therefore have low variance but high bias. This model would be *underfit*. The best model needs to balance bias and variance [66].

As this work is concerned with seizure detection, ML classification can be a useful tool in determining whether an observed pattern of movement constitutes a seizure. For this purpose, beyond the choice of an adequate *classifier*, it is also be important to extract useful features from the video signal. This will be discussed in chapter 5.

Three classification methods were tested and compared to ascertain which one would deliver the best performance. These classification methods (Support Vector Machine (SVM), Multilayer Perceptron (MLP) and Gaussian Process Classification (GPC)) will be detailed in this section.

Arguably the simplest classification method that was used was a Multilayer Perceptron. MLPs are neural networks with input and output nodes and multiple hidden layers of nodes connecting them. Each node has an activation function. The most common activation function is the Rectified Linear Unit, or ReLU, which is $f(x) = 0$ if $x < 0$ and $f(x) = x$ if $x >= 0$. Each connection between nodes has weights and biases, which must be trained using *backpropagation* when the model is fit to data, which is an iterative method based on gradient descent whereby the weights and biases are calculated starting from the output layer [66]. Support Vector Machine classifiers, in their most basic form, are linear classifiers (classifiers that separate two classes with a line called a hyperplane) that select a hyperplane that maximizes the margin between itself and the two closest data points, called *support vectors*. For linearly separable data, this is a hard margin, and for non-linearly separable data this is a soft margin, controlled by the parameter ”C”. For non-linearly separable data, non-linear SVMs can also be used, which use kernel functions to map input data points into a higher dimensional feature space where the data is linearly separable, and define the hyperplane that maximizes the margin in this space. This is called the *kernel trick*, and is a very computationally efficient way to transform the input data into a higher dimensional space [67]. The most popular kernel is the Radial Basis Function (RBF) kernel, which is also used in Gaussian Process classification (GPC). This is a method of probabilistic classification based on gaussian processes, which are stochastic process with random variables that are jointly normally distributed. Gaussian processes have a mean function and a covariance function, the latter of which is effectively a kernel. This kernel is used to compute probabilistic predictions on a set of prior functions estimated from the training data. GPCs are useful for small datasets and the choice of possible kernels enables the use of information about the dataset in the choice of model [68].

Chapter 3

State of the Art

3.1 Patient Views and Opinions

According to multiple studies [17, 18, 21], most patients with epilepsy, particularly those with long-term or refractory epilepsy, say that they are interested in seizure detection devices (SDDs). According to an article published in *Epilepsy & Behaviour* in 2021, the percentage of patients interested in SDDs was 97%, while 99% of caregivers said they were interested [21]. Another paper put this percentage at a mere 65%, although most patients that were not interested in seizure detection attributed it to being seizure-free. This paper also found that the need for SDDs was higher among patients with higher seizure frequency. These patients cited having dangerous seizures, including those with "*loss of consciousness, fall, long duration, intense movements, respiratory problems, cluster or status [epilepticus], salivation/vomiting, oxygen drop, migraine, or visual problems*" as some of the chief factors that contributed to their interest in acquiring an SDD [69].

While, according to one study, 44% of patients and 62% of caregivers had heard of SDDs, only 8% and 15%, respectively, had actually used them (18% of the patients and 24% of the caregivers that had heard of SDDs had experience with them) [21]. Various studies have inquired about the reasons for the generalized hesitance in purchasing these devices. Low affordability was mentioned consistently across studies as a significant deterrent in purchasing an SDD [16–18, 21]. In one inquiry, 38% of respondents mentioned cost as the reason they had not acquired an SDD, and more than half would **not spend over $200** on one [21]. Most devices in the market cost between $200 and $600 [21], but some devices can cost over $1000 [1]. These prices can be cost prohibitive for most patients and caregivers. Similarly, a different inquiry found that while 82.5% would spend over $100 on an SDD, 58% of respondents would **not** spend more than $200 [16]. Another paper found that a majority of patients stated that they would not purchase an SDD if it was not covered by insurance [17]. Caregivers seem to be more enthusiastic and accepting of SDDs, and would pay more than the inquired patients would to acquire them [21]. Additionally, perceived notions about SDDs' effectiveness, accuracy (sensitivity and false alarm rate) and reliability, as well as worries that focal and non-motor seizures would not be detected, colored

[1] List of popular seizure alert systems on the *Epilepsy Action* UK-based epilepsy charity website: https://www.epilepsy.org.uk/info/daily-life/safety-aids-equipment/alarms-monitors, last accessed on 22/10/2021.

respondents' views on these devices and contributed to their hesitation in purchasing one [17, 18, 21].

Although studies have found some reluctance in acquiring SDDs, there was, nonetheless, a considerable interest in these devices. Patients and caregivers thought that SDDs could help increase autonomy and freedom, decrease stress and anxiety levels, alert caregivers or family members, and help patients more objectively log seizures and share this information with health professionals. Respondents also hoped SDDs could help determine seizure severity and localization, identify seizure triggers and predict them, and improve diagnosis and treatment [18, 19, 21, 69]. For patients and caregivers, the most important factor in an SDD was accuracy, especially sensitivity (the ability to detect all seizures), and they would sacrifice other desired features for higher accuracy. Desired sensitivity varied, with some studies reporting that patients required a sensitivity of $>$75%, while others claimed that patients and caregivers required between 90% and 100% of seizures being detected. This discrepancy was often caused by the demands of seizure-free patients and those with refractory (pharmacoresistant) epilepsy varying wildly, with seizure-free patients naturally requiring much better accuracy and patients with many seizures being more lenient. Similarly, specificity requirements varied. One study claimed that 75% specificity would suffice. Other papers inquired instead about false alarm rate (FAR) on a per-week or per-seizure basis, with one study determining that for seizure-free patients, the FAR should be 1 false alarm per week, while for patients with more seizures this should be 1 false alarm per seizure [16, 17, 20, 21, 69].

Patients and caregivers valued automated alarms. One study found that 95% of caregivers and 85% of patients with epilepsy agreed with the importance of alarms [21]. While another paper reported a similar figure (i.e. 85% of patients and caregivers valued alarms) [69], a third paper reported that only 20% of patients and 60% of caregivers would find alarms useful [18]. Herrera-Fortin et al. [21] found that it was important to allow user-customization of alarms, such as, for example, enabling text-message based alarms at work and more audible alarms at night. The desired latency for alarms varied between studies, with some patients and caregivers requiring a latency of under 10 seconds [69], and others finding latencies of under 30s and even up to 1 minute acceptable [17, 21]. When one study asked medical doctors about acceptable latency, they responded that between 2 and 5 minutes was acceptable [69].

Beyond the aforementioned affordability and accuracy, when asked about what characteristics should be present in an SDD, most patients and caregivers also mentioned various other factors such as comfort and usability, privacy, confidentiality of data, low intrusiveness, low visibility, low obtrusiveness, attractive appearance, and good technical support [16, 18, 20–22].

Patients and caregivers valued devices that could be used continuously during the day and at night according to some papers [17, 21], with night-time use being rated as slightly more important; a different article found that patients and caregivers preferred night-time use over continuous monitoring [69].

Respondents of these studies emphasized that devices needed to be comfortable and unobtrusive enough as to not interfere with either their sleeping quality or daily activities, and should work unnoticeably [16, 19, 21]. In line with this required unobtrusiveness, patients preferred devices that would be discreet and not look like medical devices, ideally having an attractive appearance [20–22], possibly due to the stigma still associated with epilepsy [2] and patients' desire to be viewed as normal and not be

defined by their condition, as well as due to fears of having their freedom restricted by not being able to drive or work. Consequently, patients generally preferred discreet and removable devices like smartwatches and bracelets, but these choices could be sacrificed for higher accuracy, which was paramount [19, 21].

Data confidentiality and privacy were also worries reported by patients and caregivers. In one study, patients and caregivers highlighted the need for explicit consent when sharing data [22], while in a separate inquiry, although a majority of patients with epilepsy were concerned about these issues, only 30% of caregivers expressed this concern [18]. This split between patients and caregivers was also present in regards to SDDs that recorded video footage, which could be considered one of the most intrusive methods, in terms of how sensitive the acquired information is regarding the patients' privacy. Patients' and caregivers' interest for video monitoring varied very significantly. One study that only included patients found that only 20% would favor "permanent optical surveillance", while another found that 55% of patients would be open to using an SDD with a camera [19, 21]. Caregivers' interest was measurably higher, with one inquiry finding 77% support [21]. Caregivers, particularly those with young children, seemed to prefer to not trust a "black-box" device which would not allow them to visually monitor their children, which is congruent with their aforementioned preference for alarms. In another paper, caregivers mentioned that video tracings were essential, even going so far as to say that the benefits of video monitoring outweighed their privacy concerns [25].

Studies were also conducted on patients and caregivers who had previous experience of using SDDs. One such inquiry found that 10% had stopped using the SDD in question due to reasons related to the device, such as high false alarm rate (FAR), discomfort, and battery life [70]. Another paper found that SDDs contributed to increased independence for adolescent patients as caregivers felt more secure [71]. Borusiak et al. found that the use of an SDD contributed to improved mental health in patients and caregivers, including reduced anxiety and fear of seizures, and also helped reduce the frequency of cosleeping (young patients with epilepsy sleeping with their parents) [72].

In summary, patients and caregivers showed significant interest in SDDs, which they reckoned could help reduce stress levels, increase autonomy and help communication with the health care system. However, cost is a significant deterrent, as most devices cost more than a majority of patients and caregivers can afford. Factors like unobtrusiveness, comfort, appearance and confidentiality are important, but not as important as high effectiveness and accuracy. Sensitivity should be at least higher than 75%, and should be higher for patients with a lower seizure frequency, while both the FAR and the alarm latency should be kept low. Ideally, a sensitivity above 90% and 1 false alarm per week should be aimed for, according to a study published in *Epilepsy & Behaviour* [69], as well as a latency of at least under 1 minute, preferably under 30 seconds [17, 21]. Although patients prefer discreet options such as wearable devices, caregivers value video monitoring highly.

3.2 Seizure Detection Devices and Methods

Many scientific papers have been published detailing novel methods for seizure detection. These methods differ in their modality, the types of seizures they detect, their accuracy, level of intrusiveness or obtrusiveness, computational complexity, and many other relevant parameters. There are also various commercial devices already on the market, some of which are designed to detect seizures, and others simply alert the user about suspicious movement, temperature, or heart rate changes.

According to various systematic reviews of seizure detection methods and devices, some common seizure detection modalities include EEG, accelerometry (ACM), surface electromyography (sEMG), electrocardiography (ECG/EKG), electrodermal activity (EDA), and video detection, among others. Some methods and devices are unimodal, using only one of these modalities, while others are multimodal, incorporating and combining multiple modalities. Existing methods have varying degrees of accuracy, from 2% to 100%, and detect a wide array of seizure types, with some methods detecting more seizure types than others. However, a lot of the most sophisticated methods published in the literature have not yet been made available to the market as seizure detection devices [20, 23, 24, 73, 74].

Electroencephalography (EEG) is, as mentioned in Section 2.1, an established medical technique that detects cortical electrical activity in the brain, using electrodes which are placed on the surface of the scalp or intracranially to measure the electric field generated by the activity in the cerebral cortex [42]. EEG-based seizure detection is the gold standard method, especially when combined with video recording (video-EEG). Various unimodal EEG-based seizure detection methods have been developed, with varying degrees of success. SDDs based on scalp EEG have demonstrated fairly reliable detection sensitivity, ranging from 74% to 100% according to review studies [23, 74]. Most EEG-based seizure detection methods rely on feature extraction, followed by classification using machine-learning methods [74]. One method developed by Xie et al. achieved an accuracy of 98% by combining a wavelet decomposition method with a simple 1-nearest-neighbours classifier [75], while another study achieved a similarly high sensitivity of 98.5% by using a method based on higher order spectra features and support vector machine (SVM) classification [76].

Intracranial EEG has a higher signal to noise ratio (SNR), as the electrical activity in the brain that it measures does not have to pass through the skull and skin to reach the electrodes, but it is far more invasive, requiring surgery [42]. Therefore, for most patients, especially in an ambulatory setting, it is arguably not justified to choose an SDD based on intracranial EEG over scalp EEG. However, intracranial EEG could be used for seizure *prediction*, and not merely detection. A study published in 2005 achieved a sensitivity of 100% in the prediction of seizures, with a 10 minute prediction horizon, albeit only in 1 patient and with a high false detection rate of 1.1 per hour, with the method failing to work in a second patient [77]. Another method for seizure detection using intracranial EEG achieved a sensitivity of around 95% but a fairly high FAR of 0.58 per hour [78].

There are many portable EEG-based SDDs on the market at the moment, some of which have had research published to assess their accuracy. One such device is *epihunter* ®, a device composed of a light headband with 3 dry EEG electrodes and an amplifier with a companion mobile app for video

logging, used for detecting absence seizures. It costs 34-39€ per month, according to its website [2], and has been approved in Europe as a Class I medical device (CE-marked). This device uses an automatic seizure detection method and its sensitivity has been reported by a study as 99.6%, with an average of 5 false alarms per day [20, 79]. The 24/7 EEG™ SubQ is another commercially available EEG-based device, which is implantable subcutaneously and has also been CE-marked for epilepsy monitoring. A study reported that sensitivity was high but FAR was also high at 1 per hour, and seizure detection requires manual visual identification of seizures with automatic assistance. However, this study did find a consistent circadian cycle for when seizures are most likely to occur for the test subjects [80].

Although EEG-based seizure detection is the gold-standard for seizure identification and epilepsy diagnosis, there aren't many EEG-based SDDs on the market for ambulatory use yet. This may be because these devices can often be bulky, uncomfortable and obtrusive, which may lead to the stigmatization of patients and discourage long-term use. The low number of electrodes can lead to poor SNR and spatial resolution, and movement can cause significant artifacts, leading to unreliable detection accuracy when compared to the methods in the literature, most of which were tested with data obtained from regular EEG sessions in a hospital setting [23].

Accelerometry (ACM), on the other hand, is one of the most common modalities used for seizure detection at home, as devices can be discreet and low-profile, leading patients to prefer them over more invasive, obtrusive or uncomfortable methods [21]. Accelerometers are sensors that can measure their own acceleration (and direction of movement) three-dimensionally in space, and can thus be used to quantify and ultimately detect seizure movement when placed in wearable devices such as bracelets or smartwatches [23, 74].

Multiple studies have developed seizure detection algorithms based on wearable accelerometers. The sensitivity achieved in these studies ranged from 80% to 100% [23, 24]. One nocturnal study used a threshold based approach and achieved a sensitivity of 92% and a specificity of 84% [81]. Another article used time domain features and support vector classification to detect convulsive seizures, and reported a sensitivity of 87% and a FAR of 1.16/24h. When only GTCS were considered, these values were better, at 95.8% and 0.64/24h respectively [82].

Additionally, there are multiple commercially available SDDs which employ this modality exclusively, as well as multimodal SDDs based on ACM and other modalities like EDA or sEMG. *Inspyre* ™ by SmartMonitor is a mobile application which uses the accelerometer in a smartwatch (such as, for example, an Apple Watch). This application costs $20 to activate, plus an additional $15 to $50 a month, depending on desired features [3]. Studies involving this device have estimated it can accurately detect between 32% and 92% of GTCS, with another study finding a sensitivity of 88% [83, 84]. The study that reported the highest sensitivity, with 12 out of 13 GTCS accurately detected, recorded a total of 81 false alarms [84]. *Epi-Care* ® is another ACM-based device costing over $1000[4], which has been classified as a Class I medical device in Europe. A paper published in *Epilepsia* found this device had a sensitivity of 90% and a false alarm rate of 0.2 per day [85].

[2] Epihunter website: www.epihunter.com/product, last accessed on 23/10/2021
[3] SmartMonitor website: https://smart-monitor.com/inspyre/, last accessed on 24/10/2021
[4] Epilepsy Alarms UK website: https://www.epilepsyalarms.co.uk/product/epi-care-standard/, last accessed on 24/10/2021

Review studies mention some significant disadvantages of ACM-based devices, namely the possibility of false alarms due to non-seizure movements, the possibility of seizures not being detected due to the limb holding the accelerometer being accidentally pinned down during a seizure and the inability to detect non-motor seizures [23, 74]. However, they remain popular choices due to their discreetness and unobtrusiveness [21].

Surface Electromyography (sEMG) is a non-invasive technique used to measure the electrical activity of skeletal muscles, often needing only one channel. By detecting the electrical activity of muscles, sEMG-based devices can detect motor seizures, especially tonic seizures [23, 86].

There have been a few studies published detailing novel methods for seizure detection using sEMG, with generally positive results [23]. One article used an sEMG-based SDD with one electrode placed on the deltoid muscle and developed an algorithm based on the frequency content of the electromyographic activity, obtaining promising results, with 100% sensitivity and a FAR of 1 false alarm per day, as well as a latency of 13.7 seconds [87]. Another paper, published in *Epilepsia*, used an sEMG sensor placed on the biceps and achieved a sensitivity of 95% for GTCS and only 1 false alarm in 1399 hours of continuous recording, a very low FAR, especially as this was detected during the post-ictal period of a GCTS [88].

This study used a device that is now commercially available and FDA approved called *SPEAC* ®, which is indicated only for periods of rest [20]. Further studies have been published which test this device. One of them reported much lower accuracy than the previous study, with a sensitivity of 76%, a FAR of 2.52 per day, and a PPV of only 0.03. However, according to the author, this might have been caused by the device being improperly attached. The system showed an average latency of 7.7 seconds [20, 89].

This highlights one of the main issues with sEMG, namely that it is very sensitive to proper placement in the muscle and can be detached [23]. Moreover, 28% of patients reported skin irritation due to the EMG electrodes used in the device, which is another disadvantage of sEMG-based methods [20, 89].

Some commercial devices use **sensors placed under the user's mattress** to detect motor seizures during sleep. The most popular devices currently commercially available are the EMFIT MM™ and the Medpage MP5 ®. The EMFIT device consists of a sensor sheet with quasi-piezoelectric material that detects motion with specific characteristics associated with seizures [90] and costs $594 [5], while the MP5 employs a microphone placed under the mattress [23] and costs £246 [6]. In three separate studies, the EMFIT monitor was shown to detect 21%, 85% and 89% of seizures, with one study reporting a FAR of 0.03 per night and another finding a FAR of 0.13 per 24 hours [90–92]. The MP5 did not show particularly reliable detection sensitivity, with one study finding 7% detection sensitivity and another finding a sensitivity of 63% and 4.22 false alarms per 24 hours [93, 94].

Overall, SDDs that use under-mattress sensors have the significant advantage of being unobtrusive, comfortable and discreet, but there could be significant issues with the performance of these devices, according to existing reports.

There is significant evidence that seizures can cause changes in heart rhythm, predominantly tachy-

[5] EMFIT website: https://emfit.com/movement-monitor/, last accessed on 24/10/2021

[6] Medpage website: https://medpage-ltd.com/epileptic-tonic-clonic-seizure-alarm-MP5, last accessed on 14/11/2021

cardia (increased heart rate), but also other changes that can be analysed and quantified using **electrocardiography (ECG/EKG)** [95, 96]. These changes can even precede the clinical manifestation of the seizure by as much as 30 minutes [97], and have been used to detect and even predict seizures. ECG based seizure detection methods have sensitivities ranging from 35% to 99.8%, according to recent reviews [23].

A study published in *Medical & Biological Engineering & Computing* used a clustering algorithm to detect seizures based on changes in heart rate variability derived from the R-R interval, and achieved a prediction sensitivity of 86% [98]. Another study also based the development of a seizure detection method on the R-R interval using ECG, achieving a sensitivity between 86% and 98% and a FAR of between 1.1 and 9.5 false alarms per hour. With one set of parameters, the system achieved seizure prediction, with an average prediction time of 0.8 seconds before clinical seizure onset. With different settings, an average latency of 13.8 seconds was noted, i.e. seizures were detected on average 13.8 seconds after onset [99].

Beyond ECG, which measures heart rate using electrodes, sometimes PPG (photoplethysmography) is used to measure heart rate, as it is a low-cost and non-invasive method that does not require electrodes [100]. Although various devices use heart rate variability as a seizure detection tool, most of which do not use it exclusively, but with a combination of other modalities such as ACM. *NightWatch* ® is a multimodal SDD based on heart rate and movement detection using PPG and ACM, which costs £1249 [7]. A paper published in *Neurology* found that this device had a sensitivity of 85% and a FAR of 0.23 per night [92].

Heart rate variability methods have some disadvantages, such as ECG electrodes possibly causing some discomfort or detaching, and large movements possibly affecting measurement, as well as physical activity or stress potentially affecting heart rate, which leads to false alarms [23].

Electrodermal activity (EDA) measures skin conductance changes due to sweating and has been associated with activation of the sympathetic nervous system [101]; it can be used (mostly as part of a multimodal device) to detect seizures, as one study found, achieving a sensitivity of 94% with a FAR of 0.74 per 24 hours with a method based on EDA and ACM [102].

Empatica ® has developed two EDA-based multimodal SDDs that are currently available, the *E4* ® and the *Embrace2* ®, both in a wristwatch form factor, and both using ACM and EDA. The *Embrace2* ® costs £249 and its alert app requires a monthly subscription that costs between £10 and £45 [8]. This device also includes a gyroscope and a temperature sensor. The *E4* ® costs $1,690 [9], and includes a PPG sensor to monitor heart rate and a temperature sensor. According to a study published in *Epilepsy Research*, the algorithm that these devices use, which is based primarily in EDA and ACM, has a sensitivity of over 92% and a FAR of 0.2-1 per day [103]. Although EDA can be susceptible to artifacts [23], the multimodal approach may enable an accuracy improvement over other devices with unimodal approaches. Moreover, these devices have the advantage of being unobtrusive and discreet, which are

[7] Epilepsy Alarms UK website: https://www.epilepsyalarms.co.uk/product/nightwatch/, last accessed on 24/10/2021

[8] Epilepsy Alarms UK website: https://www.epilepsyalarms.co.uk/product/embrace2-epilepsy-sensor/, last accessed on 25/10/2021

[9] Empatica website: https://store.empatica.com/products/e4-wristband?variant=17719039950953, last accessed on 25/10/2021

desired characteristics for patients and caregivers [21].

Table A.1 in Appendix A summarizes the above-mentioned commercially available SDDs, while table A.2, also in Appendix A, summarizes the above-mentioned seizure detection methods.

3.3 Video-based Seizure Detection

Video-based seizure detection has been studied as an unobtrusive alternative to existing seizure detection methods. As previously mentioned, patients with epilepsy and caregivers value unobtrusiveness and comfort [21]. While video-EEG is the gold standard for epilepsy diagnosis and seizure identification, EEG acquisition can be too obtrusive or uncomfortable for an ambulatory setting [23]. Detecting seizures using only video footage could not only solve these problems, but also provide valuable information towards seizure registration and classification as a part of the diagnosis and management process. Video-based seizure detection and monitoring devices for ambulatory use could also offer a direct remote monitoring option for caregivers, who emphasized the importance of video monitoring (see Section 3.1) [21, 25].

Seizure detection methods based on video fall into two major categories: marker-based and marker-less. Marker-based seizure detection methods use physical markers to highlight body parts, in order to exclude signals from other people in the video footage, and discriminate between motion from different limbs. One such method that was developed using markers that reflected infrared light managed to differentiate between hyperkinetic and non-hyperkinetic seizures, and has been shown to accurately detect hyperkinetic seizures with 80% probability [104]. Another method didn't require physical markers, but did require patients to wear pyjamas with specific colors. This method used these pyjamas in order to isolate limb movement, and extracted features from it to detect seizures. It achieved a sensitivity and specificity of 93% when using a "oscillation" feature, with a latency of 21.4 seconds. Latency could be lowered to 4s with the use of a "displacement" feature, at the expense of specificity, which was lowered to 53%. This study also found that a video/EEG hybrid system achieved higher accuracy, detecting 100% of seizures with the same specificity of 93% [105]. Although these methods are promising, marker-based approaches have the disadvantage of not being very practical, due to markers possibly being occluded by blankets, for example.

Various seizure detection algorithms using marker-less video footage have been published in the last 20 years. These methods tend to calculate motion features such as displacement, duration, angle, velocity and others [106]. Most of these approaches are based on the calculation of optical flow [107–109], which is a technique used to determine the apparent velocity field of a video segment [54], while other techniques use deep learning methods like Convolutional Neural Networks (CNN) [110, 111].

One of the most accurate nocturnal motor seizure detection methods obtained 100% sensitivity and a moderate FAR of 0.78 per night in a residential care setting [112]. This algorithm detected clonic movement by calculating global motion parameters from the optical flow data and classifying based on what they defined as "*spectral contrast*", a ratio between the spectral power in the 2-6 Hz frequency range and the power outside of this frequency range [107]. This algorithm was later also validated for

children with epilepsy, and obtained a lower sensitivity of 78% and a FAR of 0.05 per night [113]. A different algorithm for detecting neonatal seizures was developed by Karayiannis et al. and published in *Clinical Neurophysiology*. It was also based on optical flow calculation and defined features ("*energy ratio*", "*variance of time intervals*", "*maximum spike duration*", and "*10% spectral power frequency*") that were used to train a feed-forward neural network. This study reported a sensitivity and specificity between 82% and 95%, depending on the features used [108]. Another method used face detection and optical flow to detect absence seizures, with a reported accuracy of between 99.86% and 99.96% [109].

Other methods that employed deep learning techniques included a study published in 2016 and a very recent study published in August 2021 [110, 111]. The latter achieved much better results than the former, with a reported sensitivity of 88% and specificity of 92%, as well as an average latency of 22 seconds, with the former achieving an AUC of 78%. The 2021 study developed a seizure detector using a combination of CNNs (Convolutional Neural Networks) for analysing each individual frame and an RNN, namely a long short-term memory (LSTM) network to analyse the time-series of frames [111].

A separate paper developed a multimodal algorithm based on a mixture of video and EEG signals. Features were extracted from video using spatio-temporal interest points (STIP) and from EEG using the entropy and energy of different brain wave frequency bands (alpha, beta and theta). For classification, different linear and non-linear classifiers were used (Naive-Bayes, linear and non-linear SVM, K-Nearest Neighbours and AdaBoost). This method achieved a sensitivity of 90% and a specificity of 92% [114].

Although video-based seizure detection has been fairly extensively studied in recent years, and various different approaches have shown generally very positive levels of accuracy, as far as could be found, **there are no video-based seizure detection devices currently commercially available**. A review study of seizure detection devices for home use published in *Epilepsia* in 2020 did not find any such devices, stating that "automated seizure detection using video [...] is not discussed below, because to the best of our knowledge there are no such devices currently on the market." [20]. Other review studies have identified devices on the market that detect "abnormal" motion, particularly during sleep, and could be used for epilepsy monitoring, but that neither claim to detect seizures nor have any published literature detailing their accuracy. Examples of devices like these are *SAMi sleep activity monitor* ®, which simply detects abnormal movement, and Vicon's Vantage+ ®, which is merely a motion detecting camera system [23, 74].

Video-based seizure detection has many advantages, such as not being obtrusive or uncomfortable, and potentially being very affordable, as video-based SDD systems essentially only need a camera and a microprocessor, to process and transmit data, both of which can be fairly inexpensive. Although this detection modality is viewed positively by caregivers [21, 25], as it enables them to personally monitor their children, patients themselves can be reluctant to be subjected to permanent surveillance, possibly due to privacy concerns and a mistrust in how safely their data will be secured [18, 21]. Table 3.1 summarizes the above-mentioned seizure detection methods based on video.

Table 3.1: The current state of the art in marker-based and markerless video-based seizure detection methods; the seizure types they claim to detect, the general functioning of each algorithm, the reported detection accuracy and latency, and the published paper describing each method.

Modality	Seizure Type	Algorithm	Detection Accuracy	Latency	Article
Marker-based w/ IR reflecting markers	Automatisms	Thresholding of movement parameters	80% probability that seizure is hyperkinetic above a threshold	Not stated	Rémi et al., 2011 [104]
Marker-based w/ colored pyjamas	Motor/ Hypermotor Seizures	Limb isolation, movement features, class. w/ threshold	Sen:93.3%, Spe:53.3-93.3%, for different video features, Sen:100%, Spe:93.3% w/ EEG	4.6-27.8s	Lu et al., 2013 [105]
Marker-less	Major Motor Seizures (nocturnal)	Opt. flow, extraction of global motion params, spectral contrast	Sen: 100%, FAR: 0.78/24h Sen: 78%, FAR: 0.05/24h (in children)	Between 7 and 35s, $\leq$10s in 78% of seizures	Kalitzin et al, 2012 [107] Geertsema et al, 2018 [112] Westrhenen et al, 2020 [113]
Marker-less	Myoclonic, Focal Clonic (Neonatal)	Feature extraction + feed-forward neural net	Sensitivity: 86.5-94.9%, Specificity: 81.5-91.5%	Not stated	Karayiannis et al., 2006 [108]
Marker-less	Absence	Optical flow + face detection + background subtraction	Accuracy: 99.86-99.96%	Not stated	Pediaditis et al., 2012 [109]
Marker-less	Seizures with at least 1 clinical sign	Convolutional Neural Network using depth and IR footage	AUC: 78.3%	Not stated	Achilles et al., 2016 [110]
Marker-less	GTCS	CNN + Long Short-Term Memory Network	Sensitivity: 88%, Specificity: 92%	22s	Yang et al., 2021 [111]
Marker-less	Nocturnal Convulsive	Video feature extraction w/ STIP + EEG feature extraction + class.	Sensitivity: 89.6%, Specificity: 91.6%, PPV: 90.7%	Not stated	Aghaei et al., 2017 [114]

Chapter 4

Video Monitoring Device

With the interest of creating an architectural basis for future implementations of the seizure detection device envisioned in this work, a working prototype for Lampsy was designed and built. Initially, this prototype was intended to be used for **acquiring raw data in a clinical epilepsy monitoring unit**, with the future goal (outside the time-frame of the present work) of eventually testing the device in an ambulatory as well as a clinical setting.

This project was developed in the context of a larger cooperative project between Instituto Superior Técnico and Hospital de Santa Maria (HSM), for which there was already an agreement with the latter's ethics commission for data acquisition at its Epilepsy Monitoring Unit, prior to the start of this work. Therefore, with the intent of recording raw video footage during long-term video-EEG sessions, an addendum was submitted to the ethics commission, which was accepted on November the 5th, 2021. The prototype was therefore placed in HSM on November the 9th, 2021.

4.1 Hardware Base

Within the time-frame of this work, the main goal in the development of this prototype was to enable it to acquire raw data, i.e. video recordings from video-EEG sessions. Nevertheless, this device was intended to, in the future, incorporate a fully fledged seizure detection algorithm. As such, it was built around the Raspberry Pi (RPI) 4 Model B single-board computer, as it satisfied the requirements defined for the prototype's processing unit, namely (a) **having enough processing/computing power** to calculate optical flow, ICA and the other methods included in the developed seizure detection algorithm in real-time (see Chapter 5); (b) being able to **connect to an internet connection**, in this case via a 4G router, with an Ethernet cable, allowing a permanent and consistent connection, therefore not requiring set-up by hospital personnel; and (c) being able to **connect to an external display**, to allow the adjustment of the camera's position.

The RPI 4 Model B, shown in *Fig. 4.1a*, has a quad-core 1.5GHz processor and between 2Gb and 8Gb of memory (RAM), which enables it to comfortably run the necessary software [1]. Moreover, it has

[1] From the Raspberry Pi website: https://www.raspberrypi.com/products/raspberry-pi-4-model-b/specifications/, last accessed on 31/10/2021

Gigabit Ethernet, enabling it to connect to a 4G router, as well as micro-HDMI and micro-USB ports, enabling it to connect to and power an external display.

Besides, there were other additional advantages in choosing the RPI. First of all, it is possible to run Python programs in it, enabling the prototype to run the algorithm. Moreover, as it runs a version of the Linux operating system (in this case, Raspbian OS), it can be configured to start a program at launch using the system daemon (`systemd`), and to encrypt and automatically transmit data to the cloud using `rclone`. Linux is very versatile, which is a significant benefit in programming such a device. Moreover, the RPI has a very small form factor, which is essential, as this enables the device to be built into a light fixture. Additionally, the fact that there are a plethora of accessories available for the RPI, such as, crucially, infrared cameras, is very advantageous. Finally, the RPI has an array of General Purpose Input-Output (GPIO) pins, which enable the addition of custom electronic circuits that add extra functionality, such as for example a power button.

The RPI is also fairly inexpensive, depending on the model choice, and affordable component choice will be essential for the final version of Lampsy, in order maximize the number of patients who will be able to afford the device. The cheapest is the RPI Pico, which starts at $4[2], while the RPI 4 Model B starts at $35[3] and was acquired for 49.99€ [4]. The version of the RPI 4 Model B with 2Gb of memory was chosen for this initial prototype. However, although all versions of RPI are affordable, further analysis will be required in future work for determining whether a cheaper (and smaller) version of RPI will provide sufficient performance when devising the final version of Lampsy.

The second most important component was, naturally, the camera. A few criteria were important regarding the camera choice, such as being able to film video footage at an adjustable resolution and at a minimum of 25 frames per second (fps), and supporting night-vision (with infrared LEDs).

Regarding the first point, it was important for the camera to allow multiple custom resolutions so that when the optimal resolution for the algorithm was chosen, this could be easily adopted. The algorithm was built around a dataset of videos with a resolution of 384×288 pixels, so the camera needed to, at a minimum, support this resolution. A minimum frame rate of 25 fps was crucial so that movements at frequencies of up to 12.5 Hz could be sampled (according to the Nyquist-Shannon Sampling Theorem [115]). A lower frame rate would not allow the sampling of some clonic movements, as these can reach high frequencies, especially in the ”vibratory” phase of tonic-clonic seizures, during the transition from tonic to clonic movement [41].

Furthermore, it was essential that some form of night-vision was supported, in order to enable night-time monitoring. This is most often achieved by including infrared LEDs (which should preferably be attachable and powered by the camera module itself). Infrared light is not within the visible spectrum, and as such the light from these LEDs is captured by the camera but only a light glow can be seen by humans. The camera module also needed to either not have an IR-cut filter (which for regular cameras removes the infrared component of video footage), or to have an IR-cut filter that was turned off when the

[2]From the Raspberry Pi® website: https://www.raspberrypi.com/products/raspberry-pi-pico/, last accessed on 31/10/2021

[3]From the Raspberry Pi® website: https://www.raspberrypi.com/products/raspberry-pi-4-model-b/, last accessed on 31/10/2021

[4]From Mauser's website: https://mauser.pt/catalog/product_info.php?products_id=096-7401, last accessed on 14/11/2021

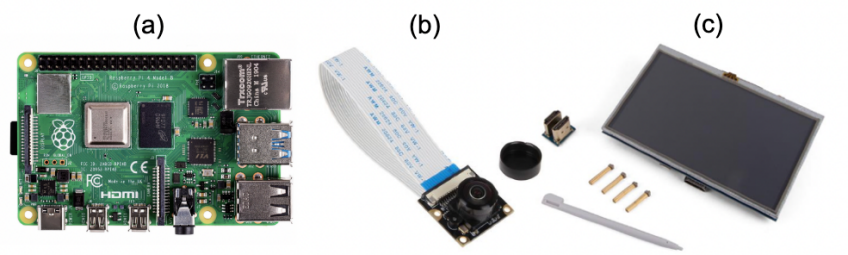

Figure 4.1: Main electronic components of the prototype: (a) – The Raspberry Pi® 4 model B was the chosen processing unit for the working prototype. (b) – ”5MP WIDE-ANGLE CAMERA 200° DEGREES” Raspberry Pi compatible camera model, produced by Joy-It®, pictured here without the attached infrared LEDs, that were included to enable night-vision capability. (c) – The ”VMP400 5” 800 x 480 HDMI-B Touchscreen for Raspberry Pi” display model, from Velleman®.

camera senses that light conditions are low. Following these criteria the ”5MP WIDE-ANGLE CAMERA 200° DEGREES” camera module produced by Joy-It® [5] (*Fig. 4.1b*) was chosen. This camera included infrared LEDs that, as mentioned, were essential to enable night-vision. This camera was acquired for 40.20€ [6].

The third and final main electronic component that was included was a small display, to facilitate the adjustment of the camera's position and viewing angle and obtain live feedback from the recording. Some requirements were defined for choosing the display, such as: (a) supporting a minimum resolution of 384 $\times$ 288; (b) having a small form factor; (c) being powered via a micro-USB or USB-C connection, to enable the RPI to power it directly; and (d) receiving signal via HDMI, to enable a direct connection to the RPI.

The chosen display was the ”VMP400 5” 800 x 480 HDMI-B Touchscreen for Raspberry Pi” model, produced by Velleman® [7], shown in *Fig. 4.1c*. This display contains a switch to turn off its backlight, which is also important for when the recording is occurring during the night, so that the light from the display does not bother the patient's sleep. The touchscreen functionality included in this display was not utilized, as this was only available if the display was connected via the GPIO pins in the Raspberry Pi, which was not ergonomic in terms of the overall structure of the prototype. This display was acquired for 65.55€ [8].

The total cost of acquiring these parts was 155.74€. It is important to mention that better parts were acquired than those that met the minimum requirements, in order to assure that there would be no technical problems.

4.2 Recording Methodology

The device needed to be able to record footage during extended time periods to be able to acquire video data from the long-lasting video-EEG sessions. Instead of only saving the footage locally, it was

[5] From Joy-It's ® website: https://joy-it.net/en/products/rb-camera-WW2, last accessed on 01/11/2021
[6] From Mauser's website: https://mauser.pt/catalog/product_info.php?cPath=1667_2620_2956&products_id=096-7651, last accessed on 14/11/2021
[7] From Velleman's® website: https://www.velleman.eu/products/view/?id=438240, last accessed on 01/11/2021
[8] From Mauser's website: https://mauser.pt/catalog/product_info.php?cPath=1667_2620_2959&products_id=096-5022, last accessed on 14/11/2021

determined that this device would need to securely transmit the data to a secure server located at Instituto de Telecomunicações at Instituto Superior Técnico, in order to ensure that it would be possible to access these files remotely without this requiring on-site access to the prototype, in the hospital (the details of the data transmission will be discussed in Section 4.3). This would also assure that the files were not lost. Therefore, the device needed an internet connection to transfer the files. As such, due to the extended nature of these recordings, the video files had to be kept at a reasonable size. This would ensure that the time it would take to upload the files would be short, which would reduce the probability of internet connection issues occurring during the upload, which could result in lost files or errors in the recording process. Beyond this, it was also important for file size to be kept small due to the storage constraints of the device itself, (though the RPI 4 has 64 Gb of storage, which would almost certainly suffice). A file size goal $<$ 100 megabytes per hour (Mb/h) was set.

Therefore, in order to ensure that the file size remained low, it was decided that the video footage would be **automatically stopped and immediately restarted at hourly intervals**. The resulting 1h long video files would be saved and uploaded individually. Furthermore, these video files (in the .h264 file format) would be converted into .mp4 files to enable playback in a computer after recording, and compressed into a .zip file, to reduce their size. However, as recording would need to be temporarily stopped every hour, it was paramount to make sure that it only stopped momentarily and immediately resumed. It was noticed that the larger the file size, the longer it took to restart the video, presumably because the file was being saved between recordings. Therefore, it was essential to adjust the resolution such that it would be high enough for the algorithm to work (at least 384$\times$288) but low enough that video files would not take up too much storage space and the recording would not be stopped for very long. This resolution also had needed to match the aspect ratio of the included display, which was 5:3 (800$\times$480). As a result, the chosen resolution was 400$\times$240 pixels.

Another issue was the need for **manual control** over the recording process, in order to enable the patient or health professional to stop the recording for privacy reasons, and later restart it. Because it was fundamental for this process to be simple and intuitive, a ”Record (**REC**)” button and a red LED were included. When the **REC** button is pressed, recording stops, the video footage since the start of recording (or since the last hourly stop) is saved and compressed into a .zip file, and the red LED is turned off, indicating that the recording is paused. As the recording stops, the live feed on the display is also stopped. When the button is pressed again, the LED also turns back on and the video recording immediately resumes, and will only restart an hour after that point, unless the recording is manually restarted again.

As an additional way of ensuring that the device’s functionality is intuitive, a power (**ON**/**OFF**) button was installed, along with a green LED to indicate that the device is on or off [9]. This was also important to guarantee that the RPI is safely powered off, as simply removing the power cable may lead to the corruption of data or even the SD card (which happened once, before this button was implemented). The two buttons and LEDs were connected to the RPI using its GPIO pins, as per the circuit diagram shown in *Fig. 4.2*.

[9] Tutorial that was followed, including cloning their repository: howchoo.com/g/mwnlytk3zmm/, last accessed on 15/11/2021

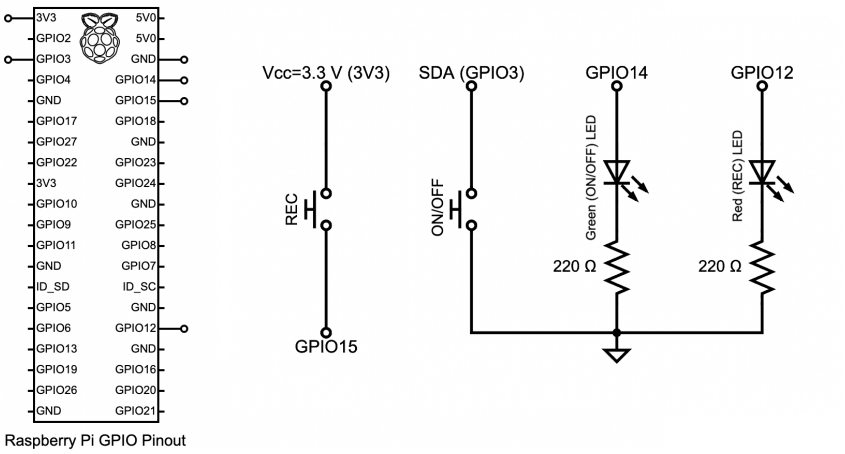

Figure 4.2: The circuit diagram for the external buttons and LEDs that were installed. On the left, the pinout for the RPI's GPIO pins can be observed, with the used pins highlighted (3V3, GPIO3, GPIO12, GPIO14, GPIO15, and GND). The circuit in the middle is the circuit for the recording (REC) button, which connects to GPIO15 for output, and 3V3 to receive an input voltage of 3.3 V. The circuit diagram on the right shows the power (ON/OFF) button and the two LEDs. The ON/OFF button is connected to GPIO3 and the ground (GND), and the two LEDs are connected to GPIO12 and GPIO14 for input and a 220 Ω resistor to control their current. These resistors are then connected to the ground (GND).

The REC button was connected to the 3V3 pin, which provides a 3.3 V voltage source, and GPIO15 for output. In the python script that was written for the operation of the recording process described above, GPIO15 was set as an input pin (as in input from the button), in pull-down mode, such that inputs would be detected when GPIO15 was set to "HIGH", i.e. when the button was pressed. This input would then control the recording as described above. When the video recording was on, a GPIO12 was set to "HIGH", which would turn on the red LED. This LED was connected in series with a 220 Ω resistor to control the current through it, such that it would not exceed its specifications.

The power button was connected to GPIO3 (SDA) and to the ground, as the Raspberry Pi can only be powered on via GPIO if GPIO3 is shorted to ground [10]. Therefore, the button was connected only to the GPIO3 and GND pins. The green LED was connected to GPIO14 for input, and to a 220 Ω resistor, such that, similarly to the red LED, the current through the LED would not exceed its specifications. This circuit was initially set up in a breadboard, and later assembled inside the official case for the RPI, so that no external wires or resistors were visible (only the buttons and LEDs were), as can be seen in *Fig. B.2* in Appendix B.

This recording methodology and buttons were, as mentioned, implemented in python. For this purpose, the `picamera` package was used for recording, the `zipfile` package was used for compressing files, the `RPi.GPIO` package was used to control the GPIO pins for the LEDs and buttons, and the `datetime` package was used to count the hourly intervals for restarting the recording. As another way to ensure that this device would function in a very intuitive manner, a `systemd` service was created to make this program run at boot-up, such that the recording was started when the RPI was booted up, without

[10] This was mentioned in the guide that was followed for the installation of this button, and confirmation from an official source was found in the documentation on Raspberry Pi's website: https://www.raspberrypi.com/documentation/computers/raspberry-pi.html#gpio-boot-mode

the need for any input from health professionals [11].

4.3 Secure Transmission and Synchronization of Data

The chosen protocol for transmission was the Secure Shell File Transfer Protocol (SFTP). Secure Shell (SSH) is a network security protocol that enables secure communication over a network by accurately authenticating the user and automatically encrypting any data that the user sends and decrypting it at the receiving side. SFTP, in turn, is a file transfer protocol based on SSH that provides an encrypted (and therefore secure) channel for transferring information over a network without the need for manual encryption [116].

Using STFP, the device could simply connect to a network and transfer the video recordings to a remote server through the encrypted channel that it provides. This network connection was achieved by connecting the device to a 4G router with an unlimited mobile data SIM card, using an Ethernet cable. The device is therefore always connected to the internet. A network attached storage (NAS) server at the Pattern and Image Analysis (PIA) lab at IT called PIANAS was chosen as a secure location for storing the data. A folder was created at this server, to which only the author of this work and the supervisors had access. In order to establish this connection between the device and the server using SFTP, the `rclone` command line program was used. This open source program primarily enables the transfer, backup, encryption and synchronization of files with cloud storage services such as Dropbox and Google Drive and many others, and can also be used to establish an SFTP connection with a server using its IP address and port number.

Using `rclone`, the folder in the RPI in which the .zip files of the video recordings were saved was connected via SFTP with the aforementioned folder in PIANAS. `Cron`, a command line service used for task scheduling, was used to synchronize the folder in PIANAS with the folder in the RPI in hourly intervals, thereby uploading any video files that had been recorded in that time through the secure SFTP channel that was created with `rclone`.

4.4 Prototype Construction

It was important to provide a fine adjustment of the camera's position. This would enable health professionals to more easily make sure that the patient's bed was well framed in the recording. For this purpose, a custom 3D model for a small camera support was designed in the 3D modelling software OpenSCAD, as can be observed in *Fig. 4.3*. This support consists of two separate parts, a camera holder and a base, which was glued to the RPI's case.

The camera holder connects to the base using a separate screw and nut, which when tightened, pull the two arms of the base together, securing the holder in place using friction. Because its held in place due to friction, it can be very easily adjusted by hand when the screw and nut have been previously

[11] The following guide was followed: https://www.dexterindustries.com/howto/run-a-program-on-your-raspberry-pi-at-startup/#systemd, last accessed on 03/11/2021

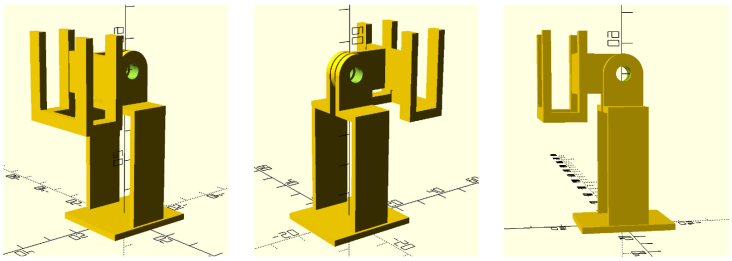

Figure 4.3: The 3D model for the camera support, designed with the 3D modelling software OpenSCAD. The two parts are connected using a separate screw and nut, which when tightened, pull the two arms of the base together, securing the holder in place using friction, enabling it to be easily adjusted, and ensuring it stays in position after it is adjusted.

Figure 4.4: The inside of the lampshade of the chosen light fixture, in which the prototype was encased. A metal support was built and screwed to the inside of the lampshade, to accommodate the RPI and keep it and the camera level.

tightened the correct amount, and stays in position after it is adjusted. The finished and 3D printed part, glued on to the official RPI case (with the integrated buttons and LEDs), can be observed in *Fig. B.1*, in Appendix B.

In order to better convey the Lampsy concept with this prototype, it was integrated within a light fixture, namely a black floor lamp with a metal lampshade that could also be adjusted to control the camera's position. The RPI, inside its case, and the camera, were placed inside the lampshade in place of the lightbulb, as can be seen in *Fig. 4.4*. A metal support was built and screwed to the inside of the lampshade, to accommodate the RPI and keep it and the camera level.

In order to feed the power, HDMI, micro-USB and Ethernet cables into the lampshade, two large holes were cut into it, one from the back for the power and HDMI cables and one from the side for the Ethernet cable and the micro-USB cable used to power the display. A support to hold the 4G router was built and attached to the lamp pole, as seen in *Fig. B.3a* in Appendix B. Both the RPI and the 4G router require power, and as such must be connected to a power outlet.

A 3D model was also designed to hold the display, as is shown in *Fig. 4.5*. This model is essentially a box with a lid. The box includes a front opening for the display, a side opening for the backlight switch, and top openings for the HDMI and micro-USB power cables. The lid slides into the box, and contains a hole for the HDMI cable to come out of. This HDMI cable then goes into the hole on the back of the lampshade, along with the power button.

The 3D printed display encasing can be observed in *Fig. B.3b* in Appendix B. The display is attached

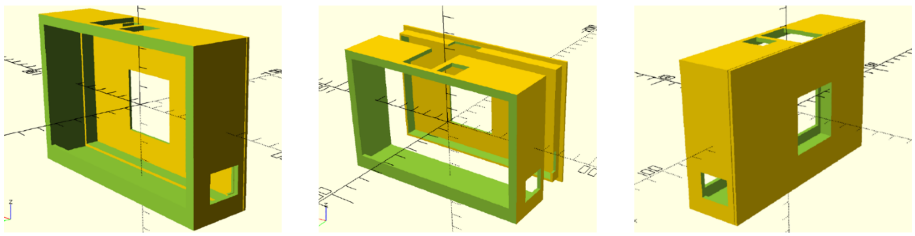

Figure 4.5: The 3D model for the display encasing, designed with the 3D modelling software OpenSCAD. The box includes a front opening for the display, a side opening for the backlight switch, and top openings for the HDMI and micro-USB power cables.

Figure 4.6: The final build of the prototype for Lampsy. The RPI was placed inside its official case and buttons and LEDs were installed in it. A support for the camera was 3D printed and glued to the case. This was then placed inside a floor lamp, to which a display was attached, inside a custom 3D printed encasing.

to the lampshade using a simple phone holder used to attach mobile phones to a car's air conditioning vent. This also holds together the lid and box of the display encasing. *Fig. 4.6* in Appendix B shows more images of the finished prototype.

This prototype was, as mentioned, eventually placed in the epilepsy monitoring unit of HSM for video acquisition during a long-term video-EEG session, on November the 9th, 2021. As of November 15th, 2021 at midnight, the device had been recording video for 130 consecutive hours without any issues. A file size per hour of 50.9 Mb/h was observed, which was lower than the maximum file size goal of 100 Mb/h.

Chapter 5

Seizure Detection Algorithm

The focal component of this work was the development of an algorithm that could accurately detect and segment seizures, as well as deliver timely alarms. First, an exploratory analysis of the data was performed. This was followed by the development of noise reduction methods, after which a seizure detection and segmentation algorithm based on Optical Flow (OF), Independent Component Analysis (ICA), novel metrics and Machine Learning (ML) classification was devised.

Over the course of the development of this algorithm, two programming languages were employed. Python 3.9 was used for calculating OF, due to the capabilities of the OpenCV library, which contains a variety of pre-built computer vision functions, including functions for calculating Dense OF using the Farnebäck algorithm (see Section 2.2.1).

The pre-processing and noise reduction, exploratory analysis of the data, and the initial version of the segmentation algorithm were performed in MATLAB R2019B. This choice hinged on the fact that MATLAB's large library of toolboxes and its plethora of built-in functions enabled a thorough exploration of the obtained signals and their features.

However, since MATLAB is not directly compatible with Raspberry Pi, once the algorithm was complete, the MATLAB code was ported and translated to Python. The Spyder IDE was used.

5.1 Datasets

Two datasets were obtained as a part of a collaboration with Hospital de Santa Maria (HSM) and Hospital Egas Moniz (HEM): a training set and a test or validation set.

The training set consisted of 7 videos containing one tonic-clonic seizure each from 3 epilepsy patients and 10 videos from 6 patients without these seizures. One of the 7 videos with tonic-clonic seizures contained two seizures instead of only one, but this second seizure was excluded from this analysis because the patient was fully restrained, and as such there were no seizure movements to detect. These videos were 2-8 minutes long. Most of them contained seizures that lasted between 20 seconds and 2 minutes. Health professionals at HSM annotated the start of every seizure, and the end was manually annotated by the author.

Other videos were also provided with myoclonic or focal seizures with automatisms. The intent of this work was to produce a detector specifically for tonic-clonic seizures, so videos without tonic-clonic seizures were not used as part of the training set, except as examples of videos without these seizures. For training this model and adjusting its parameters, 6 videos with normal, every-day movements were added, in order to help the model better discriminate between seizures and non-seizure movements. Most videos were in 384 $\times$ 288 resolution, in 4:3 aspect ratio and in 25 frames per second.

The validation (or test) set consisted of 21 videos from HEM, all of which contained 1 tonic-clonic seizure, but also significant time periods without seizures, as these videos were between 5 and 30 minutes long. The beginning and ending of each seizure was annotated with the help of Dr. Francisca Sá of HEM and video-EEG technician Octávia Brás. Resolutions were variable, with some videos being in 360 $\times$ 288 resolution while others were in 720 $\times$ 576. The latter were resized to 360 $\times$ 288, in order to be consistent with the resolution of the videos in the training set.

In this chapter, in order to make comprehension easier, the same seizure video was used for most examples and plots, namely patient 101's third seizure, henceforth referred to as *p101_s3*.

5.2 Optical Flow Calculation

The choice to calculate OF for the elaboration of this algorithm was based on the following reasons: (a) It one of the most popular and established methods for detecting and tracking movement in video sequences; and (b) It provides information on the spatial location, direction, frequency, and magnitude of movements, as opposed to more basic methods like calculating frame differences.

Dense OF was calculated with the aforementioned (in Chapter 2.2.1) Farnebäck algorithm [58], an established and computationally efficient method for calculating dense Optical Flow. This method was chosen because it enables accurate OF calculation with only two frames, enabling it to be run in real-time. This was essential in order to assure that OF was calculated between frames, such that it would not be necessary to compute OF after every 45 second period in which the video was divided (see Section 5.6), as this would be very computationally demanding and would significantly worsen latency. **This method was tested on the RPI, and it ran in real-time without issues**, indicating that it is not too computationally demanding for this device. Moreover, it has the advantage of being adequate for calculating very fast movement, such as what is required in this work, as it does not employ spatiotemporal smoothing. Nevertheless, because of this, this method can be somewhat noisy, which will be addressed in Section 5.4.2 [58].

OF was computed using a modified script from OpenCV [1], which generated **optical flow movement vectors (OF vectors)** for equidistant points in the video frame. These were split into their horizontal and vertical components and exported as arrays into Comma Separated Value (.csv) files. For all available videos, according to their resolution, the "step" variable in the OF algorithm, which represents the size of the pixel neighborhood around which the vectors are calculated, was set such that around 200 OF vectors were obtained. For most videos, which were in 384 $\times$ 288 resolution, a step value of 24 was

[1] From OpenCV's github: https://github.com/opencv/opencv/blob/4.x/samples/python/opt_flow.py, last accessed on 14/11/2021

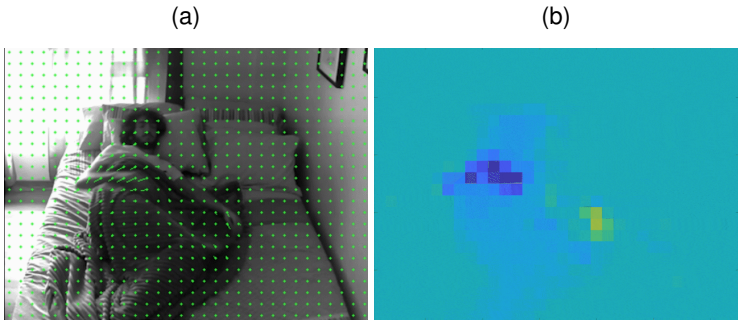

Figure 5.1: (a) – Screenshot of OF vectors obtained from a test video sequence featuring the author (b) – Still frame from the animation generated in MATLAB from the X components of the OF vectors obtained from the test video sequence in (a).

used, resulting in 16 by 12 vectors (192 vectors in total). In *Fig. 5.1*, the application of this technique to a video of the author can be seen, as well as screenshots from animation generated from these vectors in MATLAB.

The parameters for the OF function were kept as they were set in the existing implementation made by OpenCV, as they already struck a good balance between accuracy and noise levels. For instance, in this implementation, the `winsize` parameter (which controls window size) was set to a larger value than the default, which enabled the algorithm to be slightly more robust to noise while not causing blurring.

5.3 Exploratory Data Analysis

To facilitate this analysis, real valued time series signals were obtained from the OF vectors, by splitting them into their horizontal and vertical components (henceforth referred to as X and Y components). By separating the X and Y components, it was possible to obtain direction and frequency information, as well as detect periodical oscillations, which would not be possible if only the magnitude or absolute value of the vectors was considered. The X and Y components were stored in separate 3-dimensional matrices – a matrix for each frame of video. This matrix preserves the spatial locations of the OF vectors in the video frame as well as their time series. Furthermore, in addition to a spatial position in the frame, each OF vector can also be plotted in function of time, generating time series signals for its X and Y components.

Roughly 400 time series signals were obtained for each seizure video, 200 for the X components and 200 for the Y components. The time series signals of all OF vector X and Y components during this seizure period are plotted in *Fig. 5.2a* and *Fig. 5.2b* for *p101_s3*. Large spikes are noticeable, as well as some periodic oscillations towards the end of these signals. These spike "artifacts" are in fact caused by the movement of people (doctors and nurses, in this case) in front of the camera, causing large amplitude movements that often partially occlude the camera's field of view. This is one of the main problems with these signals, and it was solved by separating these movements from the seizure movements. This will be explored further ahead in this chapter.

To obtain basic frequency information from these signals, their Fourier Transform was computed

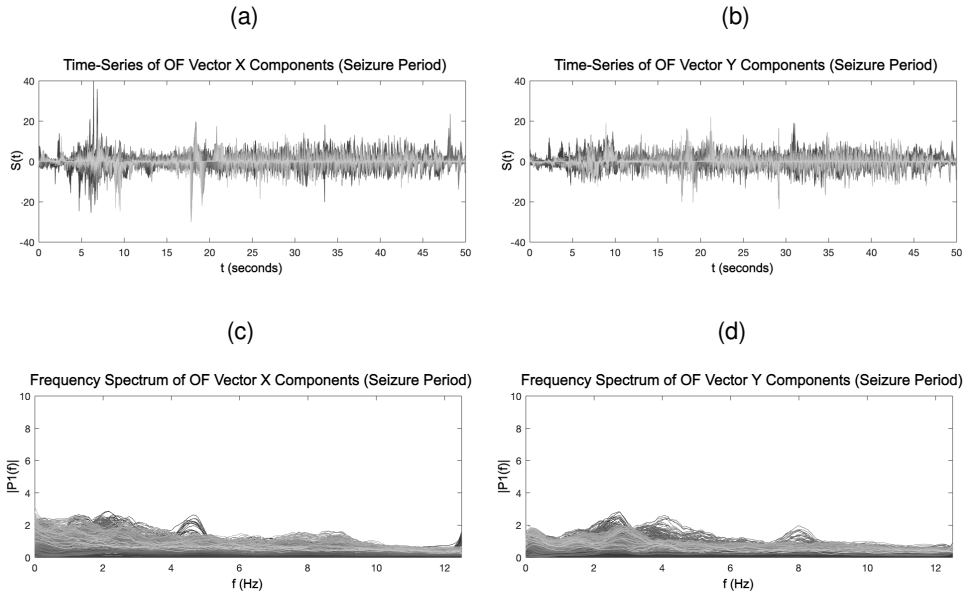

Figure 5.2: (a) – Time series of X OF vectors during the seizure period, each plotted in a different gray tone to facilitate comprehension. Positive signals indicate left-right movement. (b) – Time series of Y OF vectors during the seizure period. Positive signals indicate upward movement. (c) Frequency spectrum of X OF vectors during the seizure period, obtained using Fast Fourier Transform (FFT) and smoothed using an FIR Savinsky-Golay filter with a polynomial order of 5 and a frame length of 101. (c) Frequency spectrum of Y OF vectors during the seizure period, obtained using FFT and smoothed using an FIR Savinsky-Golay filter.

and is plotted in *Fig. 5.2c* and *Fig. 5.2d*. There seems to be significant activity in the 1-6 Hz range. However, the large number of signals and the presence of artifacts does not enable a thorough and accurate analysis. Therefore, a variety of processing steps were performed on this data. Initially, some pre-processing was performed, which included the development of a noise removal method based on a calibration period, a selection of a region-of-interest (ROI) around the patient's bed, and a low-pass filter. This denoising was then visualized in different plots to confirm its effectiveness. Afterwards, to separate the artifactual components and reduce dimensionality, Principal Component Analysis (PCA) and Independent Component Analysis (ICA) were performed. As multiple independent sources were obtained from this process, there was a source selection process using novel seizure metrics, and the selected sources were combined into one signal, to enable a much more thorough analysis. These steps will be explained in full throughout this chapter.

5.4 Pre-Processing

5.4.1 Noise Characterization

To examine the characteristics of the noise present in these signals, a 5 to 10 second period of each video segment in which there was no movement was selected. As these segments do not contain any actual movement, the only signal present is noise. In *Fig. 5.3a* and *Fig. 5.3b*, the time series of such a segment can be observed. When comparing this noise with the signal in *Fig. 5.2a* and *Fig. 5.2b*, it

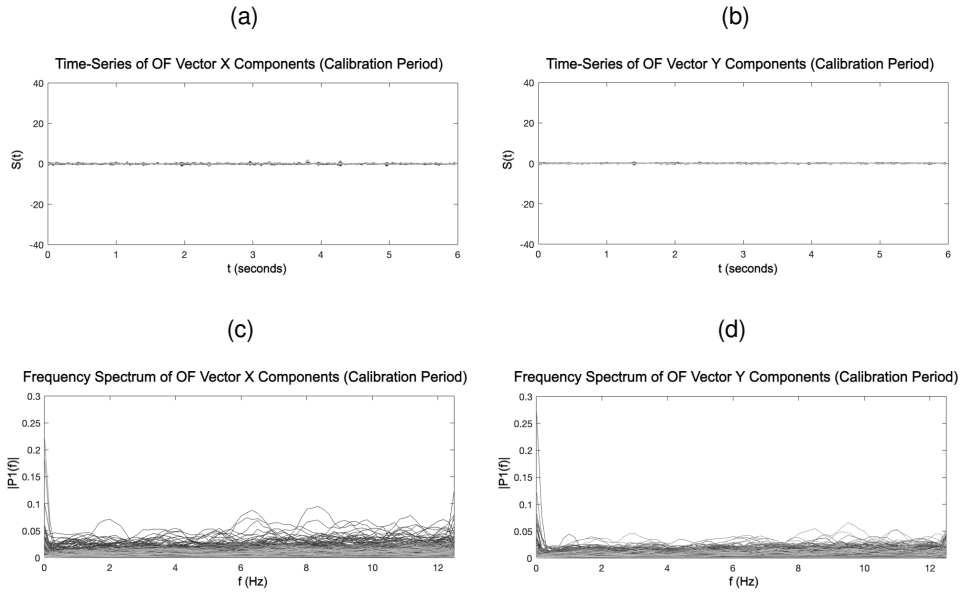

Figure 5.3: (a) – Time series of X OF vectors during the calibration period, each plotted in a different gray tone to facilitate comprehension. Positive signals indicate left-right movement. (b) – Time series of Y OF vectors during the calibration period, each plotted in a different gray tone to facilitate comprehension. Positive signals indicate upward movement. (c) Frequency spectrum of the X OF vectors during the calibration period, obtained using Fast Fourier Transform (FFT) and smoothed using an FIR Savinsky-Golay filter with a polynomial order of 5 and a frame length of 15. (c) Frequency spectrum of the Y OF vectors during the calibration period, obtained using FFT and smoothed using an FIR Savinsky-Golay filter.

becomes apparent one of the key characteristics of the noise is its **low amplitude**.

As can be observed in *Fig. 5.3c* and *Fig. 5.3d*, which shows the Fourier Transform for all OF vectors during this calibration period, the frequency content of the noise is fairly equally distributed within the frequency spectrum for most OF vectors. It is also apparent that some signals have much higher power. Additionally, there are spikes at 0 Hz and 12.5 Hz. The 0 Hz spike is the DC term, which is the mean of the signal, and is not particularly relevant for this analysis. The 12.5 Hz spike only occurs in some isolated OF vectors on the left margin of the frame. It seems to be caused by this particular video having very slight black bars on the sides, which when analysed by the OF cause problems, possibly due to the *aperture problem* of OF. This problem is explained in Section 2.2.1 and will be explored further ahead in this Section, but in general terms, when OF is computed locally, clear discontinuities in brightness can generate problems. One can look at the distribution of these noise vectors to attempt to ascertain their origin. Video and image noise tends to be mostly Additive White Gaussian Noise, which follows a normal distribution and is equally distributed within the frequency spectrum [117].

The distribution of the noise of all vectors for both X and Y components can be observed in *Fig. 5.4*. As expected, the majority of the noise is concentrated around 0, as there is, in fact, no movement. However, there are a large number of outliers, ranging from roughly -2.7 to +3.0. In comparison, signals from the seizure period (i.e. the period in the video sequence in which a seizure occurred) ranged from -40 to +40, so even these outliers have quite a low amplitude.

A way of quantifying this is by calculating the distribution's kurtosis. Kurtosis is a measure of the size

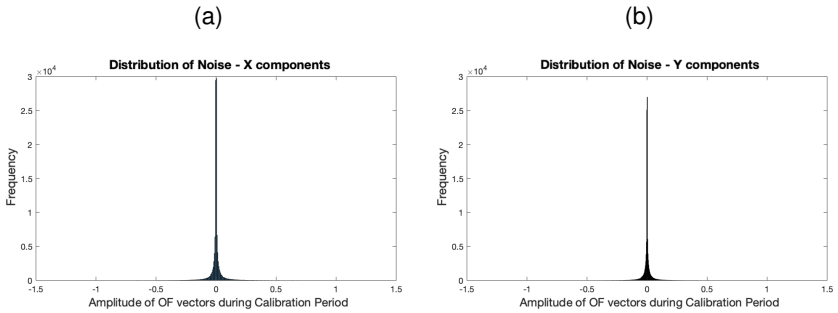

Figure 5.4: (a) – Histogram of the distribution of the X components of all OF vector signals during the calibration period. (b) – Histogram of the distribution of the Y components of all OF vector signals during the calibration period.

of the "tails" of a probability distribution, or, more precisely, according to Peter H. Westfall [118], it is a measure of a probability distribution's "propensity to produce outliers". A distribution is considered to be leptokurtic when its kurtosis is larger than 3, i.e. when its excess kurtosis – its kurtosis minus 3 – is positive. In this case, for *p101_s3*, (with step=12, so that we obtain 768 vectors), the excess kurtosis of the noise's distribution is 65.8 for the X components and 38.1 for the Y components. This very high kurtosis indicates an extraordinary propensity towards producing outliers.

A better way to show the distribution of the individual vectors and determine if there are patterns is to use a boxplot. Boxplots, often also referred to as box-and-whisker plots, are used to visualize statistical distributions. The boxes show the interquartile range between the upper and lower quartiles, and the whiskers show the variability outside of these. Outliers are also generally plotted. A boxplot (for step=12 and 768 vectors) showing each vector separately, is presented in *Fig. 5.5*, and shows that almost every vector in this noise period has a median of 0, but for some there is clearly higher variability.

A zoomed in section of this plot in *Fig. C.2* in Appendix C enables a clearer view of this pattern. For some OF vectors, the boxes and whiskers are orders of magnitude larger than others. There are also some clear outliers in some vectors. This pattern of vectors with substantially larger variability seems to be localized – vectors close to each other seem to have similar patterns – which hints at a specific spatial distribution of this behaviour. To determine if the spatial location of the vectors in the video frame was related to this divergence, the standard deviation of each vector was plotted in *Fig. 5.6*.

It becomes clear that this noise is not random but localized in specific areas of the image. It is also apparent that some areas contain noise in both X and Y components, but others are specific to either X or Y. The author proposes an explanation for this, namely that this noise is caused by the fact that OF, to address the brightness constancy constraint, often requires brightness to be differentiable, meaning it should vary smoothly across surfaces. This generates a problem at boundaries or edges where an object is occluding another or there is a significant difference in brightness. Moreover, the **aperture problem** complicates the estimation of motion at edges in the direction parallel to the edge, when estimating flow locally [54] (see Section 2.2.1).

In *Fig. 5.7*, one can observe a still frame of the video in question, *p101_s3*, with the patient blurred out to preserve their privacy, superimposed with a quiver plot of all OF vectors during this calibration period.

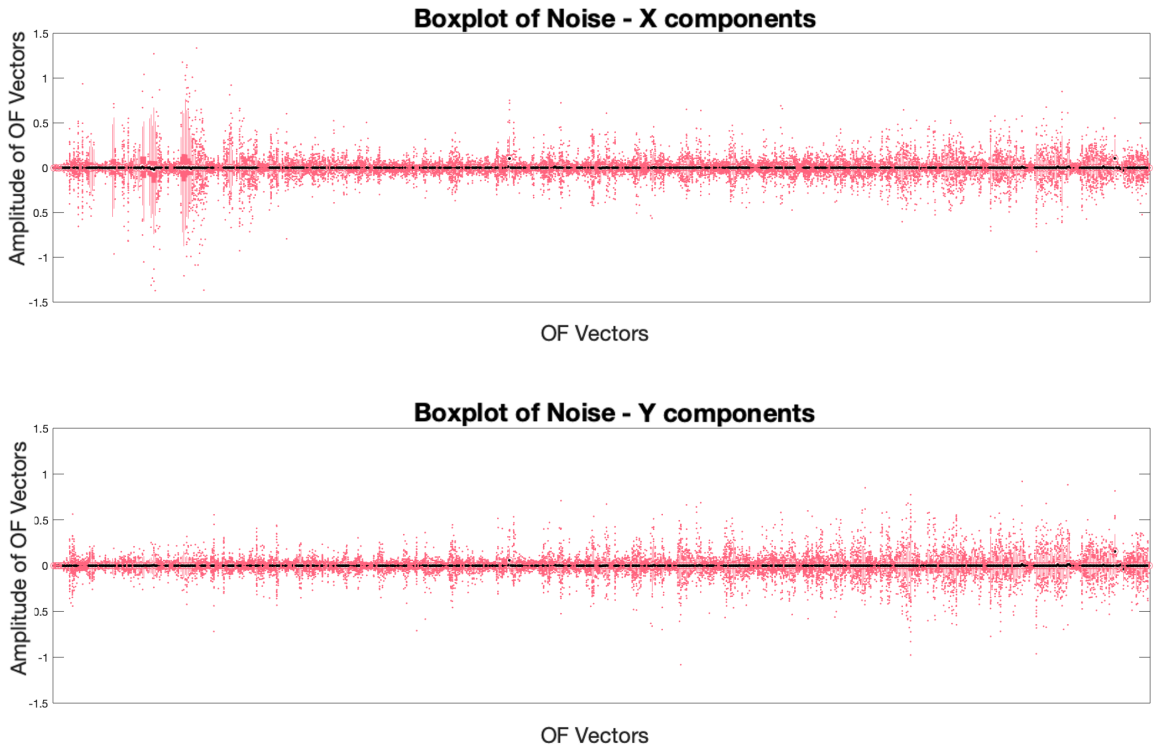

Figure 5.5: Boxplot showing the distribution of the X and Y components of all OF vectors individually. The black dots in the middle denote the median, the boxes denote the interquartile range and the whiskers represent the variability outside of it. The whisker length is specified as 3 times the interquartile range. Outliers are plotted as dots if they fall outside this range.

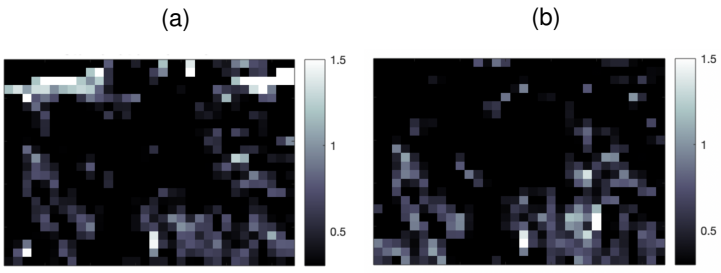

Figure 5.6: (a) – Standard deviation of the distribution of the X components of every OF vector, plotted in the spatial organization in the video footage. (b) – Standard deviation of the distribution of the Y components of every OF vector, plotted in the spatial organization in the video footage.

It becomes clear that the most prominent vectors in this calibration period are present near boundaries where there is a clear discontinuity in brightness, such as on the side of the bed or on the wall, and appear to be oriented in a direction parallel to the edges present in the video. This confirms that this noise is most likely not caused by ordinary image noise, but by the the aperture problem and by how OF addresses the brightness constancy constraint, requiring brightness patterns to be differentiable. Moreover, because the OF method used in this work was only based on two consecutive frames, it did not employ spatiotemporal smoothing to improve accuracy, which is a possible reason for the existence of this noise.

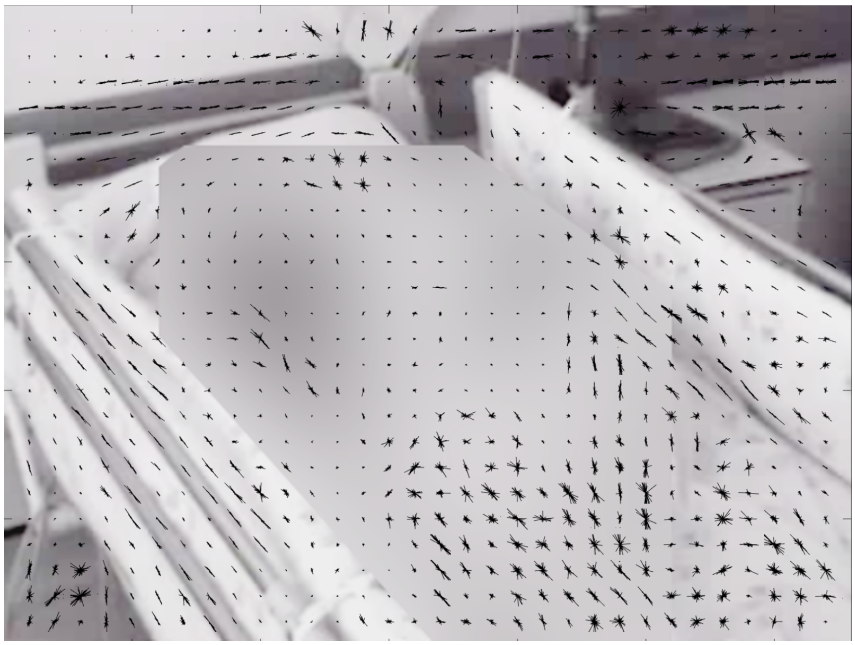

Figure 5.7: Still image of the video sequence for *p101_s3*, with the patient blurred out to preserve their privacy, superimposed with a quiver plot of all OF vectors during the calibration period, showing these vectors as lines, with their amplitude dramatically increased, such that their direction and relative magnitude becomes visible.

5.4.2 Noise Reduction

As it was previously mentioned, this noise is fairly **low amplitude**, so it should not present a very large confound to the data, and the signal to noise ratio (SNR) remains high. However, it is still important to implement noise reduction, as any improvements to the SNR will improve the quality of the data, even if this improvement is largely marginal. Moreover, ICA may isolate noise as an independent component, which could decrease specificity.

The main noise reduction method that was implemented was a combination of fixed and adaptive thresholds. The adaptive threshold is based on the standard deviation of each OF vector during the noise period, such that when its signal is lower than its threshold, it is nulled. Several factors justified this approach, namely: (a) seizure movements are higher amplitude when compared to this noise; (b) the amplitude, variability and standard deviation of this noise is far larger in some OF vectors; (c) most of the OF vectors with more significant noise components are not in the same location in the frame as the patient, as they are in the clearest boundaries, which will often be in the patient's bed or in their surroundings; and (d) this noise will be consistent throughout the video (if the camera is not moved, which it is assumed to be the case). This threshold was combined with another threshold with a fixed value for all OF vectors. The intent of this was to remove unwanted signals with low amplitude, namely the movement of the patient's bed or the slight jitter of the camera, which are not OF noise, as they represent actual movement, but are not part of the desired signal (the movements of the epilepsy patient), and as such can be nulled to increase SNR.

The adaptive threshold was initially set at a temporary value of 8 $\times$ the standard deviation, while the fixed threshold was set at 0.7. These temporary values were set such that normal or even slight

movements would clear them, and they would only remove noise and not signal. Later, once the algorithm was finished, these thresholds were calibrated according to what would maximize sensitivity and specificity, while reducing segmentation error and latency. This will be further explored in Chapter 6.

The combined threshold can therefore be defined as follows, with $T_{noise}(v)$ representing the combined noise reduction threshold for the X or Y components of the OF vector $x(t)$. There are, of course, separate thresholds for the X and Y components of each vector, as OF vectors have different standard deviation σ for different orientations.

$$T_{noise}(v) \begin{cases} 8 \times \sigma(v), & \text{if } 8 \times \sigma(v) > 0.7 \\ 0.7, & \text{otherwise} \end{cases}$$

This combined threshold superimposed with the previously shown boxplot is visible in *Fig. C.3* in Appendix C, and in *Fig. 5.8* one can observe the thresholds' spatial organization in the frame.

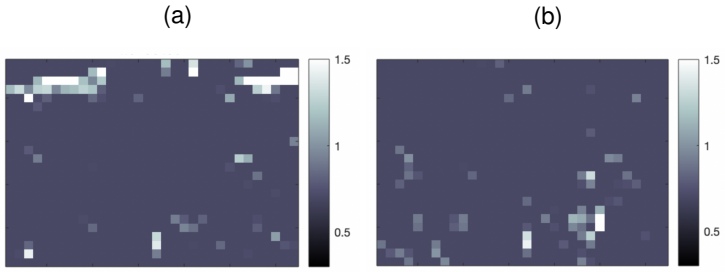

Figure 5.8: (a) – Threshold value for the X components of all OF vectors, plotted in the spatial organization in the video footage. (b) – Threshold value for the X components of all OF vectors, plotted in the spatial organization in the video footage.

Another important step in the noise reduction pipeline was to remove the 12.5 Hz spike in the frequency spectrum. This spike is not within the frequency range of normal or seizure movements, so it can be safely removed with a low-pass filter without affecting the desired signals. A digital low-pass Butterworth filter with a cutoff frequency of 11.5 Hz and order 3 was implemented.

Furthermore, another important question is the discrimination between the movements made by the patient and those made by other people in the frame, such as for example health professionals in a hospital setting or the patient's family in an ambulatory setting. The simplest way to remove a large amount of these movements is to isolate the patient's bed and general range of movement within the frame by selecting a Region-Of-Interest (ROI), whereby all OF vectors outside it are discarded. This step is important as it isolates many of the seizure movements from other movements.

This was implemented with a very simple ROI selection tool such as MATLAB's `roipoly` function (and later a Python equivalent - as shown in *Fig. 5.9c*), thereby generating a binary mask, which can be seen in *Fig. 5.9a*. In *Fig. 5.9b*, this mask is superimposed with a still image from the video sequence.

In the Chapter 6, the impact of these noise reduction steps is analyzed.

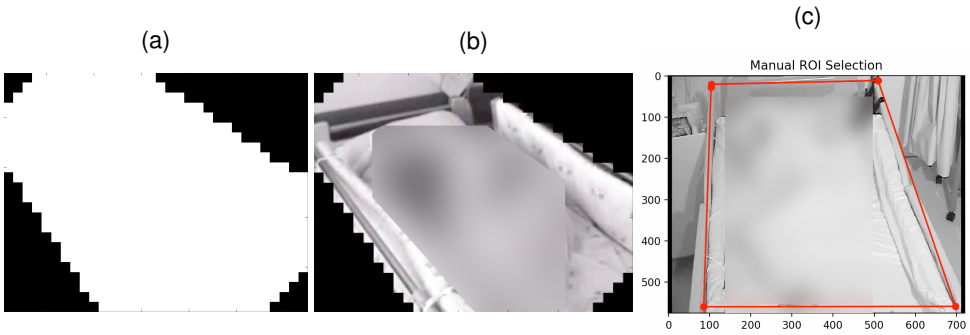

Figure 5.9: (a) – Binary mask generated from the ROI selection tool roipoly such that the patient's bed and general range of motion is isolated from other movement in the frame. (b) – Binary mask in (a) superimposed with still image of the *p101_s3* video sequence, with the patient blurred out to preserve privacy. (c) Manual ROI selection tool, implemented with MATLAB `roipoly` and then (pictured here) with a Python package of the same name that implements the MATLAB `roipoly` function.

5.5 Isolation of Seizure Movement with PCA and ICA

The main challenge in this work was to distinguish the patient's tonic-clonic seizure activity from all other movements. Its trivial to simply detect the existence of motion; however, classifying this motion as seizure movement or non-seizure movement is more difficult.

As previously mentioned, for a step size of 24, the data was comprised of the time series of 192 OF vectors for both their horizontal and vertical orientations. The information for seizure and non-seizure movements is encoded within these 400 signals (approx.), but many of these signals contain both types of movements, since both the patient and other people in the frame move in and out of the pixel neighbourhoods where each OF vector is detecting movement. However, a crucial piece of information we have about these movements is that they are independent from each other. Naturally, signals derived from movements by the patient and other people should be independent from each other, as they are performed by different people. Moreover, signals derived from movements performed by the patient during a seizure episode should be independent from other non-seizure movements. Furthermore, the data contains more OF vectors than there are sources of movement. Even if the patient's arms and legs, for example, are considered as different sources of movement, as well as those from other people in the frame, there certainly aren't 200 independent sources of movement. It can be safely considered that there are many more data points (or observations) than there are independent sources of movement.

Therefore, it makes sense to use dimensionality reduction methods such as Principal Component Analysis (PCA) to reduce the large number of data points while maintaining the maximum amount of the variability and information in the data, followed by a source separation method, such as Independent Component Analysis (ICA), that will identify the independent sources of movement and separate them. One of the main steps in the seizure detection algorithm that was developed in this work was therefore the extraction of **one single signal** (more accurately, one such signal for both the horizontal and vertical components of the OF vectors) that was a sum of the different sources of seizure movement, and as such **described the seizure activity in a video segment**.

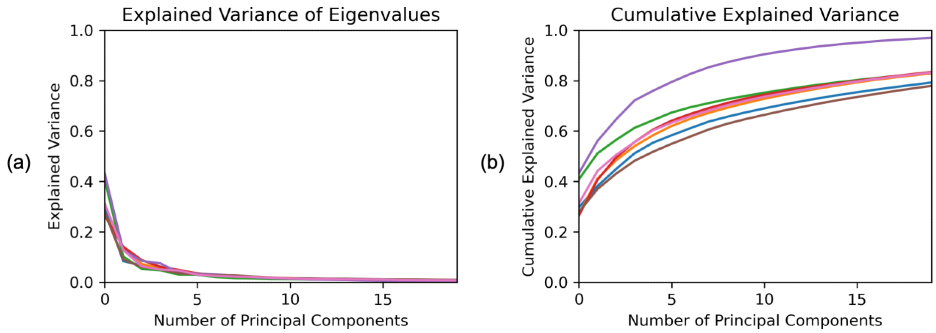

Figure 5.10: (a) – Eigenvalues from the first 20 PCs generated by performing PCA on videos in the training set, sorted by explained variance. PC #20 only accounts for 0.8% of the total variance of the signals. (b) – Cumulative explained variance of the first 20 PCs generated by performing PCA on videos in the training set. The cumulative variance explained by the first 20 PCs is 84%.

However, ICA can only identify the independent sources and cannot separate the ones that stem from seizure activity from those caused by other movements. For that purpose, seizure metrics were developed in this work. These metrics were used to differentiate seizure sources from non-seizure sources, and were based on an analysis of what constitutes a seizure signal and what differentiates it from the other sources typically detected by ICA in this setting. The following sections concern the implementation of PCA and ICA, as well as the source selection method and the development of the aforementioned seizure metrics.

5.5.1 Dimensionality Reduction

The use of PCA as a pre-processing step before ICA is a fairly common step. For example, a study published in the Computer Vision and Image Understanding journal on face recognition using PCA and ICA found that "pre-applying PCA enhances ICA performance by (1) discarding small trailing eigenvalues before whitening and (2) reducing computational complexity by minimizing pair-wise dependencies. PCA decorrelates the input data; the remaining higher-order dependencies are separated by ICA." [62]. Therefore, PCA was implemented, generating principal components (PCs), each of which explained a certain percentage of the variability in the signals. A set number of PCs that explained the most variability was kept and used to perform ICA, while the rest were discarded.

Before applying ICA to datasets with as many signals as this one, it is important to select the number of independent sources *a priori*. If this number of sources was not set beforehand, ICA generated a very large number of sources. Additionally, it was found that the quality of the sources was vastly better when the number of PCs was the same as the number of independent components (ICs). Otherwise, ICA generated sources that were all very similar to each other, which would not allow for proper identification of seizure movements.

Fig. 5.10a shows the eigenvalues from the first PCs generated by performing PCA on the seizure videos in the training set, sorted by the variance of the original signal they explain. *Fig. 5.10b* shows the cumulative variance explained by the first PCs, up to PC #20. On average, the first 20 PCs account for 84% of the variance of the OF vectors. By removing all PCs past PC #20, the dimensionality of the data

can be reduced ten-fold (from approx. 200 vectors to 20 signals), reducing the computational complexity of ICA while keeping 84% of the information in the data. In practice, choosing more PCs does not make a difference in source quality, as each additional PC would contribute less than 1% of the total variance of the original signals (PC #20 only explains 0.8% of the variance). Moreover, as mentioned, source selection worked much better when the number of PCs and ICs was equal. When this was not the case, the obtained sources were very homogeneous, which did not enable proper separation. Therefore, the chosen number of PCs and ICs for most segmentation steps in the algorithm (see Section 5.6) was 20.

5.5.2 ICA And Source Selection

After PCA was applied, ICA was performed on the obtained PCs. In order to isolate seizure movement, only a select number of sources needed to be chosen, and these would be added together to reduce the multivariate signal into one source of seizure movement. This would be performed for the horizontal (X) and vertical (Y) components of the OF vectors, like every step that has been described so far. ICA can be calculated in a "deflationary" or "parallel" approach, which concerns the order in which the sources are calculated. The deflationary approach was chosen because an exploratory analysis found it was faster and provided better results. For the purpose of selecting the sources that most resembled seizure movements, a novel metric is herein proposed, called the ***Temporal Consistency Factor (TCF)***. Before defining this metric, it is useful to look at an example of the sources estimated using ICA. *Fig. 5.11* shows 20 ICs obtained by calculating ICA for the first 20 PCs (obtained from previously applying PCA) of the OF vectors of a seizure segmented from the video of *p101_s3*.

ICA produced independent sources with practically identical standard deviation, which is expected (as whitening the data entails ensuring the input signals have equal and unitary variance). However, as can be observed in *Fig. 5.11*, some of the estimated sources have very high and isolated peaks, while others contain very consistent signals. The sources with these peaks, such as source #2, #3 or #5, contain high amplitude movements likely caused by health professionals walking quickly in front of the camera, while more consistent sources most likely contain the very consistent and oscillatory clonic movements present in the focal to bilateral tonic-clonic seizure that patient 101 had in this video. Beyond the aforementioned need to separate seizure movements from other movements, its also important to remove these peaks, because they could be confounds in any further analysis that takes into account movement frequency or amplitude. It is then essential to develop a way of separating and isolating the best sources. It is also important to mention that, for this particular seizure, seizure-like movement can be fairly clearly identified in most independent sources. This is not always the case, especially when there are large obstructions of the camera's view or a large number of health professionals surrounding the patient. Therefore, for most seizures, as there are often fewer good sources, the process of isolating them becomes even more important.

The ***Temporal Consistency Factor*** is the metric that was developed for this purpose. It can be defined as a product of two metrics, the *Peak to SSR Ratio* **(P2SSR)** and the *Percentage of Low and High Amplitude Samples* **(L/H%)**. The Peak to SSR ratio is a novel metric which was based on the *Peak*

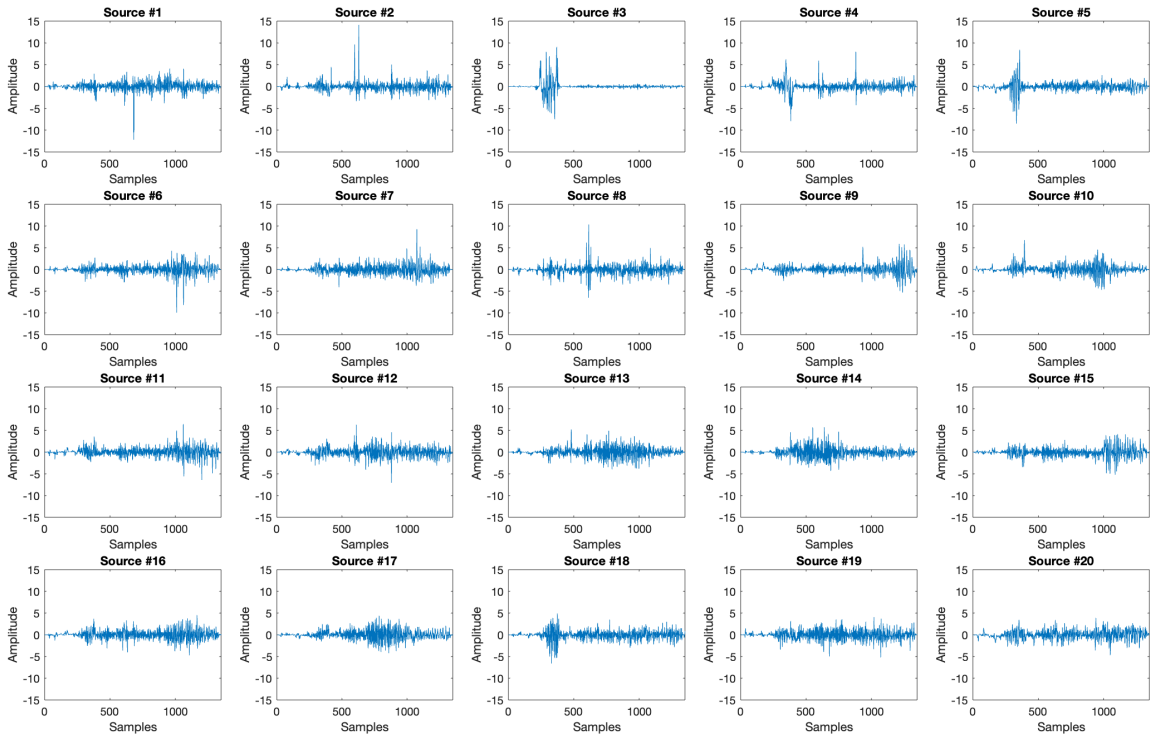

Figure 5.11: 20 independent components (or sources) estimated from the first 20 principal components obtained using ICA, after applying PCA to the OF vectors of a segment of the video of *p101_s3* which contains a seizure. Sources are numbered from 1 to 20 in the order they were calculated by ICA.

to RMS ratio (or Peak-Magnitude to RMS ratio), also known as the Crest Factor, which is used in audio and mechanical engineering applications [119]. The Peak to RMS ratio (P2RMS) is the ratio between the maximum (or peak) amplitude of the signal and its Root Mean Square level (RMS), which is the square root of the mean of the squared samples of the signal, as defined in Eq. 5.1:

$$RMS = \sqrt{\frac{\sum_{i=1}^{N} x_i^2}{N}} \tag{5.1}$$

Sources corresponding to seizure movement (e.g. source #19) have lower peaks, leading them to have a low P2RMS, while sources like source 3 have fairly low amplitude signals with the exception of a very high peak, and as such have high P2RMS. This metric can therefore be regarded as a proxy for consistency. However, because the sources have the same variance and standard deviation, the RMS of all sources is practically identical. Therefore, if P2RMS was used as a metric, it would effectively be the same as only using the source's peak. A further problem is that higher values contribute more towards a source's RMS value, because the individual samples are squared first (and higher values become comparatively even higher when squared). The denominator in this metric (RMS, in this case) should have higher values for more consistent signals with median amplitude, such that sources that contain seizure movement have a lower value in the overall metric. Therefore, if this is reversed (i.e. taking the **root of each sample and then squaring the sum of these roots**), higher values would contribute

less to the denominator, increasing the overall ratio for signals with higher, more isolated peaks, and decreasing it for more consistent signals, improving the differentiation between seizure movement and artifactual sources. This altered metric can be called Peak to *SSR* ratio (P2SSR), with SSR referring to Squared Sum of Roots. Lower P2SSR values would refer to more consistent sources, while higher values would refer to sources with more higher and more prominent peaks. Moreover, instead of simply dividing the peak by the SSR, the SSR value was effectively normalized by dividing it by its mean over all sources, ensuring that it would only be a comparative measure between the sources, and would not depend on source length or any such factors. SSR is defined mathematically in Eq. 5.2 and the P2SSR metric is defined in Eq. 5.3, with $s(t)$ denoting the sources.

$$\boldsymbol{SSR}(s(t)) = \left(\sum_{t=1}^{N_{Samples}} \sqrt{(s(t))} \right)^2 \tag{5.2}$$

$$\boldsymbol{P2SSR}(s(t)) = \frac{Max(|s(t)|)}{SSR(s(t))} \times \frac{\sum_{n=1}^{N} SSR(s_n(t))}{N} \tag{5.3}$$

The sources in *Fig. 5.11* can be sorted according to this metric, in order to visually inspect whether it is properly isolating the best sources. This can be seen in *Fig. 5.12*. The sources with the lowest P2SSR, namely sources #20, #17, #16 and #19, are vastly more consistent than those with the highest values (such as sources #3, #2 and #1, which show high peaks), and seem to include movement consistent with clonic movements. Therefore, it is safe to assume that P2SSR is a good proxy for consistency, and could be used as part of a metric that should differentiate seizure sources from other movements.

However, as it was previously mentioned, **TCF** was defined as the product of two metrics, namely P2SSR and the *Percentage of Low and High Amplitude Samples* **(L/H%)**. This was calculated as the percentage of samples in a source that had higher or lower (absolute) values than two set thresholds. Therefore, sources that had consistently median amplitudes, without many peaks such as those present in sources #1 and #2, and also without many empty periods such as in source #3, would have a **lower** L/H%, so that when this was multiplied with the P2SSR, it would amplify the differentiation between temporally consistent sources and sources with high "peakedness". These two thresholds that defined "high" and "low" values were defined visually as equalling 0.5 and 1.8 after applying a moving average filter with a window of L$\times$ 0.015 (with L being the length of the sources) to the absolute values of the signals, to smooth out the oscillatory movements and obtain the average amplitude of these signals over time. The estimated sources with this moving average filter and the two thresholds can be observed in *Fig. 5.13*.

When the order of the sources according to L/H% is compared to the order obtained with P2SSR (in *Fig. 5.12*), it is noticeable that some sources, namely source #7 and #12, are given much better ratings, while others such as sources #5, #9 and #10 are placed much lower on this list. The sources that are placed higher are fairly consistent and only have very short peaks that are smoothed out by

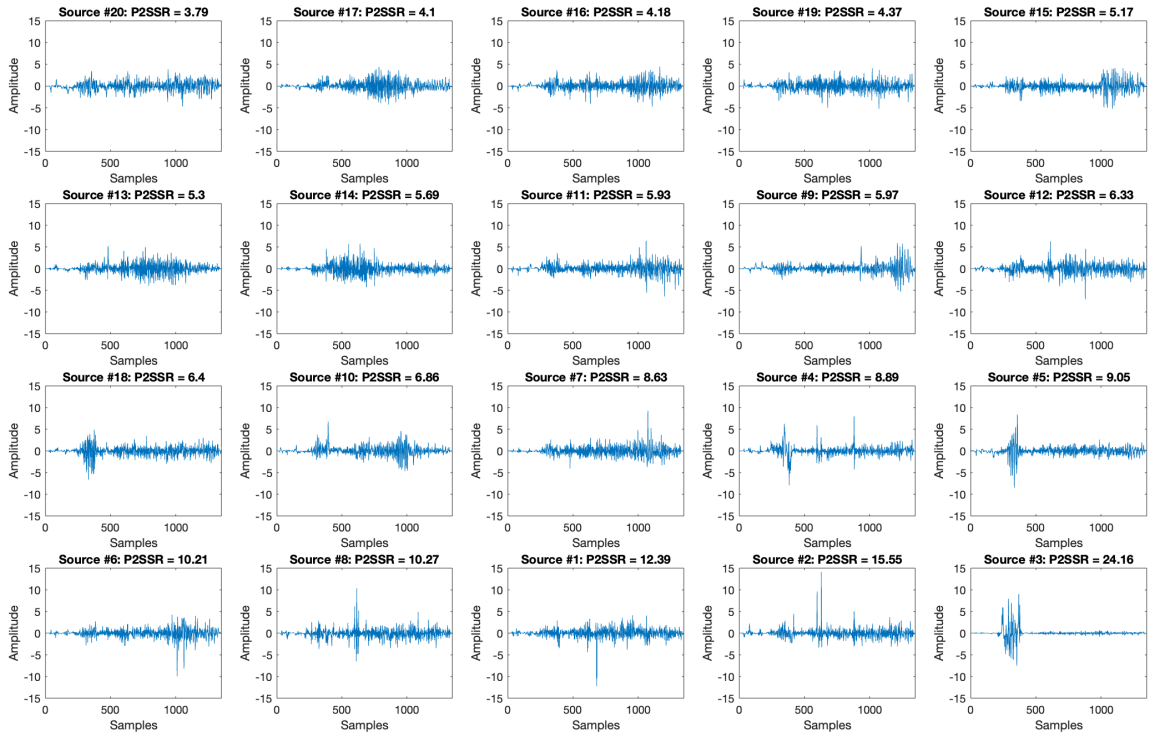

Figure 5.12: The 20 ICs estimated using ICA on *p101_s3*, ordered according to their Peak to SSR ratio (P2SSR - defined in Eq. 5.3), from the lowest to the highest. Lower P2SSR values indicate less "peakedness", i.e. more temporal consistency. Sources are numbered from 1 to 20 in the order they were calculated by ICA.

the moving average. Although the sources that are placed lower according to L/H% are not necessarily inconsistent, they seem to show short periods with very high amplitudes, such as what can be observed at the start of source #5 or at the end of source #9. These could be artifacts, or just parts of the signal that ICA considered to be independent. However, it is useful to rate these sources lower as they are not very consistent overall. It is again important to mention that this is a very clear seizure, and in more ambiguous cases these metrics have a higher impact. L/H% can be defined mathematically as follows, in Eq. 5.4 (with $s(t)$ being the ICs):

$$\boldsymbol{L/H\%}(s(t)) = \frac{\sum_{t=0}^{L} LH(t)}{L}, \; with \; LH(t) = \begin{cases} 1, & \text{if } Mov.Avg.(s(t),\, w = 0.015 \times L)(t) < 0.5 \\ 1, & \text{if } Mov.Avg.(s(t),\, w = 0.015 \times L)(t) > 1.8 \\ 0, & \text{otherwise} \end{cases} \tag{5.4}$$

The product of P2SSR and L/H% was used as the main metric for ICA source selection. This was named, as mentioned, the ***Temporal Consistency Factor***, and is defined in Eq. 5.5.

$$\boldsymbol{TCF}(s(t)) = P2SSR(s(t)) \times L/H\%(s(t)) \tag{5.5}$$

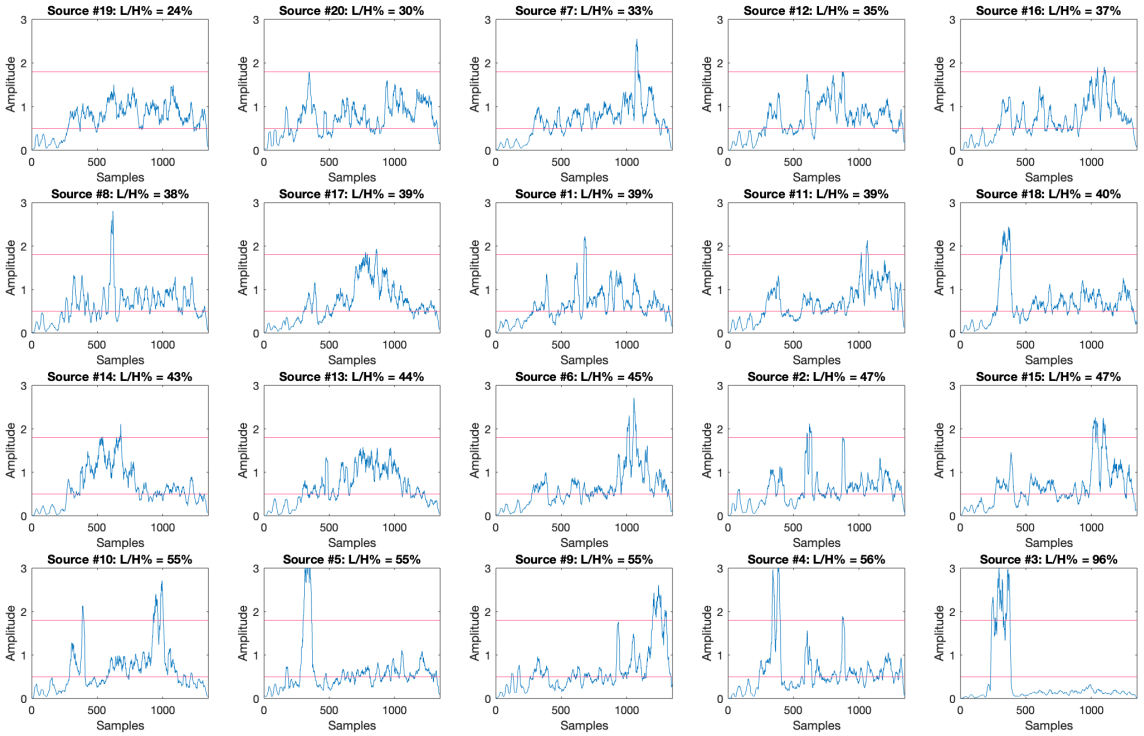

Figure 5.13: The 20 ICs estimated using ICA on *p101_s3*, ordered according to their *Percentage of Low and High Amplitude Samples* (L/H% - defined in Eq. 5.4) from lowest to the highest and passed through a moving average filter with a window of size 0.015 $\times$ the length of the sources; and two thresholds with values of 0.5 and 1.8. Lower L/H% indicates that the source contains fewer or shorter high peaks and periods of low amplitude. Sources are numbered from 1 to 20 in the order they were calculated by ICA.

In *Fig. 5.14a*, the sources are plotted and ordered according to their TCF values.

Visually, it seems that this metric has been successful in locating the most seizure-like sources, such as source #19 and #20, and the most artifactual, such as sources #2, #3, and #5.

The actual selection of the correct sources was performed using a threshold. Sources with a lower TCF than the threshold were selected and summed together into one signal that described the seizure. The seizure detection and segmentation algorithm was then largely based on evaluating these resulting signals on their TCF and classifying them using other features. This will be described further in Section 5.6. In Chapter 6, the quality of the source selection will be analysed further by comparing the sum of the selected sources with the sum of all sources using time-frequency analysis.

It is important to mention that ICA (due to the whitening step) and PCA[2] have a degree of randomness involved. Therefore, although results are fairly consistent, they are not always exactly repeatable. This influenced the training and testing of the classifier (Section 5.7), as well as how the results (Chapter 6) were calculated.

[2]PCA is a deterministic method, but the `sklearn` implementation that was used, which calculates PCA with Singular Value Decomposition (SVD), utilizes a randomized solver for large enough data: https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.PCA.html

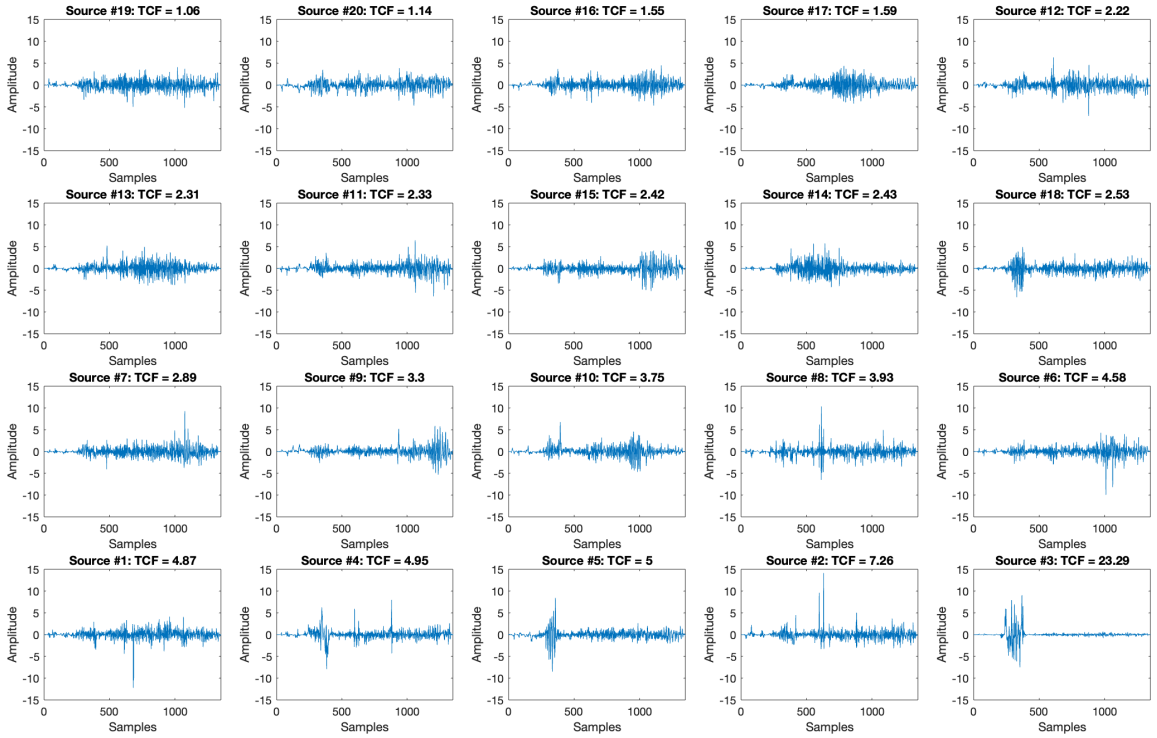

(a) The 20 ICs estimated using ICA on *p101_s3*, ordered (from the lowest to the highest) according to their Temporal Consistency Factor, defined as the product of the Peak to SSR ratio (P2SSR - defined in Eq. 5.3) and the Percentage of Low and High Amplitude Samples (L/H% - defined in Eq. 5.4). Lower TCF values indicate less "peakedness", i.e. more temporal consistency. Sources are numbered from 1 to 20 in the order they were calculated by ICA.

5.6 Seizure Segmentation Algorithm

The ICA-aided source selection plays a central role in the seizure detection algorithm. In general terms, this algorithm followed the structure that can be seen in the flowchart in *Fig. 5.15*. In Appendix C, a much more detailed flowchart describes the functioning of the algorithm in a more comprehensive manner.

This algorithm was built to accomplish two main goals, namely issuing timely and accurate alarms when seizures are detected, and accurately annotating the beginning and end of the seizure, for seizure segmentation purposes. It works by combining **two ICA-aided segmentation steps** and a final classification step.

Regarding the first segmentation step, after calculating and denoising the OF vectors resulting from a 45 second long video segment, PCA and ICA are performed in succession. The sources that have a lower TCF than a set threshold (henceforth referred to as the **Source Selection Threshold**, or T1) are selected and summed together. This source selection step, like the noise reduction steps, is performed for both the X and Y OF vectors. Then, if the resulting signal has a lower TCF than a second threshold (henceforth referred to as the **Episode Selection Threshold**, or T2) for either the X **or** the Y vectors, this 45s long episode is considered to be successfully selected. If both signals resulting from the source selection of X and Y OF vectors have a TCF above the Episode Selection Threshold for this step, then it is determined that a seizure is not occurring in this 45s long period. These thresholds were tuned using

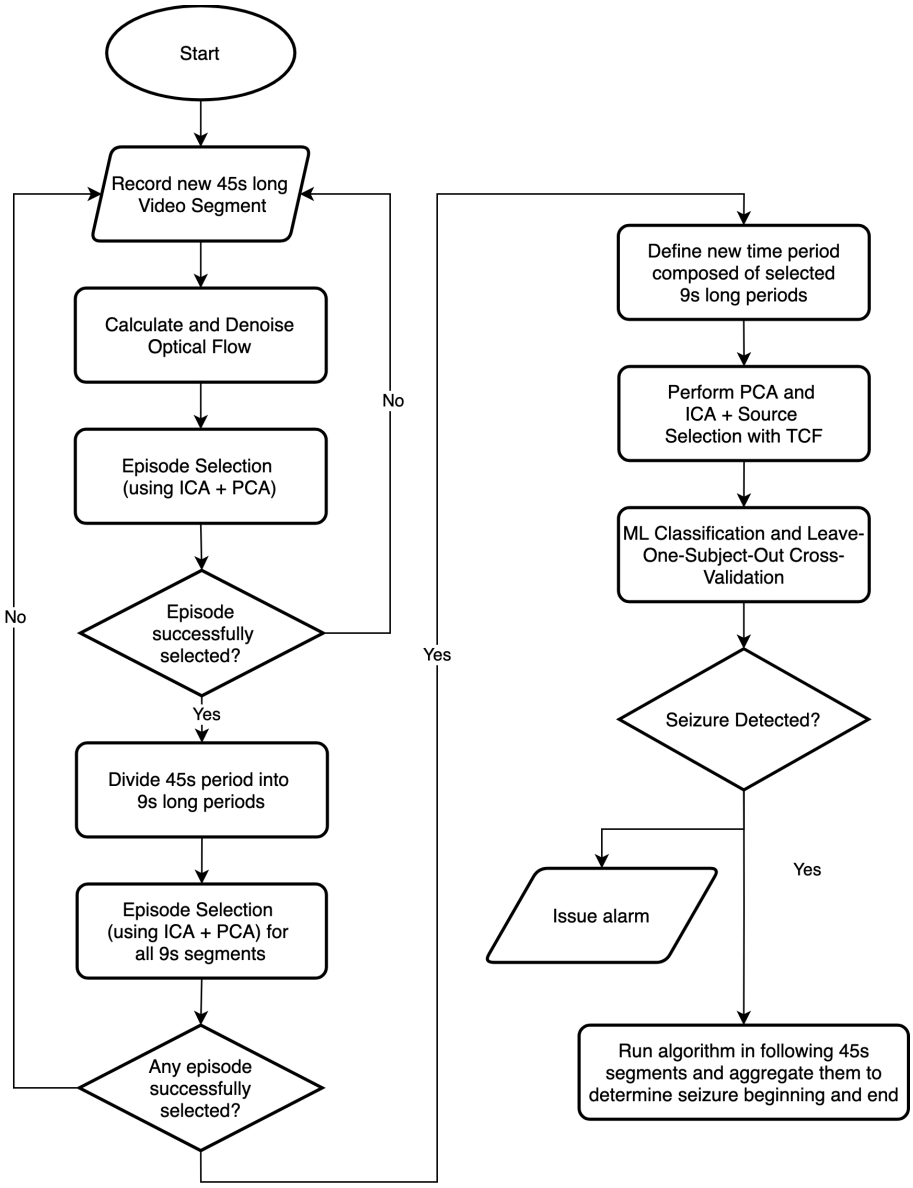

Figure 5.15: Simplified flowchart detailing the structure of the seizure detection algorithm. After the X and Y OF vectors are denoised according to the process defined in Section 5.4.2, two successive segmentation steps are performed. Then, if a possible seizure episode is segmented, ICA source selection is performed on the OF vectors during this period to obtain a time series, from which features are calculated. These features are then used to perform classification; if a seizure is detected from this classification, an alarm is issued, and then this algorithm is run again for the following 45s periods to determine the beginning and end of the seizure.

the training set that is detailed in Section 5.1. This tuning is described in Chapter 6. 45 seconds was chosen as the length for the first step as seizures often lasted between 45 seconds and 2 minutes, and as such most seizures would be as long as one or two of these periods.

To simplify, TCF was not only used as a metric to select the best sources, but also to select potential seizure episodes, i.e. time periods where the signal is temporally consistent enough to possibly constitute a seizure. Because of this step, only consistent movement will be able to generate false alarms (albeit not often, due to other metrics that were introduced in later steps and contribute to a fairly low false alarm rate). The justification for the use of this metric as an episode selection tool is that the clonic

movement present in clonic and tonic-clonic seizures (see Section 2.1.2) generally maintains a constant amplitude over time, only stopping once the seizure ends. This very consistent nature of seizure episodes further reinforces the use of a metric like TCF for episode selection.

After this step, the second segmentation step is performed for every 45s long segment that passed the first step. While the objective of the first step was to discriminate whether a seizure could be occurring, the goal of the second step is different, namely to more finely segment the selected 45s periods, in order to remove any parts of these periods which could serve as confounds in further evaluation of these signals using other features. For example, if a seizure starts in the middle of such a 45s segment, it is useful to remove the first part in order to improve the accuracy of the classification.

This step was very similar to the first segmentation step, in that after the 45s period was divided in five 9 second periods, these underwent PCA and ICA source selection with a *Source Selection Threshold* (T1) and then episode selection with an *Episode Selection Threshold* (T2), with the selected 9s episodes later being aggregated together to form a new time period. As an example, if the first 9s episode is not selected, but every following one is, then the newly segmented time period starts 9 seconds into the previously selected 45s period, and ends at its end, namely at 45 seconds.

In case a 9s episode in the middle of the 45s period is not selected, but the periods before and after are, then there is a case of having two segmented periods, one before and one after this episode that was not selected. In this case, the average TCF for each of these is compared, and the segment with the lowest TCF is chosen. The segment that is chosen from this second segmentation step is then used for classification.

After this new time period is defined, PCA and ICA source selection are performed on the denoised OF vectors, in order to obtain a univariate time series signal from which features for classification can be calculated. Before this step, however, a value of TCF is calculated for all X and Y OF vectors during the selected time period. Only those vectors with a TCF lower than $1.3\times$ the mean TCF over all OF vectors are selected to be used for source selection and classification (which in practice means only a few noisy vectors will be excluded). This ensures that a cleaner signal is obtained, improving the quality of the features that will be used for classification.

If the classification step (that will be discussed in Section 5.7, and uses features related to movement frequency and the total amount of movement) is successful, and the detected seizure period is **longer than the minimum time period of 9 seconds**, an alarm is issued. Therefore, this alarm is only issued for segmented time periods of between 18 and 45 seconds, as 9 seconds was deemed too short to constitute a real tonic-clonic seizure. Moreover, alarms are only issued in case the previous segment did not generate an alarm, as issuing repeated alarms for the same seizure is, naturally, not useful. Afterwards, these segmentation steps are run again for the following 45 second periods. If more time periods are selected and classified as seizures, these are aggregated together to determine the beginning and end of the seizure, for seizure segmentation purposes.

5.7 Classification

Although the segmentation steps do a good job of detecting most, if not all, seizure episodes, they detect a fair number of non-seizure episodes too. It is therefore essential to improve specificity (i.e. decrease the false alarm rate), for which machine learning was herein explored. The approach that was initially devised consisted of a Convolutional Neural Network for classifying the resulting spectrogram of the selected time periods (as an image). However, the available datasets did not lend themselves to this approach due to the very small number of data points available (the training set only consisted of 7 tonic-clonic seizures, while the validation set had 21, as is discussed in Section 5.1). Therefore, other classification approaches were considered. The chosen approach was based on the calculation of two features, "*spectral contrast*" and "*amount of movement*" and classification was tested with three methods: support vector machine classification (SVM), gaussian process classification (GPC), and multilayer perceptron classification (MLP). More information about these methods is presented in Section 2.2.3.

One of the features that was developed ended up being very similar to a feature that was used in a paper published by Kalitzin et al. in 2012 [107] and was named "*spectral contrast*". Although their methodology was very different[3], due to the similarities present in this particular feature, their name was adopted and credit is given to them for developing this feature.

Spectral contrast is defined in this work as the ratio between the power of the signal in the frequency band between 2 and 7 Hz, and its power in the frequency band between 0 and 2 Hz. It is a measure of the rhythmicity and high frequency of clonic seizure movements when compared to other, more natural movements, that have higher power in lower frequencies. The power of each frequency band (or "*band power*") is calculated by determining the area of the signal's *periodogram* (which plots power spectral density over different frequencies)[4] between these two frequencies. The general definition of this feature can be seen in Eq. 5.6, with P referring to the signal power in a certain frequency band.

$$\boldsymbol{Spectral\ Contrast} = \frac{P(2Hz - 7Hz)}{P(0Hz - 2Hz)} \tag{5.6}$$

The second feature that was implemented was named "*amount of movement*". It is defined as a product of the standard deviation of the total movement and the area it occupies within the frame.

The standard deviation $\boldsymbol{\sigma}$ of the movement during the selected time period is calculated by multiplying the ICA selected sources by the *unmixing matrix* W (see Section 2.2.2) to obtain the reconstructed OF vector components that correspond to these sources, and then calculating the standard deviation of their sum. Equation 2.4, in Section 2.2.2, describes the aforementioned first step of computing the observations from the sources.

The standard deviation of the total movement is then calculated as follows (Eq. 5.7), with L being the length of the time period, μ being the mean of the sum of observations $\sum_{i=1}^{N} x_i(t)$ and N being the

[3] They did not use ICA or PCA, instead reconstructing group velocities from OF, and also did not use machine learning classification, instead using thresholding. The only relevant similarity with this work is therefore the use of this feature and OF, which is a very common tool for seizure detection using video, as is explained in Chapter 3

[4] Calculated with scipy's `scipy.signal.periodogram`

number of reconstructed observations (i.e. the number of reconstructed OF vectors, as defined by the `step` value chosen when calculating OF):

$$\boldsymbol{\sigma} = \sqrt{\frac{\sum_{t=0}^{L}\left(\sum_{i=1}^{N} x_i(t) - \mu\right)^2}{L}} \times \frac{1}{N} \tag{5.7}$$

The area $\boldsymbol{A}$ of the movement within the frame is calculated by obtaining the reconstructed observations with the unmixing matrix W like it was done for the previous feature, and then computing the percentage of reconstructed OF vectors (which are uniformely spatially distributed in the frame) whose total movement (i.e. their sum over the selected time period) is larger this sum's average for all reconstructed OF vectors. Therefore, mathematically, it can be defined as it is in Eq. 5.8, with m_i being the number of reconstructed OF vectors with more movement than the average reconstructed OF vector, $\sum_{t=0}^{L} x_i(t)$ being the sum of the movement in one OF vector x_i over the selected time period, N being the number of reconstructed OF vectors and L being the length of the time period:

$$\boldsymbol{A} = \frac{\sum_{i=1}^{N} m_i}{N}, with \tag{5.8}$$
$$m_i = \begin{cases} 1, & \text{if } \sum_{t=0}^{L} x_i(t) > \frac{\sum_{i=1}^{N}\left(\sum_{t=0}^{L} x_i(t)\right)}{N} \\ 0, & \text{otherwise} \end{cases}$$

As mentioned, these features were used to classify the selected seizure episodes using three machine learning classification methods, namely SVM, GPC and MLP. With the training set, the optimal hyperparameters for these classifiers were found with a grid search using a 10-fold cross-validation.

It is worth repeating that this classification step is simply a way of attempting to increase specificity. As such, in training, the negatively labelled data points will be episodes that have been selected as part of the episode selection and segmentation algorithm described in Section 5.6 but don't represent real seizures, therefore being false positives. The positively labelled data points are episodes which have been selected by the algorithm and do correspond to real seizures (during the selected time period a seizure was occurring).

To obtain a training dataset for the hyperparameter choice, the two aforementioned features were calculated with videos from the training set (see Section 5.1). However, this training set was very small, only consisting of 7 videos of seizures and 19 videos without seizures. As mentioned in Section 5.5.2, ICA has a degree of randomness, which leads the source selection and segmentation process to produce slightly different results for every time the algorithm is run. This can be taken advantage of in order to increase the size of the dataset without leading to overfitting. This was done by running the segmentation algorithm 3 times and using the three different values of the two features for these three runs for each video, leading to a dataset that is three times larger than what was available. In order to increase the number of false positives to ensure the dataset is balanced, the episode selection threshold was increased for the calculation of these features, increasing the number of false positives. This set can be

Figure 5.16: The training set used for classification. It was obtained by running the segmentation and episode selection algorithm described in section 5.6 three times on the videos in the training set described in 5.1, and calculating the *spectral contrast* and *amount of movement* features for all selected episodes. These data points were labeled as *negative* if they were false alarms and *positive* if a seizure did indeed occur during the selected time period. The negative points are in red and the positive points are in blue.

seen in *Fig. 5.16*. In an initial visual inspection, it seems that the separation between the two classes is fairly good, with the exception of a few positively labelled data points which show similar values to those present in the negative set.

The first classifier that was tested was a non-linear SVM with an RBF kernel [5]. For this classifier, the hyperparameters to be found were the "γ" and "C" parameters, which, in general terms, can affect the bias and variance of the classifier, affecting its accuracy and how over- or underfit it is. This is explained in further detail in Section 2.2.3. A grid search was conducted using 10-fold cross-validation to estimate classification accuracy. The γ parameter was varied between 10^{-9} and 10^{3}, while the C parameter was varied between 10^{-2} and 10^{9} [6]. The best parameters were $C = 10$ and $\gamma = 1$, with an accuracy score of 96%, as can be seen in *Fig. 5.17a*.

However, as can be observed in *Fig. 5.17b*, these parameters seem to cause the model to be overfit to this training set. Therefore, because the accuracy score $C = 10$ and $\gamma = 0.1$ was very similar, at 95.3%, and it caused the model to be much less overfit, a compromise was reached by choosing the parameter values $C = 10$ and $\gamma = 0.2$.

The hyperparameter that was tuned for the GPC was the kernel, as the other hyperparameters did not have any influence on the accuracy scores of this classifier in the training set. A mean accuracy score was computed for each kernel using 10-fold cross-validation[7]. The `Matern`, `RBF` and `Rational Quadratic` kernels obtained the best results, with an accuracy of 97.3%, while the second-best performing kernel was the `Dot product` and kernel, with 94% accuracy. The `White` kernel only managed an accuracy score of 51%.

The decision function for the three best kernels (RBF, Rational Quadratic, Matérn) can be seen in

[5] The `SVC` model from the `sklearn` package was used

[6] This was done with the help of an existing script from scikit learn's website, which was modified to suit this specific use case: https://scikit-learn.org/stable/auto_examples/svm/plot_rbf_parameters.html#sphx-glr-download-auto-examples-svm-plot-rbf-parameters-py, last accessed on 07/11/2021

[7] This was done with the help of an existing script from an online tutorial, which was modified to suit this specific use case: https://machinelearningmastery.com/gaussian-processes-for-classification-with-python/, last accessed on 07/11/2021

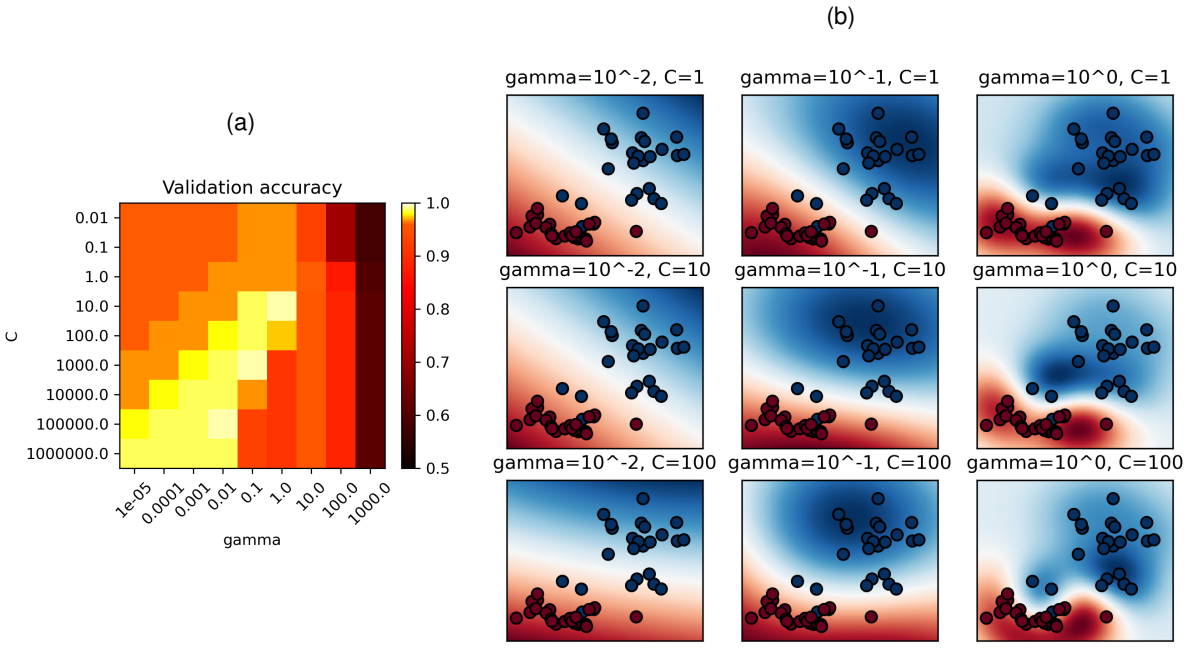

Figure 5.17: (a) – The accuracy scores for the different parameter (γ and C) combinations for a non-linear SVM using an RBF kernel. The parameter combination with the best score was $C = 10$ and $\gamma = 1$, with an accuracy score of 96%, but many other parameter combinations provided accuracy scores above 95%. (b) – The decision function for the implemented non-linear RBF SVM plotted over the training set, for different parameter combination. The combination with the best score was $C = 10$ and $\gamma = 1$, which can be seen in the center-right plot. However, visually, this seems overfit. The center plot, corresponding to $C = 10$ and $\gamma = 0.1$ seems to show lower variance towards this training set, with a simpler curve.

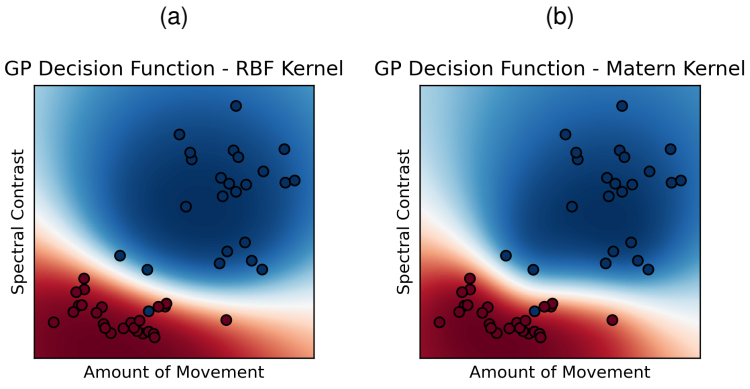

Figure 5.18: (a) – The decision function that resulted from fitting the training data to the GPC using the RBF and Rational Quadratic kernels. These two kernels produced a decision function with same shape. This shape seems to indicate this model would have quite a low variance, as it does not seem to be overfitting to the training data. (b) – The decision function that resulted from fitting the training data to the GPC using the Matérn kernel. This shape seems to indicate this model would have a fairly low variance, albeit a slightly higher one than the RBF and Rational Quadratic kernel models.

Fig. 5.18. However, as the RBF and Rational Quadratic kernels produced the exact same decision function shape, only the the RBF and Matérn kernel decision functions are shown. Using the Matérn kernel seems to lead to slightly higher variance, but not to the point of the model seeming overfit. In any case, the RBF kernel was chosen, as a slightly lower variance could lead to better generalization.

Finally, the MLP was the last classifier that was implemented [8]. This classifier takes a few parameters, namely the `solver`, `activation`, α and maximum number of iterations. The latter was defined as 1000, as larger numbers could compromise computational efficiency, while the `solver` was chosen as "lbfgs", as it is indicated for smaller datasets [9]. A grid search was conducted to determine the type of activation and the value for α, with 10-fold cross-validation. The possible activation types were "ReLU" (the most common type), "logistic", and "tanh". The parameter α was varied between 10^{-5} and 0.1. The parameters which yielded the best accuracy score were all activations when paired with $\alpha = 0.1$. *Fig. 5.19* shows the different accuracy scores and decision functions for the different parameter combinations. From the parameter combinations that yield the highest scores (as can be noted in *Fig. 5.19a*), "ReLU" activation also yields possibly the least overfit decision function, as can be seen in *Fig. 5.19b*.

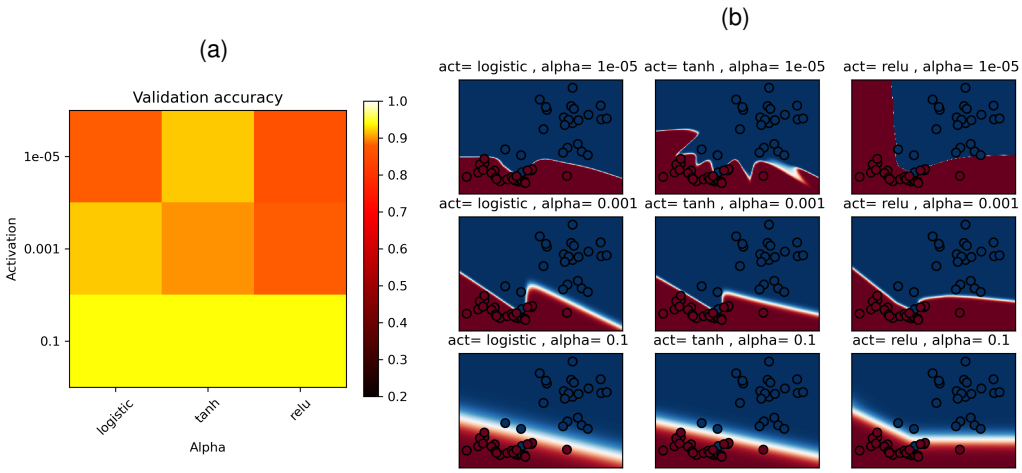

Figure 5.19: (a) – The accuracy scores for the different parameter combinations of the MLP classifier, assessed using 10-fold cross-validation on the training set. The highest scoring parameter combinations, at 94%, were all three activations when paired with $\alpha = 0.1$. (b) – The decision functions that resulted from fitting the training data to the MLP classifier using different parameter combinations, indicated above each plot. Some seem underfit, such as the combinations of the "logistic" and "tanh" activations with $\alpha = 0.1$, only having a linear separation. Others seem overfit, with overly complex decision functions, such as the combination of "tanh" activation and $\alpha = 10^{-5}$. The most balanced parameter combination, is also one of the combinations with the best score ("ReLU" activation and $\alpha = 0.1$), seen in the bottom right plot.

The parameters chosen for the three classifiers were $C = 10$ and $\gamma = 0.2$ for the SVM, the RBF kernel for the GPC, and $\alpha = 0.1$ and ReLU activation for the MLP. These were tested with the videos in the validation set, along with the full seizure detection algorithm. The results from these tests will be presented next, in Chapter 6.

[8] Implemented in Python using sklearn's `MLPClassifier` function

[9] According to the page on sklearn's website referring to the `MLPClassifier`: https://scikit-learn.org/stable/modules/generated/sklearn.neural_network.MLPClassifier.html?highlight=mlp#sklearn.neural_network.MLPClassifier, last accessed on 08/11/2021

Chapter 6

Results And Discussion

In this chapter, the results of the implemented methods, namely the noise reduction, source selection, and seizure segmentation and classification algorithm will be discussed and their performance will be evaluated. This evaluation will be based on objective metrics that measure the usefulness of the implemented methods and their accuracy. Notwithstanding this quantitative assessment, this analysis will be complemented with visual metrics such as plots and graphs that show the qualitative impact of the different processes.

Only the training set was used for the purposes of choosing source selection, episode selection and noise reduction threshold values. Since these do not concern the final classification step, for the purpose of choosing their values, the algorithm was tested without the classification. The validation set was withheld from all steps except for the final evaluation of the algorithm, in order to fairly determine its performance without delivering overly optimistic or unfair results. An exception to this was the choice of classifier, between SVM, GPC and MLP, which was assessed using the validation set, leave-one-subject-out cross-validation (LOSOCV), and Receiver Operating Characteristic (ROC) curves. The metrics used to objectively judge performance were sensitivity, specificity, root mean square segmentation error (segmentation RMSE) and alarm latency.

The **sensitivity**, or true positive rate, was defined as the percentage of annotated seizures which were detected. A seizure was considered to be correctly detected when this detection corresponded to a time period (of between 18 and 45 seconds) that at least partly overlapped with the annotated seizure period. As previously mentioned, seizure starts and endings for the validation set were annotated by Dr. Francisca Sá and video-EEG technician Octávia Brás of HEM, while for the training set the start of the seizures was annotated by health professionals at HSM and the ending was manually annotated by the author.

Specificity was defined as $1 - False\,Positive\,Rate$, with the false positive rate being defined as the number of false positives per analysed video sequence. Although a similar approach has been used in other studies [105], this is not an ideal definition, and it would be better to specify the false alarm rate per 24h. However, this would not be a fair measure to compare with other studies which used continuous, 24h video monitoring, because the available videos in the training and validation sets were only at most

30 minutes long and contained a significant amount of movement compared to an average 30 minute segment (which, for instance, could contain large time periods with no movement, at night).

The **alarm latency** was defined as the difference between the annotated seizure start and the earliest moment in time after this seizure start in which this seizure could be detected by this algorithm, which would be at the end of a selected 45s period that has been correctly classified as a seizure. Therefore, if the seizure starts in the middle of a selected and correctly classified 45s period, the latency is the time between the annotated seizure start and the end of said 45s period, being therefore lower than 45 seconds. If only the next period is classified as a seizure (and the seizure has not yet ended by the start of this second period, making this classification correct), the latency will be higher than 45 seconds.

The segmentation error is the difference between the annotated and detected start and finish of a seizure, measured in seconds. The metric used to quantify this error was the **Root Mean Square Error (RMSE)** of the segmentation, defined as the square root of the squared segmentation errors. This measure is helpful because it more heavily penalizes outliers than a simple average.

6.1 Noise Reduction and Calibration

In this section, the noise reduction methods detailed in Section 5.4.2, namely the adaptive and fixed thresholds, the low-pass filter, and the ROI, will be evaluated in their capability to attenuate or remove the noise present in the OF vectors, as well as their role in improving the accuracy of the seizure detection and segmentation. Initially, after an exploratory analysis of the training data, the adaptive threshold was set at a value of 8 and the fixed threshold was set at 0.7. Later, the optimal value for these thresholds was computed using the training set.

However, initially it is important to assess these methods' performance during the calibration period, which, as mentioned, only contains noise, as no detectable movement is present. The time series signals of the horizontal (X) Optical Flow (OF) vectors and their respective Fourier transforms during this calibration period before the noise reduction, as well as the same time series signals and Fourier transforms for this period post-denoising can be seen in *Fig. 6.1* [1].

The noise has been completely nulled, which is a positive indicator of the performance of this method outside of the calibration period. It then becomes important to assure that outside of this period, these methods do not diminish the quality of the optical flow signals. To better assess this, the time series signals and respective frequency spectra for all X OF vectors are plotted in *Fig. 6.2*.

There are very few differences between the original and denoised signals. Some peaks appear to be reduced, but this is a very slight effect. In regards to the FFT frequency spectra, the low-pass filter has successfully removed the high frequency spikes at 12.5 Hz, but otherwise the difference is fairly unnoticeable. In order to more closely analyse these spectra on an individual basis, each OF vector FFT spectrum was plotted in *Fig. 6.3* in a surface plot (a 3-dimensional plot, with the x-axis corresponding to the frequency in Hz, the y-axis corresponding to the OF vectors and the z-axis showing the spectrum's magnitude). In these plots it becomes slightly easier to notice differences between the original and

[1] In this section, all plots for the Y components are omitted, as they show very similar results to the X components.

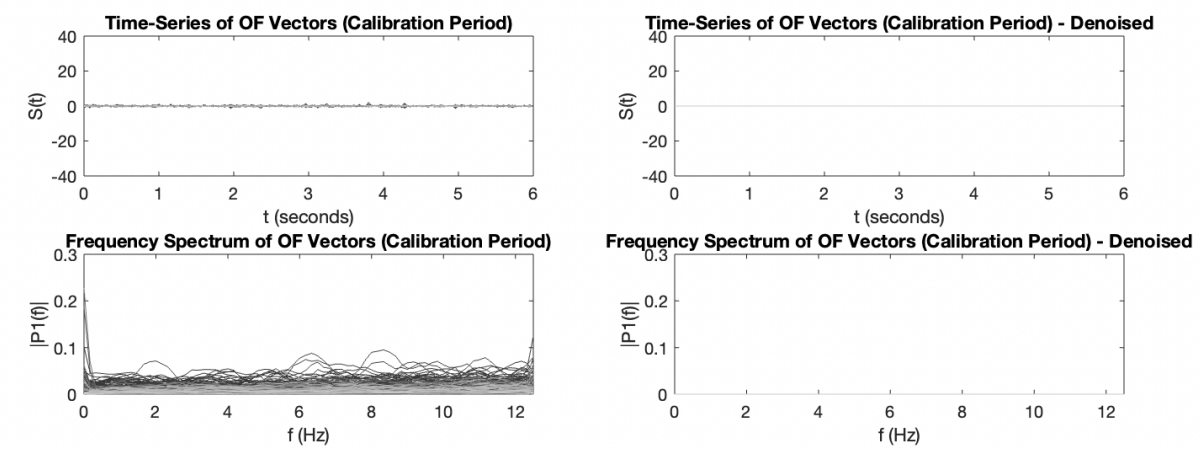

Figure 6.1: Time series signals and FFT frequency spectra of all OF vector X components during the calibration period, each plotted in a different gray tone to facilitate comprehension, before and after noise reduction. Positive signals indicate left to right movement, and negative signals indicate right to left movement (in the reference frame of the camera). The frequency spectra were obtained using Fast Fourier Transform and smoothed using an FIR Savinsky-Golay filter with a polynomial order of 5 and a frame length of 15.

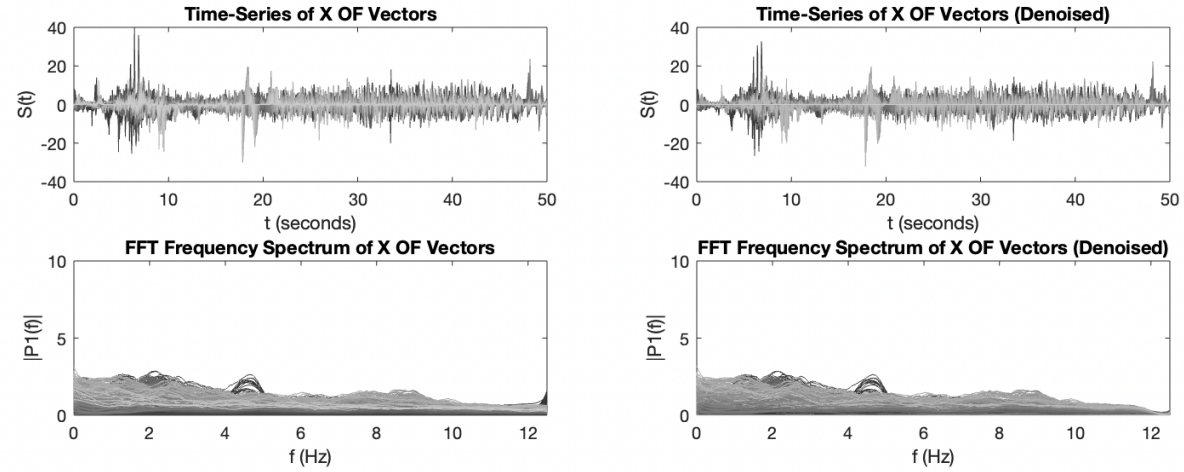

Figure 6.2: Time series signals and FFT frequency spectra of all OF vector X components during the seizure period, each plotted in a different gray tone to facilitate comprehension, before and after noise reduction. Positive signals indicate left to right movement, and negative signals indicate right to left movement (in the reference frame of the camera). The frequency spectra were obtained using Fast Fourier Transform and smoothed using an FIR Savinsky-Golay filter with a polynomial order of 5 and a frame length of 101.

denoised OF vectors. For instance, some spectra show a lower amplitude in lower frequencies, between 0 and 1 Hz. The 12.5 Hz spikes, which are only present in some OF vectors, were successfully removed, as could be seen in *Fig. 6.2*. The most important takeaway is that, so far, the signals don't seem to be severely affected by these methods, and the noise reduction does not seem to be destructive.

After assessing the performance of the implemented noise reduction methods in the time and frequency domains, it is then important to ascertain their impact regarding the video frame's spatial domain. *Fig. D.1* (in Appendix D) shows the standard deviation of all OF vector X components during the seizure period, in the spatial organization of the frame. It can be observed that the OF vectors in the center

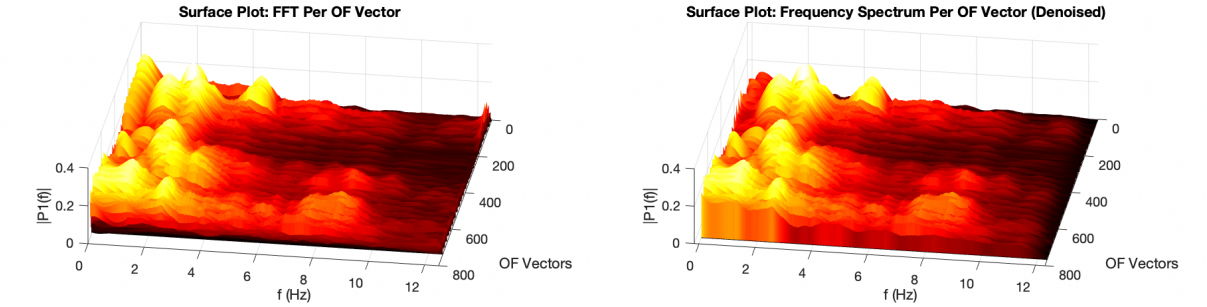

Figure 6.3: Surface plot of FFT frequency spectra of all X OF vectors during the seizure period, before and after noise reduction. The x-axis corresponds to the frequency in Hz, the y-axis corresponds to the OF vectors and the z-axis shows the magnitude of the frequency spectra. The frequency spectra were obtained using Fast Fourier Transform and smoothed using an FIR Savinsky-Golay filter with a polynomial order of 5 and a frame length of 101.

of the frame (where the patient is located) have standard deviation because of the seizure movements. The denoised OF vectors can also be seen in *Fig. D.1*, and show a very similar distribution, with the obvious difference that the OF vectors outside of the ROI are nulled. It is clear that spatially, from the implemented noise reduction methods, the ROI mask has had by far the biggest impact in the standard deviation of the OF vectors, and contributes to a better isolation of seizure movement by removing movement from other people in the frame.

These noise reduction methods were evaluated using further metrics obtained from their frequency spectra, namely the center-of-mass (COM) and the peak frequency of the spectra. The values for these metrics for each OF vector are plotted in the spatial organization of the video frame in *Fig. D.2* in Appendix D.

Consistent with what can be observed in *Fig. 6.1* and *Fig. 6.3*, most peak frequency values are fairly low. This is most likely due to non-seizure movements that have low frequency components with high power and amplitude, such as, for instance, health professionals walking in front of the camera. However, on the top left of the video frame, a few OF vectors have a peak frequency at 12.5 Hz. This 12.5 Hz spike can be observed in *Fig. 6.1* and was previously mentioned in Chapter 5. This spike is removed by the noise reduction methods. Regarding the COM of the frequency spectra, which is arguably a better measure of the frequency content than the peak frequency, it can be noted that some OF vectors have higher values, with COM's of around 5 Hz. These seem to be mainly located in the area of the frame corresponding to the movement of the patient's bed. A fair number of these vectors are nulled by the ROI.

Finally, and most importantly, it is necessary to evaluate whether these methods positively affected the accuracy of the seizure detection or quality of the segmentation, in order to determine analytically whether or not they are useful. In order to assess this, the algorithm's performance (without the classification step) on the training set with all noise reduction methods was compared to its performance without each of the individual methods. For this purpose, it is useful to take advantage of the randomness introduced by PCA and ICA to take multiple data points for each of these 4 versions of the algorithm (one with all methods, one with the filter and thresholds but without the ROI, one with the ROI and thresholds but without the filter, and one with the filter and ROI but without the thresholds). Therefore, the

model was run 60 times for each of these 4 versions, which yielded 4 distributions. Then, by comparing these distributions using box plots and violin plots, the specific contribution of each method to the overall performance of the algorithm could be evaluated. Furthermore, by calculating p-values with the non-parametric Wilcoxon Rank-Sum Test, it could be determined whether each contribution was statistically significant. *Fig. 6.4* shows the box plots and violin plots for these 4 distributions for all 4 aforementioned performance metrics, namely sensitivity, specificity, segmentation RMSE and latency. On the left the control can be noted, which is the algorithm with all noise reduction methods. The next distributions correspond to the model without the ROI (woROI), the model without the low-pass filter (woFilter) and without the thresholds (woThresh). Above these plots, the p-value calculated with the Wilcoxon Rank-Sum Test is visible. This test was used in lieu of a student t-test due to the fact that distributions are arguably not normal, and this test is non-parametric, as opposed to the t-test. This p-value corresponds to the comparison between the control distribution and the other individual ones. A p-value of under 0.05 was considered a statistically significant difference between the distributions. The mean values of the distributions are also visible in this plot as diamond-shaped points.

It becomes clear that the fixed and adaptive noise reduction thresholds are by far the noise reduction method that most significantly impacts the algorithms performance. The p-values for the difference between the distributions obtained with and without these thresholds are at their lowest possible value ($p - value < 2.22 \times 10^{-16}$), indicating a distinct difference in these distributions. Without the thresholds, the model performs worse in all metrics except latency. When these thresholds are excluded, the mean segmentation RMSE rises from 15.63 to 60.41; the mean sensitivity falls from 95% to 66%, and the mean specificity falls from 72% to 5%. Average latency also falls from 44.28s to 37.36s, possibly simply as a consequence of only detecting the seizures that are the easiest to detect. Both the ROI and low-pass filter seem to affect sensitivity in a statistically measurable way, with the filter improving it ($p - value = 0.00039$) and the ROI slightly decreasing it ($p - value = 0.0026$). This very slight reduction in sensitivity (the mean sensitivity decreases from 98% to 95% when the ROI is applied) may be due to one specific seizure in the training set, namely *p103_s1*. This seizure seems to be detected fewer times when the ROI is included in the model, while this behaviour is reversed for the other seizures, that are detected more often with the ROI. This seems to be an isolated behaviour, as far as it could be analysed. The ROI and filter do not seem to affect latency or specificity. However, although the ROI leads to a statistically significant improvement in segmentation error ($p - value = 0.0029$), the filter does not significantly impact this metric ($p - value = 0.1$).

Due to the considerable effect that the thresholding had on the model's performance, it was then important to calibrate these thresholds to their optimal values. *Fig. 6.5* shows the results for the 4 performance metrics (sensitivity, specificity, latency and segmentation RMSE) for adaptive thresholds of between 0 and 10 $\times$ the standard deviation of each OF vector during the calibration period and fixed thresholds of between 0 and 1.6.

In general terms, it seems that the sensitivity is very high for most values of the fixed threshold between 0.4 and 1.2. It seems slightly lower for values higher than 1.2, and significantly lower for values lower than 0.4. The adaptive threshold does not seem to seriously affect these values, but there is a

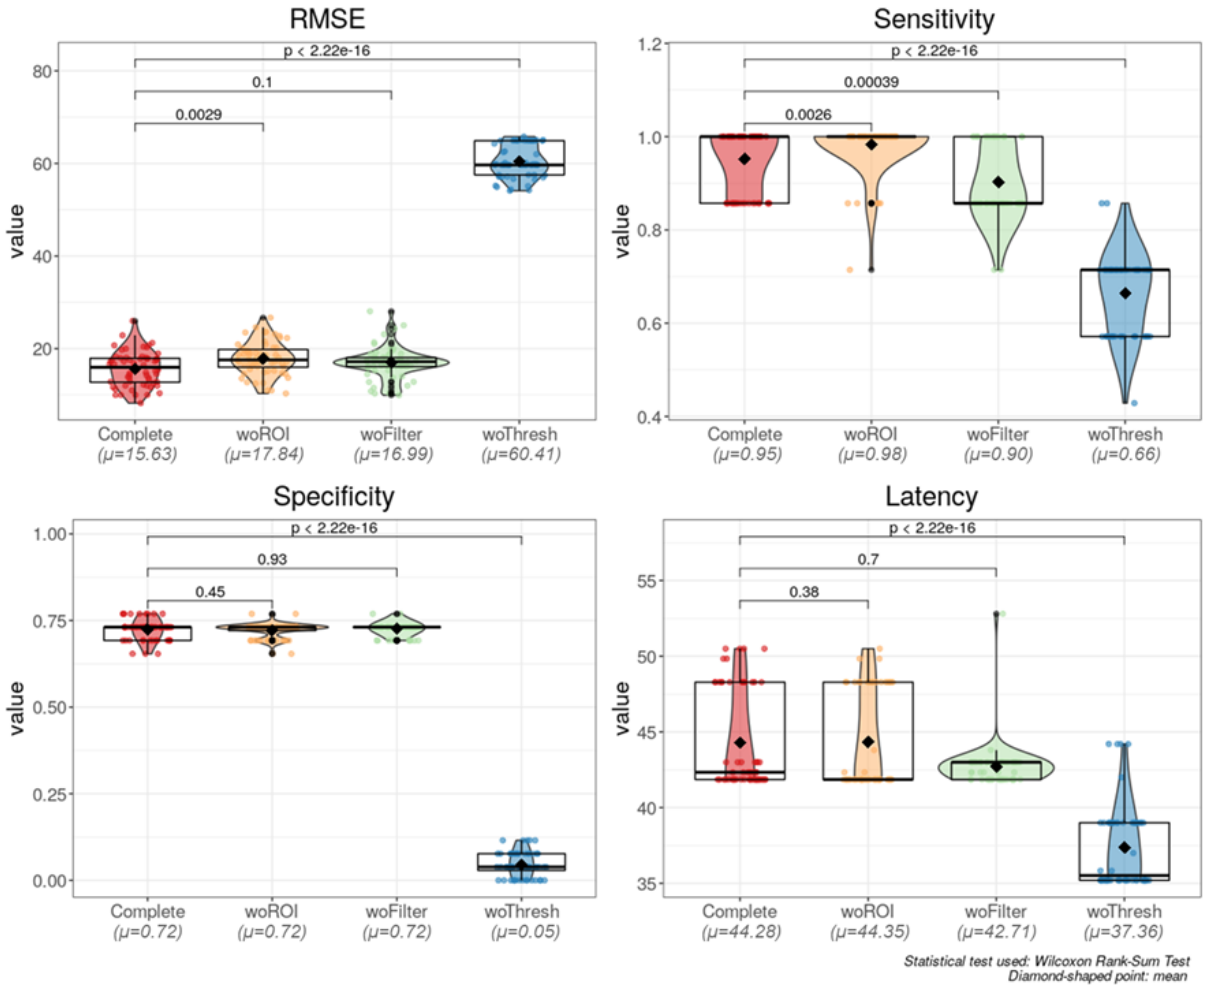

Figure 6.4: Boxplots, Violin plots, and Wilcoxon Rank-Sum test results showing the difference in performance (as measured by segmentation RMSE, sensitivity, specificity, and latency) between a version of the algorithm all noise reduction methods ("Complete") and versions without each individual method (woROI - without ROI, woFilter - without the low-pass filter and woThresh - without the fixed and adaptive thresholds). These distributions were sampled by running the algorithm 60 times for each of the 4 versions of the model and measuring performance metrics. Boxplots indicate the median, interquartile range and variability of the distributions, while violin plots more accurately describe the shapes of the distributions. Above these is the Wilcoxon Rank-Sum Test result for the comparison of the "complete" model with all noise reduction methods and the models without each of the methods. Generally, p-values below 0.05 indicate that the distributions are statistically different.

slightly worse performance for low adaptive thresholds, particularly below 4. Regarding the specificity, however, both thresholds seem to play a significant part, as lower values for both thresholds significantly hamper specificity, and higher values significantly increase it. It seems that a fixed threshold above 0.6 and an adaptive threshold above 4 will yield the best specificity. Latency is also slightly affected by the fixed threshold, as there seems to be about a 10 second difference between lower and higher fixed threshold values, with the former yielding better latency. The segmentation error also seems to be mostly affected by the fixed threshold, with lower threshold values leading to far higher errors. To obtain the lowest segmentation error, the optimal values for this threshold seem to be between 0.8 and 1.2.

In summary, it seems that the fixed threshold has a higher impact on the results. This is most likely because the fixed threshold is applied to every OF vector, and not just the vectors that are most affected

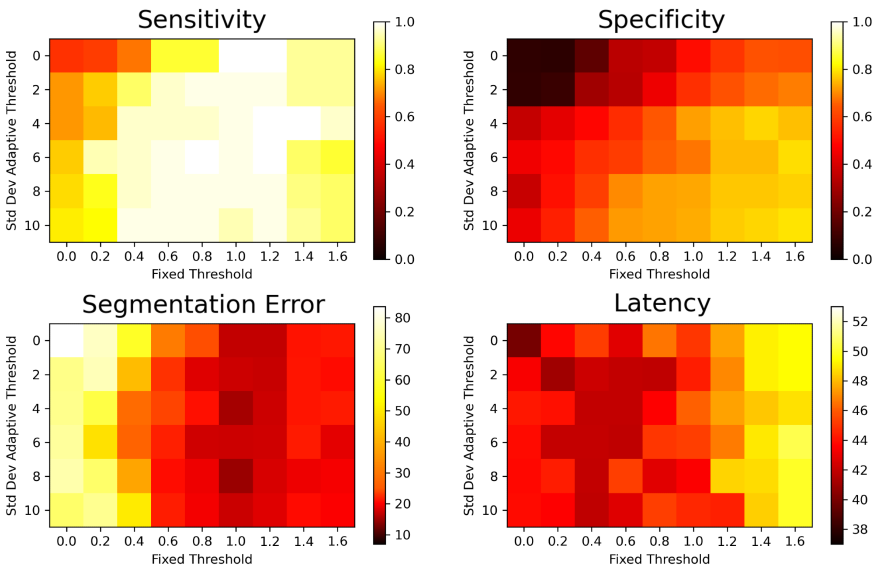

Figure 6.5: The performance of the algorithm for different noise reduction adaptive and fixed thresholds, measured by its sensitivity, specificity, latency (in seconds) and RMSE (in seconds). The model was run 8 times for each threshold value pair, and the mean of these 8 runs was taken.

by the OF-related noise that was detailed in Chapter 5. This is relevant because the OF vectors that are most affected by the OF-related noise are those that correspond to clear discontinuities in brightness (or "edges"), which are more likely to happen in OF vectors corresponding to the side of the patient's bed or furniture in the room, and not the patient themselves. Moreover, the fixed threshold removes low amplitude signals such as, for instance, the jittering of the patient's bed. These low amplitude signals could be isolated by ICA as an independent source and selected as a seizure episode, leading to lower specificity. However, most of these low amplitude occurrences that would be eliminated by the fixed threshold would not ultimately be classified as a seizure in the final algorithm, because the classification step is partly based on the "amount of movement" feature, which excludes movements with a low amplitude. This is not held into account here, since classification is not included in this model, as previously mentioned. Nevertheless, this threshold's ability to exclude these possible false episode selections will most likely be useful to improve the quality of the source and episode selection process. Even so, if the fixed threshold is set too high, it may worsen sensitivity by excluding seizure signals that might have a lower amplitude.

A fixed threshold of 0.8 and an adaptive threshold of 8 seems to be a combination that maximizes sensitivity while obtaining good specificity, segmentation error and latency. These were the final values chosen for these thresholds.

Overall, although the ROI and low-pass filter seem to have a noticeable impact in visual inspections of the signal, when performance metrics are analysed, their impact, although statistically significant, seems fairly small in comparison to the considerable impact of the fixed and adaptive threshold. Nonetheless, these methods do seem to have a positive impact, overall. As such, the ROI and low-pass filter have been kept as a part of the final algorithm. Following a comparison of different threshold values, these have been set at 8 for the adaptive threshold and 0.8 for the fixed threshold.

6.2 Source Selection with Seizure Metrics

The source selection process served two main purposes, as was previously explained: reducing dimensionality and isolating seizure movement. By estimating, selecting and summing together selected independent sources, this process enables the reduction of the number of signals from the large number of OF vectors to a single signal. Regarding seizure movement isolation, PCA isolates a set number of principal components with most of the variability in the signal and ICA identifies independent sources of movement that are then selected with the help of TCF thresholds. This section must therefore ascertain if the time series signals obtained from this process succeed in encapsulating the main movements in the video sequence into univariate signals, and whether the seizure movements are successfully isolated. Non-seizure movements made by other people often have high, narrow peaks with significant lower frequency components, as opposed to the consistent and high frequency clonic movements present in GTCS and other seizures.

For this purpose, it is useful to compare three signals, namely the average of all OF vectors, the sum of all ICA sources (without selecting the best ones) and the sum of the selected sources with lower TCF than the source selection threshold T1. The differences between these three signals will indicate whether TCF was useful for selecting seizure movement, and whether PCA and ICA are better methods for reducing dimensionality into one time series signal than a simple average. *Fig. 6.6* shows these signals for *p101_s3*, as well as two other seizure videos in the training set (*p103_s1* and *p109_s3*) for which this method's impact was particularly noticeable.

There are very noticeable peaks in the average of all OF vectors and in the sum of all sources, particularly in *p103_s1* and *p109_s3*. Because *p101_s3* did not contain many non-seizure movements, and the signal was already fairly clean, this method's impact is reduced. Nevertheless, the isolated signal clearly corresponds to a seizure-like event, with consistent high frequency oscillation. On the other hand, *p103_s1* and *p109_s3* contained significant and prominent non-seizure movement, as there were multiple people in the video sequence running in front of the camera while attempting to help and secure the patient. Because of this, source selection also has a much higher impact for these videos, with high peaks being removed and the oscillatory components corresponding to clonic movement being heightened.

Interestingly, because *p103_s1* and *p109_s3* are videos of focal to bilateral tonic-clonic seizures, they start with focal clonic movement, transition to a tonic phase as generalization occurs, and end in the clonic phase. This clonic-tonic-clonic behaviour is clearly observable in these signals - at the start there is a short clonic phase with low amplitude oscillatory movement, followed by a tonic phase without these clonic oscillations, and then followed by the clonic phase, which is the most noticeable, with noticeable medium amplitude and high frequency movements lasting roughly 30-40 seconds in both cases.

Notably, even the sum of all ICA sources removes some of the non-seizure movements that can be seen in the average of all OF vectors, indicating that computing PCA and ICA and summing the obtained sources leads to better results than taking a simple average of all vectors, most likely because the sources are normalized, leading to high amplitude components like these peaks to not be as prominent.

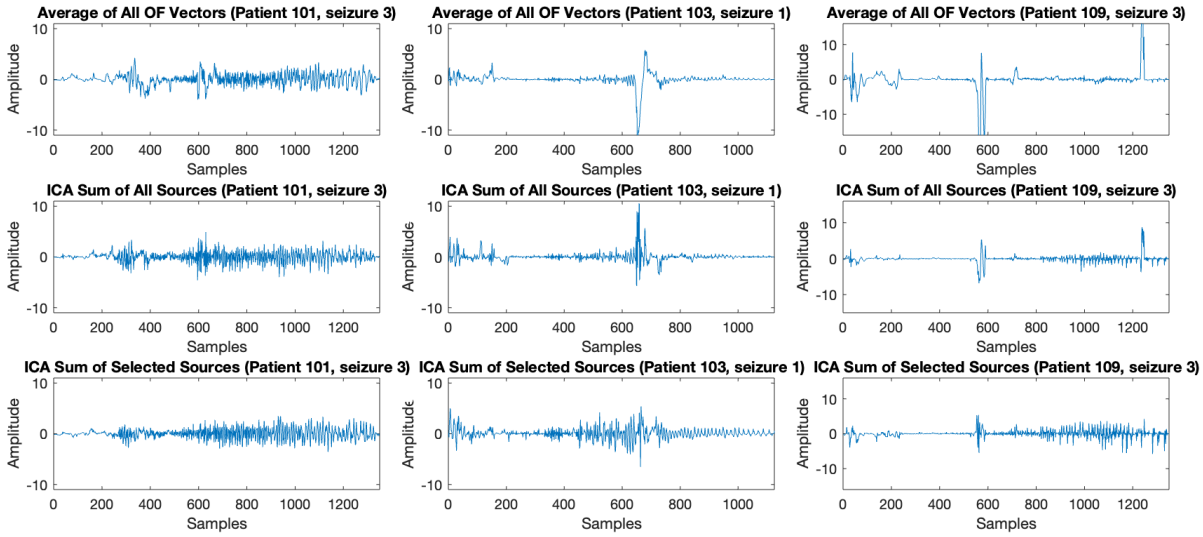

Figure 6.6: The average of all OF vectors (above), the sum of all sources estimated by ICA (in the middle), and the sum of the selected sources with a TCF lower than T1 (below), for *p101_s3*, *p103_s1*, and *p109_s3*. There are very noticeable peaks in the average of all OF vectors and in the sum of all sources. Meanwhile, in the signals deriving from the source selection process using TCF, these artefacts seem significantly reduced, and the oscillatory components seem to be heightened.

To more closely inspect these movements when it concerns their frequency contents, time-frequency analysis was performed on these signals using Short-time Fourier transform (STFT) spectrograms (here showing the power spectral density or PSD), which are shown in *Fig. 6.7*. This is important because one of the features used for classification, *spectral contrast*, measures the power between 2 and 7 Hz in comparison to the power between 0 and 2 Hz. If the isolation of seizure movement was successful, a noted increase in high frequency oscillations and a decrease in lower frequency ones should be observed.

From these spectrograms, it seems clear that the source selection process has had a significant impact in removing non-seizure movement. The aforementioned peaks, which are distinguishable in these spectrograms as vertical spikes, seem to be noticeably attenuated in the spectrograms of the signals that underwent source selection. Conversely, the oscillatory components that correspond to clonic movements, noticeable here as diagonal streaks (because clonic movements have a very defined frequency that decreases towards the end of the seizure) seem to be substantially heightened, meaning that seizure movements have been successfully isolated.

It then becomes important to calibrate the source selection and episode selection thresholds (T1 and T2), as they have a considerable effect on performance. These were calibrated in a similar way to the noise thresholds, by running the model multiple times (5, in this case) on the training set for each threshold combination and evaluating detection and segmentation accuracy on the 4 established performance metrics. Since there are actually 4 thresholds (a source and episode selection threshold for both the first and second segmentation step), this had to be done twice: first for the 45s step, as these thresholds will have a higher impact on sensitivity and specificity, and then for the 9s step. This can be observed in *Fig. 6.8*.

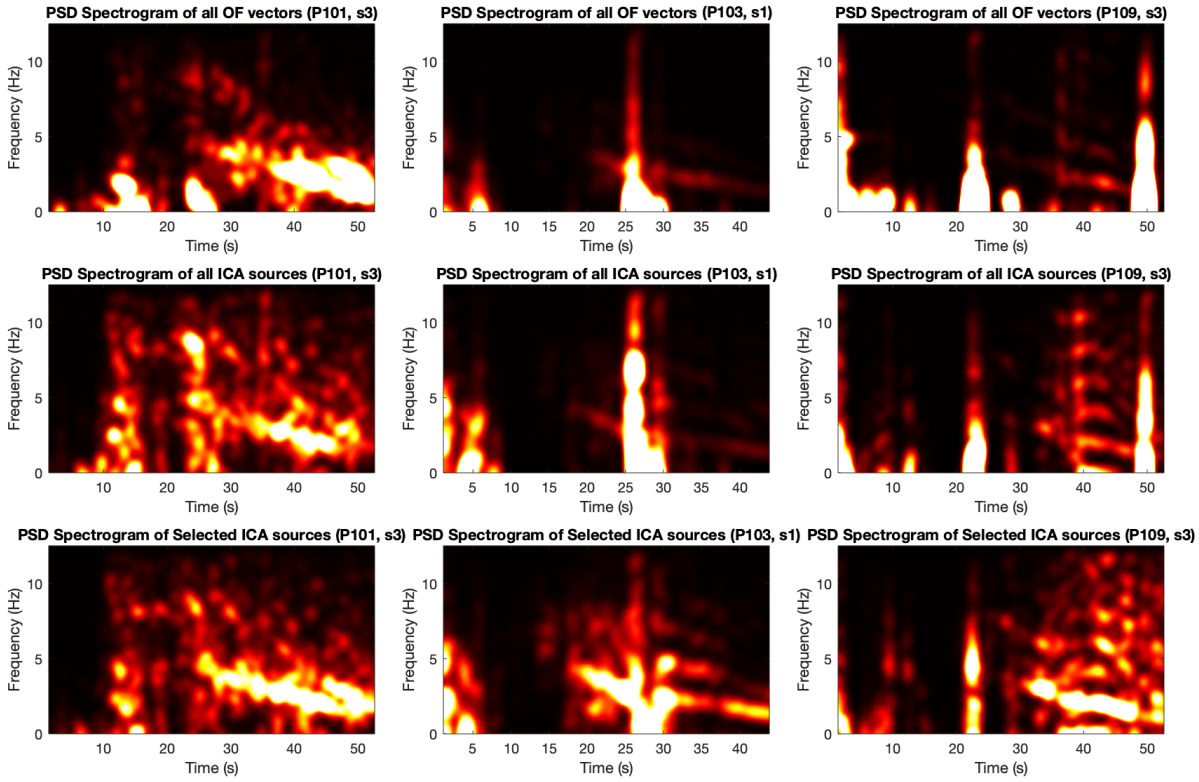

Figure 6.7: The STFT spectrograms showing the PSD of the average of all OF vectors and the sum of sources estimated by ICA, without source selection (middle) and with source selection (below), for *p101_s3*, *p103_s1*, and *p109_s3*. The peaks in the signals caused by non-seizure movements are noticeable here as vertical spikes. These seem to be reduced or almost eliminated in the signals deriving from the source selection process (below) using TCF. The oscillatory components, noticeable here as diagonal streaks, are substantially heightened.

Some patterns are apparent, namely that all 4 performance metrics are affected by both threshold values in very similar ways. While sensitivity and latency are better for threshold values above 2 and 2.5 respectively, specificity and RMSE are better for lower values. Both specificity and RMSE seem to be more affected by the episode selection threshold (T2) than the source selection threshold (T1). For specificity, this indicates that most false positives have fairly consistent sources, but as a whole tend to not be as consistent, leading some episodes to be excluded due to having a TCF that is higher than T2. Because sensitivity is the most important metric (especially since, as mentioned, specificity will be improved by the classifier), it seems natural to choose threshold values that maximize it. Choosing thresholds that are only slightly higher than 2 might be risky, since there could be seizures that don't clear this threshold (it is important to note that the training set is very reduced, only containing 7 seizures). Therefore, threshold values should be at least somewhat larger than 2. Since specificity and RMSE are more dependent on T2, it makes sense to choose a T2 that is lower than T1. The chosen thresholds were $T1 = 4$ and $T2 = 3.5$.

Fig. D.3 in Appendix D shows the results of the calibration of T1 and T2 for the second (9 second) segmentation step. Sensitivity and specificity are not affected by this step, which is to be expected, as it is very rare for a 45s segment to be selected and for all of its 9s parts to not be selected, which is

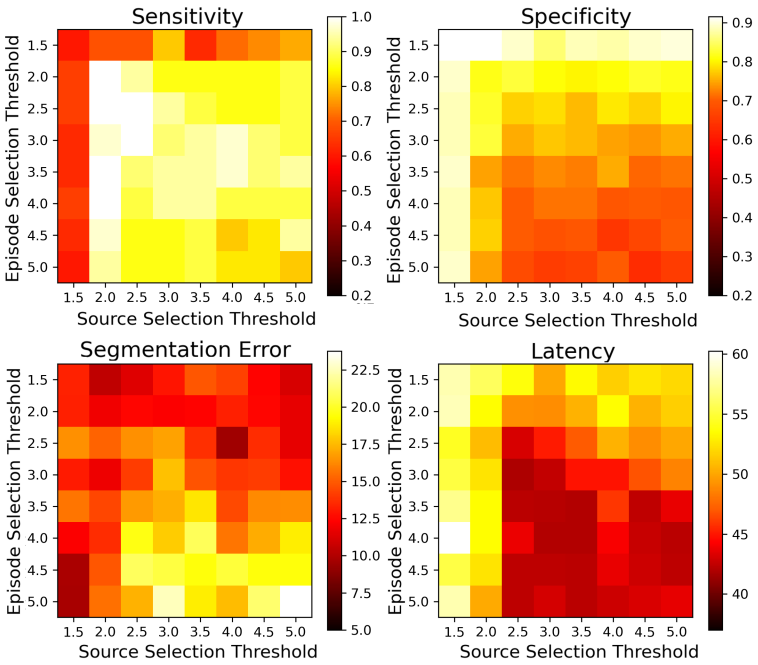

Figure 6.8: The performance of the algorithm for different source and episode selection thresholds for the first (45s) segmentation step, measured by its sensitivity, specificity, latency (in seconds) and RMSE (in seconds). The model was run 5 times for each threshold value pair, and the mean of these 5 runs was taken.

the only way for these thresholds to affect sensitivity. Mostly, this step will contribute to the accuracy of the segmentation. However, as it can be observed in *Fig. D.3*, although the only affected parameter seems to be the segmentation RMSE, with slightly higher values for lower thresholds, the impact seems to be rather small. However, it must be noted that it is possible that once the classification step is introduced, these thresholds could have an impact, as better segmentation could affect the values of the classification features. In any case, it makes sense to choose values for these thresholds that contribute to the lowest possible segmentation error. The chosen values were $T2 = 11$ and $T1 = 4$. T2 was set at 11 because 9s segments are likely to not be as consistent as 45s segments, and as such it makes sense for this threshold to be more lenient. T1, the source selection threshold, was set at 4, because, without there being any evidence for or against it, it is the same value as the source selection threshold in the previous step.

In summary, it is apparent that PCA and ICA source selection have succeeded in reducing the dimensionality of the signal, and TCF has been successful in isolating seizure-like movement and eliminating other sources of movement. However, this can only be determined objectively by evaluating the performance of the final algorithm, with the classification step, on the validation set.

6.3 Algorithm Evaluation

Perhaps the single most important component of this work was to deliver the most accurate seizure detection algorithm possible. As it was previously explored, the metrics that were chosen to evaluate this

performance were sensitivity, specificity, segmentation error (RMSE) and latency. Sensitivity, specificity and latency are very important to patients and caregivers (see Chapter 3), while segmentation error is an important metric of the quality of the isolation of seizure episodes, which could be critical in developing tools to help health professionals in the lengthy process of annotating seizures in long-term video-EEG sessions.

6.3.1 Detection Accuracy

To assess the seizure detection accuracy, the testing set was first used to compare the different machine learning classification methods detailed in chapter 5 (non-linear SVM with $C = 10$ and $\gamma = 0.2$, GPC with RBF kernel, and MLP with "ReLU" activation and $\alpha = 0.1$). Then, the best performing classifier was chosen, and its detection and segmentation performance was evaluated using the previously mentioned performance metrics.

The comparison between classification methods will be discussed first. For this purpose, leave-one-subject-out cross-validation (LOSOCV) was used, which increases the size of the data by training the algorithm for each seizure video with every video in the training and testing set except the specific video which will be classified. For small datasets such as the one at hand, this is a good way of obtaining better estimations for the performance of a classifier. In order to further increase the dataset, the model was run twice for each video, such that, by taking advantage of the randomness introduced by ICA and PCA, each video could contribute two different data points, doubling the size of the dataset. It is crucial to stress that videos were **not trained using data points generated by themselves** (as this would naturally lead to overly optimistic performance results).

An example of a resulting dataset obtained for *p388_s1*, as well as the corresponding decision functions generated by the three classifiers during the cross-validation process can be observed in *Fig. 6.9*. It seems that all three classifiers have generated similar decision functions for this particular dataset, with different levels of bias and variance. Visually, it seems that the SVM has the lowest variance, while the MLP has the highest. However, to analytically compare their performance, their classification thresholds were varied in order to generate Receiver Operating Characteristic (ROC) curves, for which the Area Under Curve (AUC) could be calculated. **These ROC curves do not simply represent the performance of the classifiers, but of the full algorithm.** This is visible in *Fig. 6.10*.

All three methods appear to have very high detection accuracy, with AUC's of 94.9% for the SVM, 96.3% for the GP, and 96.5% for the MLP. The best performing classifier seems to be the MLP, but the difference between them is fairly small. Crucially, because this represents the performance of the complete algorithm, and not only the classifiers, it can be said that **the algorithm delivers excellent detection performance with any of the three methods**. In order to decide the optimal threshold for these classifiers, the Equal Error Rate was computed by running the seizure detection algorithm with different classification thresholds for all three classifiers. This is the point at which specificity and sensitivity intersect, balancing false positives and false negatives. The algorithm was run 10 times per threshold, and the mean sensitivity and specificity for each were taken and plotted *Fig. 6.11*.

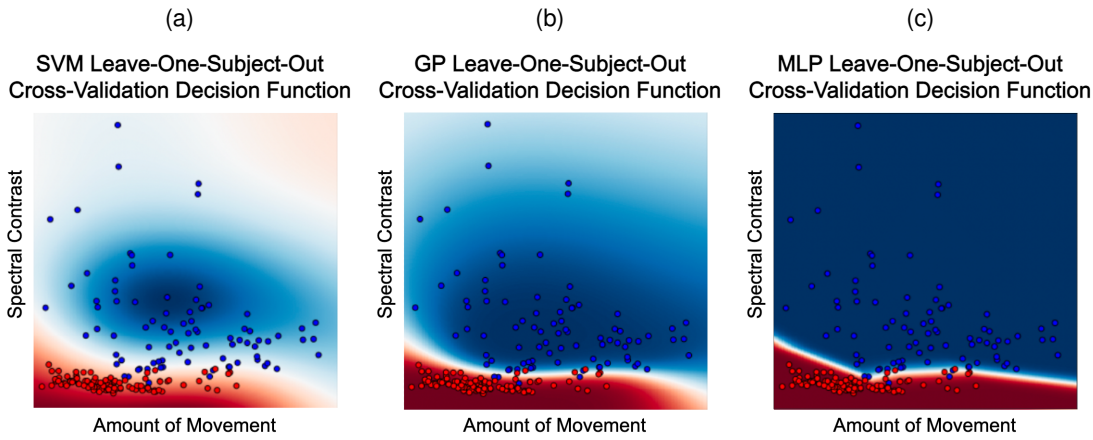

Figure 6.9: (a) – The decision function generated by the SVM classifier and the set of data points used for the classification step for *p388_s1*. The set was obtained with leave-one-subject-out cross-validation (i.e. with all data points except those generated by *p388_s1*).(b) – The decision function generated by the GP classifier and the set of data points used for the classification step for *p388_s1*. The set was obtained with leave-one-subject-out cross-validation (i.e. with all data points except those generated by *p388_s1*). (c) – The decision function generated by the MLP classifier and the set of data points used for the classification step for *p388_s1*. The set was obtained with leave-one-subject-out cross-validation (i.e. with all data points except those generated by *p388_s1*).

For all three classifiers, the threshold which positively classified a test data point if the prediction probability was above 50% (which is 0 for the SVM, and 0.5 for the GPC and MLP, due to the different methods that these classifiers have of calculating detection probability), would lead to nearly 100% sensitivity but only between 80% and 90% specificity. The threshold for the equal error rate was higher for all three classifiers, indicating that a "stricter" threshold, which classifies fewer points as positive, led to a better balance between sensitivity and specificity. For the SVM, this equal error rate was 7.7% (meaning a 92.3% sensitivity and specificity were achieved) at a threshold of 0.85 (with 0 being the threshold at which a point is classified as positive if the classification probability is 50%, for the SVM). For the GPC and MLP, the threshold that yielded equal specificity and sensitivity was the same, at 0.8 (meaning that only points with an outputted prediction probability of over 80% would be classified as positive). This generated a sensitivity and specificity of 94.0% and 94.8%, respectively, at equal error rates of 6.0% and 5.2%.

In summary, it seems that **this seizure detection algorithm delivers excellent accuracy, with the best classification method, the Multilayer Perceptron neural network, delivering an AUC of 96.5%** and an equal error rate of 5.2% (94.8% sensitivity and specificity) at a classification threshold of 0.8. This is **superior to most methods and devices in the state of the art, and above what patients claim is sufficient sensitivity and specificity**. Nonetheless, the false alarm rate (FAR) per hour could not be reliably estimated for such short videos with unrepresentative amounts of movement. Therefore, it is unclear whether this algorithm would achieve an FAR of under 1 false alarm per week, which patients and caregivers state is an acceptable amount (especially for patients with a low number of seizures per week). Still, these results are very promising and a good indicator that the combination of optical flow, ICA, source selection with TCF and MLP classification can accurately be used to detect seizures in

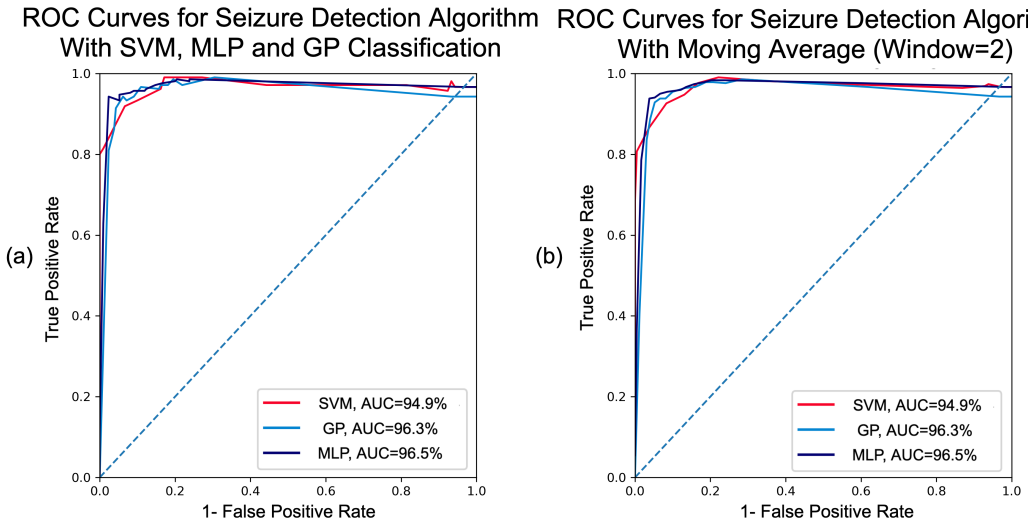

Figure 6.10: (a) – Receiver Operating Characteristic (ROC) curves for the seizure detection algorithm using the three tested classifiers: a non-linear Support Vector Machine with $C = 10$ and $\gamma = 0.2$, a Gaussian Process Classifier with an RBF kernel, and a Multilayer Perceptron with "ReLU" activation and $\alpha = 0.1$. (b) – The same curves are shown with a moving average filter (window=2) to increase interpretability.

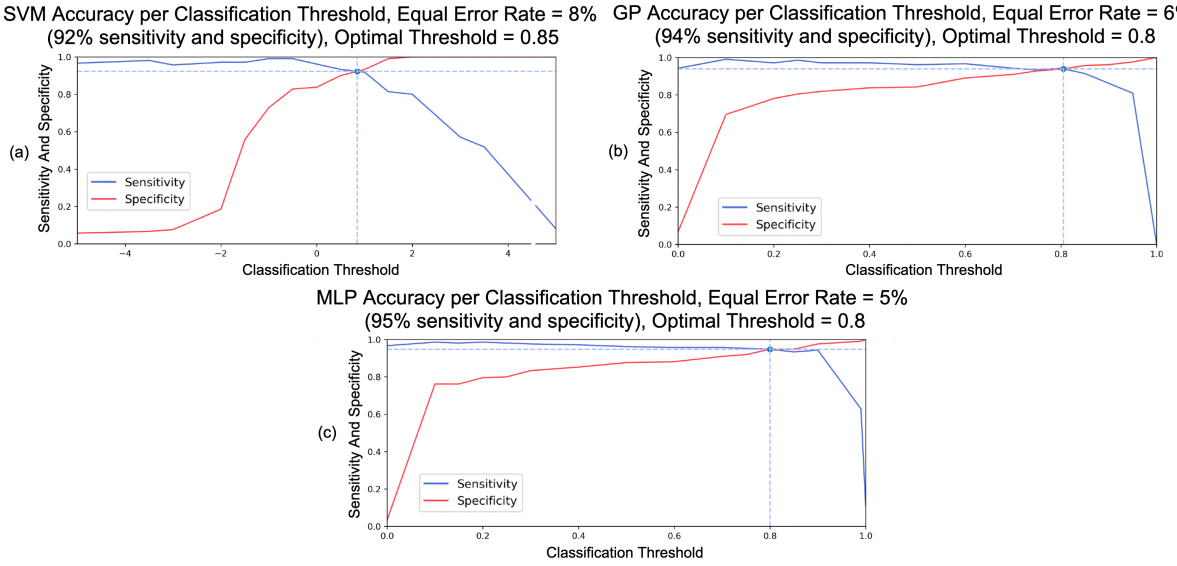

Figure 6.11: (a) – The sensitivity and specificity of the seizure detection algorithm with SVM classification generated by running it with different classification thresholds. The algorithm was run 10 times per threshold, and the mean sensitivity and specificity for each were taken. The equal error rate, i.e. the point at which the sensitivity and specificity intersect, is also visible. For the SVM, it is 7.7%, meaning that the algorithm achieves 92% sensitivity and 92.3% specificity with this classifier, at the equal error rate. (b) – The sensitivity and specificity of the seizure detection algorithm with GPC generated by running it with different classification thresholds. The equal error rate for the GPC is 6.0%, meaning it achieves 94.0% sensitivity and specificity. (c) – The sensitivity and specificity of the seizure detection algorithm with MLP classification generated by running it with different classification thresholds. The equal error rate for the MLP is 5.2%, meaning it achieves 94.8% sensitivity and specificity.

real-time from video.

6.3.2 Segmentation Accuracy

The algorithm with the best performing classifier (MLP) was tested (by running it 20 times) with the optimal threshold to assess the quality of the segmentation of seizure episodes in comparison to the seizure start and ending times annotated by health professionals. **The average segmentation error was 17.7 seconds (and the segmentation RMSE was 23.5 seconds), which corresponds to an error of $\pm$ 8.85 seconds between the detected and annotated beginning and ending of seizures**, which is very similar to the length of the 9s segmentation step, meaning most seizures were segmented nearly as accurately as possible. The average latency was 52.9 seconds, which is somewhat sub-par, when comparing to the current methods in the state of the art, and is a metric that must be improved in the future.

To further ascertain the quality of this segmentation, the algorithm was run 20 times and the true and false positive segmented seizure episodes were plotted alongside the seizure episodes annotated by Dr. Francisca Sá of Hospital Egas Moniz in *Fig. 6.12*.

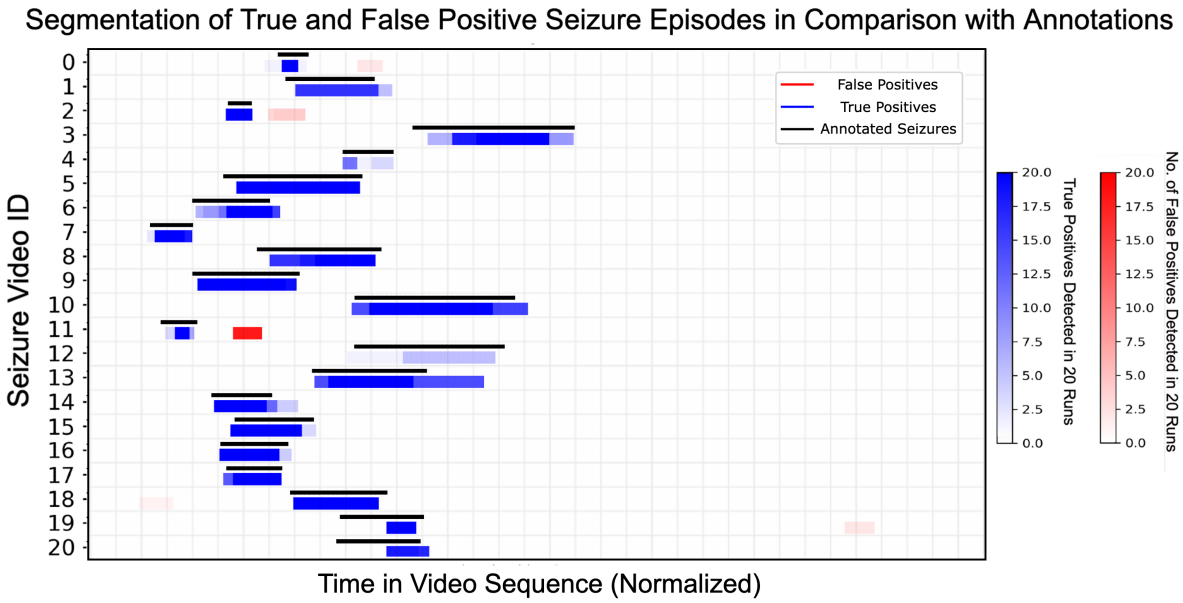

Figure 6.12: The true positive (blue) and false positive (red) segmented seizure episodes, compared to the seizure episodes annotated by Dr. Francisca Sá of Hospital Egas Moniz (black). The algorithm was run 20 times and the resulting segmentations were plotted. For the true and false positive detections, strength of the color indicates the number of times that specific segmentation occurred in the 20 runs, with lighter colors a segmentation that only occurred a few times, and starker colors indicating very common segmentations. The x-axis corresponds to the time in the video sequence, normalized so that all sequences have the same length.

Most segmented seizure episodes were detected in all 20 runs, with seizures #4 and #12 being detected slightly less frequently than the others. When compared to the annotated seizure episodes (in black), it can be noted that the segmentation was very accurate for the majority of the seizures in the testing set, with high similarity between manually annotated and automatically detected seizure episodes. False alarms were detected in five seizures, although very infrequently for most. An exception to this occurred in seizure video #11, where a false alarm was detected in multiple runs in the post-ictal

period (as indicated by the starker shade of red), shortly after the seizure had ended. This was caused by the patient moving erratically after the seizure due to them being disoriented following the seizure. False alarms were detected in the post-ictal periods of two more seizure videos, while in two others false alarms were detected elsewhere in the video sequences. Nonetheless, these false alarms were not detected in the vast majority of runs, as indicated by their lighter shade of red.

This high level of detection and segmentation accuracy enables this algorithm to be used not only for real-time seizure detection in an ambulatory setting, but also for automatic registration and segmentation of seizure episodes to help health professionals annotate seizures in long-term video-EEG sessions.

Chapter 7

Conclusions and Outlook

7.1 Concluding Remarks

This project's main ambition was to develop a device that could address some of the various challenges that epilepsy patients and caregivers face. While almost all patients and caregivers are interested in SDDs and believe they could lower their stress levels and improve their quality of life, most don't own one because they are either unsure about their reliability or because they simply can't afford one. Seizure detection devices on the market cost upwards of 200$, which for a lot of lower income families is unfeasible, especially when they are unfamiliar with the associated benefits.

The benefits of seizure detection are extensive, especially for the roughly 30% of patients with refractory epilepsy. On average, patients are only aware of half of the seizures they have, which can lead to poor diagnosis accuracy, leading patients to have their medication poorly adjusted. As such, seizure detection and registration devices, besides giving patients and caregivers much needed peace of mind, can also shorten and simplify the diagnosis process. Furthermore, such devices can increase diagnostic accuracy by helping classify seizure types and determine seizure frequency. SDDs can also reduce the need for lengthy and uncomfortable long-term video-EEG sessions, and help adjust medication regimens, which can contribute towards achieving lasting remission. Potentially, seizure detection devices could also enable fast responses to serious seizure episodes, reducing the probability of SUDEP.

Knowing these benefits, the vast literature covering seizure detection methods, and the large number of devices on the market, it is somewhat disquieting that most, if not all devices on the market are unaffordable or unreliable. Moreover, it is astonishing that, although caregivers claim to highly value video-based monitoring, **there is no video-based seizure detection device currently commercially available**. Video-based seizure monitoring and detection would be particularly valuable for caregivers of children with epilepsy, as they could avoid needing to sleep with their children or having their quality of sleep severely worsened by constantly worrying about potential seizures. Additionally, video monitoring also contributes toward affordability, as no special sensors are needed, and small cameras compatible with single-board computers or micro-controllers are very affordable. Developing an affordable seizure detection device that used this modality and did not compromise on accuracy, while ensuring accessi-

bility, unobtrusiveness and privacy by design was the guiding light of this project.

The concept that guided the development of this device was that of *"invisibles"*, which are devices that fully integrate into the user's day-to-day life without being noticeable or obtrusive. This idea is innovative in that it looks towards a future where humans are increasingly integrated with technology in ways that enable the improvement of health and quality of life in a way that prioritizes user well-being and privacy. This lead to the integration of this device into a light fixture, from which it got its name, **Lampsy**. This integration contributed towards the unobtrusiveness of the device, which was essential in order to facilitate ambulatory monitoring and day-to-day use, as well as avoid the possible stigma that patients with epilepsy can suffer as a result of someone identifying their condition.

A prototype was built for Lampsy using a Raspberry Pi 4, an RPI-compatible camera and a small display. It was integrated within a floor lamp and placed in Hospital de Santa Maria for raw data acquisition. To ensure that this was as easy to use as possible, two buttons were installed, a power (ON/OFF) button and a record (REC) button, which could stop and restart the video recording (for example if the patient needed a moment of privacy). The prototype was also programmed to automatically stop and restart the recording in order to reduce file size and enable data transmission. The acquired video data was transmitted automatically to a secure server in IT using a 4G router, after being converted into .mp4 files and then compressed into .zip files. In order to protect patient privacy, data transfer was done via SFTP, which is a method of file transfer based on SSH that uses a secure and encrypted channel for transmission. As of November 15th, this device had recorded 130 consecutive hours of video footage at HSM without any issues.

A seizure detection and segmentation algorithm was developed, based on the calculation of Optical Flow, followed by a novel noise reduction method based on a calibration period, the calculation of PCA and ICA, the selection of ICA sources using a novel metric named the *Temporal Consistency Factor* (TCF), the segmentation of possible seizure episodes using this source selection and TCF, and the classification of these episodes using any of three classification methods: Support Vector Machine classification, Gaussian Process classification and Multilayer Perceptron classification.

It was always crucial to ensure that any method would ensure patient privacy by design. Calculating optical flow vectors stores the video data in a completely unidentifiable way, which assures that even if there was a security breach, this data could not be used to identify the user.

Leave-one-subject-out cross-validation was used to evaluate the performance of this algorithm. From the results explored in Chapter 6, it seems clear that, as far as it could be tested with the available data, **the developed model manages to reach very high accuracy levels. With the best classification method, a Multilayer Perceptron, this seizure detection algorithm achieves an AUC of 96.5% and a sensitivity and specificity of 94.8%** (an equal error rate of 5.2%). While the false alarm rate per 24h could not be accurately tested, the specificity remains high, even for videos with large amounts of movement. Indeed, this model not only enables night-time monitoring, but also day-time and continuous monitoring, by the virtue of the success of the source selection process. When comparing these results with the state of the art, it is fair to say that it stands on par or above the majority of seizure detection algorithms. The fact that it also enables seizure segmentation, by detecting the beginning and ending of

the seizure with fairly low errors (at an average of $\pm$ 8.85 seconds) is a significant advantage, whether it be to help health professionals annotate seizures in video-EEG sessions or to help seizure logging to improve diagnosis accuracy and facilitate this process. A performance metric that falls short of the state of the art is the latency, which at an average 52.9s is higher than most published methods. Nevertheless, as detailed in Chapter 3, some patients and health professionals do accept a latency of under 1 minute.

In conclusion, the results obtained from this work are very promising. The integration into a light fixture and the choice of a Raspberry Pi, as well as the choice of modality, contribute towards keeping this device unobtrusive and affordable, which was essential. The principles of privacy by design were followed throughout this process, both in the data transmission and in the algorithm. Furthermore, the developed algorithm yielded very promising results, which in large part surpassed the current state of the art.

7.2 Limitations and Future Work

The prototype and algorithm developed within this project constitute the first steps towards developing a fully fledged seizure detection device, named **Lampsy**, for which a patent application has been filed and approved by Instituto Superior Técnico, with development being expected to continue in the following years. Nevertheless, further steps need to be completed before this device could be made available to patients and caregivers.

First of all, it would be important to test this device and algorithm together. Although the prototype was installed at HSM, it did not yet include the algorithm, only being there to record raw data. By placing the finished algorithm in the prototype, the algorithm could be tested in real-time, which would enable latency to be calculated in a real world setting, instead of simply assuming that computations would be essentially instant. Although could effectively be the case for most normal computers, it likely won't be the case for a Raspberry Pi, which is why it is also important for the RPI to be benchmarked in regards to its ability to efficiently run the algorithm.

Although the size of the training and testing sets remains a limitation regarding the obtained results, the number of videos that were tested remains in line with what is used for most state of the art studies. Nonetheless, with more data and longer videos, the performance of this algorithm could be more accurately and reliably estimated. More data could also enable the use of more complex machine learning methods to, for example, distinguish focal to bilateral tonic-clonic seizures from generalized tonic-clonic seizures, which would provide important information towards diagnosis, or implement deep-learning OF methods. However, it is not trivial to acquire large numbers of annotated seizure videos, as there are complex processes regarding data protection laws that must be diligently followed.

Moreover, there are further techniques that could be integrated into the algorithm. A robust person-detection algorithm could be used to select only the specific OF vectors that corresponded to the patient, improving the isolation of seizure movement. It would also be interesting to find and integrate more features into the classification, which could improve the algorithm's accuracy, and use regularization methods such as LASSO (least absolute shrinkage and selection operator) [120] to perform feature

selection.

Further work could also include building a more advanced prototype. The ideal version of Lampsy would not be merely integrated into a light fixture, but into a lightbulb, maximizing its unobtrusiveness and integration into the user's daily life. This way, Lampsy could serve as both a regular light source and a seizure detector, and the infrared LEDs could be integrated into the lightbulb for a night-time mode. This version of Lampsy could use a Raspberry Pi Zero, which is small and affordable[1], to process video on-site, or an even smaller microcontroller such as a Raspberry Pi Pico, which starts at $4[2] and could be used to simply record and transmit video which would be processed elsewhere. The data transmission would be done via power-line communication, which could enable the transmission of data over the electrical grid. This way, a microcontroller could be used to process and transmit video and this video could be processed in the cloud or on a separate device connected to the power grid. *Fig. 7.1a* shows a rendering of how this device could look like from the outside, and *Fig. 7.1b* shows the possible components that would be included, such as a small camera, a microprocessor, a power-line transcreiver, IR LEDs, and components present in normal lightblubs such as an AC/DC converter and a thermal conduction plate.

(a) (b)

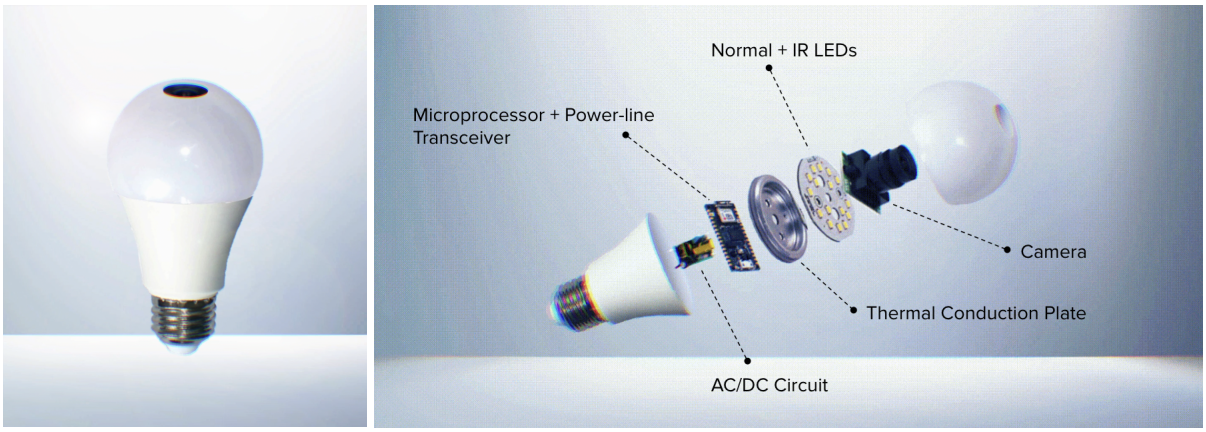

Figure 7.1: (a) – Rendering of Lampsy as a final product, integrated within a lightbulb. (b) – The possible final components for Lampsy, which include a lightbulb's normal components, such as LEDs, a thermal conduction plate and an AC/DC circuit, as well as a microcontroller, a power-line transcreiver, a camera and IR LEDs.

In summary, although Lampsy is still in development, the promising results shown in this work shine a bright light into its future, by setting the base for highly accurate seizure detection without being prohibitively expensive or affecting day-to-day life. Hopefully, in the future, Lampsy can play a part in improving the quality of life of patients and caregivers who suffer from one of the most common and often debilitating neurological conditions, and provide them with the peace of mind that they so desperately need.

[1] From Mauser's website: https://mauser.pt/catalog/product_info.php?products_id=096-4559, last accessed on 11/11/2021
[2] From the Raspberry Pi® website: https://www.raspberrypi.com/products/raspberry-pi-pico/, last accessed on 31/10/2021

Bibliography

[1] R. B. D. et al. *Bradley's neurology in clinical practice*, volume 2, chapter III.G 101 – Epilepsies, pages 1563–1614. Elsevier, London, 2016. ISBN 0323287832.

[2] Epilepsy: a public health imperative. Technical report, World Health Organization, Geneva, 2019.

[3] C. H. et al. Practice guideline summary: Sudden unexpected death in epilepsy incidence rates and risk factors. *Neurology*, 88(17):1674–1680, Apr 2017. doi: 10.1212/wnl.0000000000003685. URL https://doi.org/10.1212/wnl.0000000000003685.

[4] R. G. Beran. SUDEP—to discuss or not discuss: that is the question. *The Lancet Neurology*, 5 (6):464–465, Jun 2006. doi: 10.1016/s1474-4422(06)70449-5. URL https://doi.org/10.1016/s1474-4422(06)70449-5.

[5] M. J. Brodie. Outcomes in newly diagnosed epilepsy in adolescents and adults: Insights across a generation in Scotland. *Seizure*, 44:206–210, Jan 2017. doi: 10.1016/j.seizure.2016.08.010. URL https://doi.org/10.1016/j.seizure.2016.08.010.

[6] J. L. Jameson, A. Fauci, D. Kasper, S. Hauser, D. Longo, and J. Loscalzo. *Harrison's principles of internal medicine*, chapter 14 – 184 – Seizures and Epilepsy, pages 987–999. McGraw-Hill Education, New York, 2018. ISBN 9781259644030.

[7] M. P. Kerr. The impact of epilepsy on patients' lives. *Acta Neurologica Scandinavica*, 126:1–9, Oct 2012. doi: 10.1111/ane.12014. URL https://doi.org/10.1111/ane.12014.

[8] P. Kwan and M. J. Brodie. Early identification of refractory epilepsy. *The New England Journal of Medicine*, 342(5):314–319, Feb 2000. doi: 10.1056/nejm200002033420503. URL https://doi.org/10.1056/nejm200002033420503.

[9] S. Spencer and L. Huh. Outcomes of epilepsy surgery in adults and children. *The Lancet Neurology*, 7(6):525–537, Jun 2008. doi: 10.1016/s1474-4422(08)70109-1. URL https://doi.org/10.1016/s1474-4422(08)70109-1.

[10] V. M. et al. Long-term EEG in adults: Sleep-deprived EEG (SDE), ambulatory EEG (amb-EEG) and long-term video-EEG recording (LTVER). *Neurophysiologie Clinique/Clinical Neurophysiology*, 45(1):47–64, Mar 2015. doi: 10.1016/j.neucli.2014.11.004. URL https://doi.org/10.1016/j.neucli.2014.11.004.

[11] J. W. Sander. The epidemiology of epilepsy revisited. *Current Opinion in Neurology*, 16(2):165–170, Apr 2003. doi: 10.1097/00019052-200304000-00008. URL https://doi.org/10.1097/00019052-200304000-00008.

[12] C. E. Elger and C. Hoppe. Diagnostic challenges in epilepsy: seizure under-reporting and seizure detection. 17(3):279–288, Mar 2018. doi: 10.1016/s1474-4422(18)30038-3. URL https://doi.org/10.1016/s1474-4422(18)30038-3.

[13] G. A. Baker, A. Jacoby, D. Buck, C. Stalgis, and D. Monnet. Quality of life of people with epilepsy: A european study. *Epilepsia*, 38(3):353–362, Mar 1997. doi: 10.1111/j.1528-1157.1997.tb01128.x. URL https://doi.org/10.1111/j.1528-1157.1997.tb01128.x.

[14] A. C. W.-G. et al. Quality of life and burden in caregivers of patients with epilepsy. *Journal of Neuroscience Nursing*, 39(6):354–360, Dec 2007. doi: 10.1097/01376517-200712000-00006. URL https://doi.org/10.1097/01376517-200712000-00006.

[15] J. D. C. et al. Assessing the impact of caring for a child with dravet syndrome: Results of a caregiver survey. *Epilepsy And Behaviour*, 80:152–156, Mar 2018. doi: 10.1016/j.yebeh.2018.01.003. URL https://doi.org/10.1016/j.yebeh.2018.01.003.

[16] A. D. P. et al. Patient-centered design criteria for wearable seizure detection devices. *Epilepsy Behaviour*, 64:116–121, Nov 2016. doi: 10.1016/j.yebeh.2016.09.012. URL https://doi.org/10.1016/j.yebeh.2016.09.012.

[17] D. F. T. Quiroga, J. W. Britton, and E. C. Wirrell. Patient and caregiver view on seizure detection devices: A survey study. *Seizure*, 41:179–181, Oct 2016. doi: 10.1016/j.seizure.2016.08.004. URL https://doi.org/10.1016/j.seizure.2016.08.004.

[18] E. B. et al. Wearable technology in epilepsy: The views of patients, caregivers, and healthcare professionals. *Epilepsy Behaviour*, 85:141–149, Aug 2018. doi: 10.1016/j.yebeh.2018.05.044. URL https://doi.org/10.1016/j.yebeh.2018.05.044.

[19] C. H. et al. Novel techniques for automated seizure registration: Patients' wants and needs. *Epilepsy Behaviour*, 52:1–7, Nov 2015. doi: 10.1016/j.yebeh.2015.08.006. URL https://doi.org/10.1016/j.yebeh.2015.08.006.

[20] E. Bruno, P. F. Viana, M. R. Sperling, and M. P. Richardson. Seizure detection at home: Do devices on the market match the needs of people living with epilepsy and their caregivers? *Epilepsia*, 61 (S1), May 2020. doi: 10.1111/epi.16521. URL https://doi.org/10.1111/epi.16521.

[21] T. Herrera-Fortin, E. B. Assi, M.-P. Gagnon, and D. K. Nguyen. Seizure detection devices: A survey of needs and preferences of patients and caregivers. *Epilepsy Behaviour*, 114:107607, Jan 2021. doi: 10.1016/j.yebeh.2020.107607. URL https://doi.org/10.1016/j.yebeh.2020.107607.

[22] S. K. S. et al. Patients' experience of wearing multimodal sensor devices intended to detect epileptic seizures: A qualitative analysis. *Epilepsy Behaviour*, 102:106717, Jan 2020. doi: 10.1016/j.yebeh.2019.106717. URL https://doi.org/10.1016/j.yebeh.2019.106717.

[23] A. U.-C. et al. Automated seizure detection systems and their effectiveness for each type of seizure. *Seizure*, 40:88–101, Aug 2016. doi: 10.1016/j.seizure.2016.06.008. URL https://doi.org/10.1016/j.seizure.2016.06.008.

[24] J. van Andel, R. D. Thijs, A. de Weerd, J. Arends, and F. Leijten. Non-EEG based ambulatory seizure detection designed for home use: What is available and how will it influence epilepsy care? *Epilepsy & Behavior*, 57:82–89, Apr 2016. doi: 10.1016/j.yebeh.2016.01.003. URL https://doi.org/10.1016/j.yebeh.2016.01.003.

[25] A. van Westrhenen, T. Souhoka, M. E. Ballieux, and R. D. Thijs. Seizure detection devices: Exploring caregivers' needs and wishes. *Epilepsy Behaviour*, 116:107723, Mar 2021. doi: 10.1016/j.yebeh.2020.107723. URL https://doi.org/10.1016/j.yebeh.2020.107723.

[26] H. Plácido da Silva. *Biomedical Sensors as Invisible Doctors*, pages 322–329. 01 2019. ISBN 978-3-9504607-3-5.

[27] Z. K. et al. Passive monitoring at home: A pilot study in parkinson disease. *Digital Biomarkers*, 3 (1):22–30, Apr. 2019. doi: 10.1159/000498922. URL https://doi.org/10.1159/000498922.

[28] P. Jallon, M. Goumaz, C. Haenggeli, and A. Morabia. Incidence of first epileptic seizures in the canton of Geneva, Switzerland. *Epilepsia*, 38(5):547–552, May 1997. doi: 10.1111/j.1528-1157.1997.tb01139.x. URL https://doi.org/10.1111/j.1528-1157.1997.tb01139.x.

[29] A. T. Berg and S. Shinnar. The risk of seizure recurrence following a first unprovoked seizure: A quantitative review. *Neurology*, 41(7):965–965, July 1991. doi: 10.1212/wnl.41.7.965. URL https://doi.org/10.1212/wnl.41.7.965.

[30] W. A. H. et al. Risk of recurrent seizures after two unprovoked seizures. *The New England Journal of Medicine*, 338(7):429–434, Feb 1998. doi: 10.1056/nejm199802123380704. URL https://doi.org/10.1056/nejm199802123380704.

[31] R. S. F. et al. ILAE official report: A practical clinical definition of epilepsy. *Epilepsia*, 55(4):475–482, Apr 2014. doi: 10.1111/epi.12550. URL https://doi.org/10.1111/epi.12550.

[32] E. B. et al. Global, regional, and national burden of epilepsy, 1990–2016: a systematic analysis for the global burden of disease study 2016. 18(4):357–375, Apr 2019. doi: 10.1016/s1474-4422(18)30454-x. URL https://doi.org/10.1016/s1474-4422(18)30454-x.

[33] E. K. et al. *Principles of neural science*, chapter VII 50 – Seizures and Epilepsy, pages 1116–1139. McGraw-Hill, New York, 2013. ISBN 9780071390118.

[34] L. G. S. et al. Automatisms in absence seizures in children with idiopathic generalized epilepsy. *Archives of neurology*, 66(6), Jun 2009. doi: 10.1001/archneurol.2009.108. URL https://doi.org/10.1001/archneurol.2009.108.

[35] A. Ebner, D. S. Dinner, S. Noachtar, and H. Luders. Automatisms with preserved responsiveness. *Neurology*, 45(1):61–64, Jan 1995. doi: 10.1212/wnl.45.1.61. URL https://doi.org/10.1212/wnl.45.1.61.

[36] R. S. F. et al. Operational classification of seizure types by the international league against epilepsy: Position paper of the ILAE commission for classification and terminology. *Epilepsia*, 58(4):522–530, Mar 2017. doi: 10.1111/epi.13670. URL https://doi.org/10.1111/epi.13670.

[37] W. T. B. et al. Glossary of descriptive terminology for ictal semiology: Report of the ILAE task force on classification and terminology. *Epilepsia*, 42(9):1212–1218, Jan 2002. doi: 10.1046/j.1528-1157.2001.22001.x. URL https://doi.org/10.1046/j.1528-1157.2001.22001.x.

[38] H. Blumenfeld. What is a seizure network? long-range network consequences of focal seizures. In *Issues in Clinical Epileptology: A View from the Bench*, pages 63–70. Springer Netherlands, May 2014. doi: 10.1007/978-94-017-8914-1_5. URL https://doi.org/10.1007/978-94-017-8914-1_5.

[39] W. A. Hauser, J. F. Annegers, and L. T. Kurland. Incidence of epilepsy and unprovoked seizures in Rochester, Minnesota: 1935-1984. *Epilepsia*, 34(3):453–458, May 1993. doi: 10.1111/j.1528-1157.1993.tb02586.x. URL https://doi.org/10.1111/j.1528-1157.1993.tb02586.x.

[40] D. G. Hirtz. Generalized tonic–clonic and febrile seizures. 36(2):365–382, Apr 1989. doi: 10.1016/s0031-3955(16)36654-8. URL https://doi.org/10.1016/s0031-3955(16)36654-8.

[41] W. H. T. et al. The secondarily generalized tonic-clonic seizure. *Neurology*, 44(8):1403–1403, Aug 1994. doi: 10.1212/wnl.44.8.1403. URL https://doi.org/10.1212/wnl.44.8.1403.

[42] W. O. T. IV. *Handbook of EEG interpretation*. Demos Medical, New York, NY, 2022. ISBN 9780826147080.

[43] O. C. et al. Remission of epilepsy: results from the national general practice study of epilepsy. *The Lancet Neurology*, 346(8968):140–144, Jul 1995. doi: 10.1016/s0140-6736(95)91208-8. URL https://doi.org/10.1016/s0140-6736(95)91208-8.

[44] D. Schmidt and S. C. Schachter. Drug treatment of epilepsy in adults. 348(feb28 2):g254–g254, Feb 2014. doi: 10.1136/bmj.g254. URL https://doi.org/10.1136/bmj.g254.

[45] J. A. French. Refractory epilepsy: Clinical overview. *Epilepsia*, 48(s1):3–7, Mar 2007. doi: 10.1111/j.1528-1167.2007.00992.x. URL https://doi.org/10.1111/j.1528-1167.2007.00992.x.

[46] L. Forsgren, E. Beghi, A. Oun, and M. Sillanpaa. The epidemiology of epilepsy in europe - a systematic review. *European Journal of Neurology*, 12(4):245–253, Apr 2005. doi: 10.1111/j.1468-1331.2004.00992.x. URL https://doi.org/10.1111/j.1468-1331.2004.00992.x.

[47] J. F. T.-Z. et al. Sudden unexpected death in epilepsy: Evidence-based analysis of incidence and risk factors. *Epilepsy Research*, 65(1-2):101–115, Jun 2005. doi: 10.1016/j.eplepsyres.2005.05.004. URL https://doi.org/10.1016/j.eplepsyres.2005.05.004.

[48] R. M. et al. Mortality in adults with newly diagnosed and chronic epilepsy: a retrospective comparative study. *The Lancet Neurology*, 5(6):481–487, Jun 2006. doi: 10.1016/s1474-4422(06)70448-3. URL https://doi.org/10.1016/s1474-4422(06)70448-3.

[49] A. Gaitatzis, K. Carroll, A. Majeed, and J. W. Sander. The epidemiology of the comorbidity of epilepsy in the general population. *Epilepsia*, 45(12):1613–1622, Dec 2004. doi: 10.1111/j.0013-9580.2004.17504.x. URL https://doi.org/10.1111/j.0013-9580.2004.17504.x.

[50] M. Manford. Recent advances in epilepsy. *Journal of Neurology*, 264(8):1811–1824, Jan 2017. doi: 10.1007/s00415-017-8394-2. URL https://doi.org/10.1007/s00415-017-8394-2.

[51] M. J. Eadie. Shortcomings in the current treatment of epilepsy. *Expert Review of Neurotherapeutics*, 12(12):1419–1427, Dec 2012. doi: 10.1586/ern.12.129. URL https://doi.org/10.1586/ern.12.129.

[52] A. Jacoby. Stigma, epilepsy, and quality of life. *Epilepsy And Behaviour*, 3(6):10–20, Dec 2002. doi: 10.1016/s1525-5050(02)00545-0. URL https://doi.org/10.1016/s1525-5050(02)00545-0.

[53] C. Dravet. The core dravet syndrome phenotype. *Epilepsia*, 52:3–9, Apr 2011. doi: 10.1111/j.1528-1167.2011.02994.x. URL https://doi.org/10.1111/j.1528-1167.2011.02994.x.

[54] B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1-3):185–203, Aug 1981. doi: 10.1016/0004-3702(81)90024-2. URL https://doi.org/10.1016/0004-3702(81)90024-2.

[55] J. Gibson. *Optical flow and trajectory estimation methods*. Springer, Cham, Switzerland, 2016. ISBN 978-3-319-44940-1.

[56] D. Marr. *Vision : a computational investigation into the human representation and processing of visual information*. MIT Press, Cambridge, Mass, 2010. ISBN 978-0-262-51462-0.

[57] P. Sand and S. Teller. Particle video: Long-range motion estimation using point trajectories. *Int J Comput Vis*, 80(1):72–91, May 2008. doi: 10.1007/s11263-008-0136-6. URL https://doi.org/10.1007/s11263-008-0136-6.

[58] G. Farnebäck. Two-frame motion estimation based on polynomial expansion. In *Image Analysis*, pages 363–370. Springer Berlin Heidelberg, 2003. doi: 10.1007/3-540-45103-x_50. URL https://doi.org/10.1007/3-540-45103-x_50.

[59] A. Hyvarinen. *Independent component analysis*. J. Wiley, New York, 2001. ISBN 9780471221319.

[60] P. Comon. *Handbook of blind source separation: independent component analysis and applications*. Elsevier, Amsterdam Boston, 2010. ISBN 978-0123747266.

[61] A. Hyvarinen. Fast and robust fixed-point algorithms for independent component analysis. 10(3): 626–634, May 1999. doi: 10.1109/72.761722. URL https://doi.org/10.1109/72.761722.

[62] B. A. D. et al. Recognizing faces with PCA and ICA. 91(1-2):115–137, Jul 2003. doi: 10.1016/s1077-3142(03)00077-8. URL https://doi.org/10.1016/s1077-3142(03)00077-8.

[63] I. T. J. et al. Principal component analysis: a review and recent developments. 374(2065): 20150202, Apr 2016. doi: 10.1098/rsta.2015.0202. URL https://doi.org/10.1098/rsta.2015.0202.

[64] T. Mitchell. *Machine Learning*. McGraw-Hill, New York, 1997. ISBN 0070428077.

[65] S. Shwartz. *Understanding machine learning : from theory to algorithms*. Cambridge University Press, New York, NY, USA, 2014. ISBN 9781107057135.

[66] T. Hastie. *The elements of statistical learning : data mining, inference, and prediction*. Springer, New York, 2009. ISBN 9780387848570.

[67] M. Mohri. *Foundations of machine learning*. The MIT Press, Cambridge, Massachusetts, 2018. ISBN 0262039400.

[68] C. Rasmussen. *Gaussian processes for machine learning*. MIT Press, Cambridge, Mass, 2006. ISBN 026218253X.

[69] A. V. de Vel, K. Smets, K. Wouters, and B. Ceulemans. Automated non-EEG based seizure detection: Do users have a say? *Epilepsy Behaviour*, 62:121–128, Sep 2016. doi: 10.1016/j.yebeh.2016.06.029. URL https://doi.org/10.1016/j.yebeh.2016.06.029.

[70] P. Meritam, P. Ryvlin, and S. Beniczky. User-based evaluation of applicability and usability of a wearable accelerometer device for detecting bilateral tonic-clonic seizures: A field study. *Epilepsia*, 59:48–52, Jun 2018. doi: 10.1111/epi.14051. URL https://doi.org/10.1111/epi.14051.

[71] M. E. Thompson, J. Langer, and M. Kinfe. Seizure detection watch improves quality of life for adolescents and their families. *Epilepsy Behaviour*, 98:188–194, Sep 2019. doi: 10.1016/j.yebeh.2019.07.028. URL https://doi.org/10.1016/j.yebeh.2019.07.028.

[72] P. B. et al. A longitudinal, randomized, and prospective study of nocturnal monitoring in children and adolescents with epilepsy: Effects on quality of life and sleep. *Epilepsy Behaviour*, 61:192–198, Aug 2016. doi: 10.1016/j.yebeh.2016.05.035. URL https://doi.org/10.1016/j.yebeh.2016.05.035.

[73] A. V. de Vel and K. C. et al. Non-EEG seizure-detection systems and potential SUDEP prevention: State of the art. *Seizure*, 22(5):345–355, Jun 2013. doi: 10.1016/j.seizure.2013.02.012. URL https://doi.org/10.1016/j.seizure.2013.02.012.

[74] S. R. et al. Seizure detection, seizure prediction, and closed-loop warning systems in epilepsy. *Epilepsy & Behavior*, 37:291–307, Aug 2014. doi: 10.1016/j.yebeh.2014.06.023. URL https://doi.org/10.1016/j.yebeh.2014.06.023.

[75] S. Xie and S. Krishnan. Wavelet-based sparse functional linear model with applications to EEGs seizure detection and epilepsy diagnosis. *Med Biol Eng Comput*, 51(1-2):49–60, Oct 2012. doi: 10.1007/s11517-012-0967-8. URL https://doi.org/10.1007/s11517-012-0967-8.

[76] U. R. Acharya, S. V. Sree, and J. S. Suri. Automatic detection of epileptic EEG signals using higher order cumulant features. *International Journal of Neural Systems*, 21(05):403–414, Oct 2011. doi: 10.1142/s0129065711002912. URL https://doi.org/10.1142/s0129065711002912.

[77] M. D. et al. A multi-feature and multi-channel univariate selection process for seizure prediction. *Clinical Neurophysiology*, 116(3):506–516, Mar 2005. doi: 10.1016/j.clinph.2004.11.014. URL https://doi.org/10.1016/j.clinph.2004.11.014.

[78] Y. Liu, W. Zhou, Q. Yuan, and S. Chen. Automatic seizure detection using wavelet transform and SVM in long-term intracranial EEG. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(6):749–755, Nov 2012. doi: 10.1109/tnsre.2012.2206054. URL https://doi.org/10.1109/tnsre.2012.2206054.

[79] D. Loeckx, T. Buckinx, and L. Lagae. Validation of automatic absence seizures detection in single-lead frontal eeg. In *Epilepsia Special Issue: 33rd International Epilepsy Congress Bangkok, Thailand 22 – 26 June 2019*, volume 60, pages 5–248. Wiley, 2019.

[80] P. F. V. et al. 230 days of ultra long-term subcutaneous EEG: seizure cycle analysis and comparison to patient diary. *Ann Clin Transl Neurol*, 8(1):288–293, Dec 2020. doi: 10.1002/acn3.51261. URL https://doi.org/10.1002/acn3.51261.

[81] K. Cuppens, L. Lagae, B. Ceulemans, S. V. Huffel, and B. Vanrumste. Detection of nocturnal frontal lobe seizures in pediatric patients by means of accelerometers: A first study. volume 2009, page 6608–6611. IEEE, Sep 2009. doi: 10.1109/iembs.2009.5332557. URL https://doi.org/10.1109/iembs.2009.5332557.

[82] S. K. et al. Automated detection of convulsive seizures using a wearable accelerometer device. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 66(2):421–432, Feb 2019. doi: 10.1109/tbme.2018.2845865. URL https://doi.org/10.1109/tbme.2018.2845865.

[83] A. L. Patterson and B. M. et al. SmartWatch by SmartMonitor: Assessment of seizure detection efficacy for various seizure types in children, a large prospective single-center study. *Pediatric Neurology*, 53(4):309–311, Oct 2015. doi: 10.1016/j.pediatrneurol.2015.07.002. URL https://doi.org/10.1016/j.pediatrneurol.2015.07.002.

[84] M. Velez, R. S. Fisher, V. Bartlett, and S. Le. Tracking generalized tonic-clonic seizures with a wrist accelerometer linked to an online database. *Seizure*, 39:13–18, Jul 2016. doi: 10.1016/j.seizure.2016.04.009. URL https://doi.org/10.1016/j.seizure.2016.04.009.

[85] S. B. et al. Detection of generalized tonic-clonic seizures by a wireless wrist accelerometer: A prospective, multicenter study. *Epilepsia*, 54(4):e58–e61, Feb 2013. doi: 10.1111/epi.12120. URL https://doi.org/10.1111/epi.12120.

[86] E. Criswell. *Cram's introduction to surface electromyography*. Jones and Bartlett, Sudbury, MA, 2011. ISBN 9780763732745.

[87] I. C. et al. Automated algorithm for generalized tonic–clonic epileptic seizure onset detection based on sEMG zero-crossing rate. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 59(2):579–585, Feb 2012. doi: 10.1109/tbme.2011.2178094. URL https://doi.org/10.1109/tbme.2011.2178094.

[88] C. Á. S. et al. Electromyography-based seizure detector: Preliminary results comparing a generalized tonic–clonic seizure detection algorithm to video-EEG recordings. *Epilepsia*, 56(9):1432–1437, July 2015. doi: 10.1111/epi.13083. URL https://doi.org/10.1111/epi.13083.

[89] J. J. H. et al. Detection of generalized tonic–clonic seizures using surface electromyographic monitoring. *Epilepsia*, 58(11):1861–1869, Oct 2017. doi: 10.1111/epi.13897. URL https://doi.org/10.1111/epi.13897.

[90] A. P. N. et al. Assessment of a quasi-piezoelectric mattress monitor as a detection system for generalized convulsions. *Epilepsy & Behaviour*, 28(2):172–176, Aug 2013. doi: 10.1016/j.yebeh.2013.04.017. URL https://doi.org/10.1016/j.yebeh.2013.04.017.

[91] K. V. P. et al. Prospective study of the Emfit movement monitor. *Journal of Child Neurology*, 28 (11):1434–1436, Jan. 2013. doi: 10.1177/0883073812471858. URL https://doi.org/10.1177/0883073812471858.

[92] J. A. et al. Multimodal nocturnal seizure detection in a residential care setting. *Neurology*, 91 (21):e2010–e2019, Oct 2018. doi: 10.1212/wnl.0000000000006545. URL https://doi.org/10.1212/wnl.0000000000006545.

[93] C. C. et al. Detecting nocturnal convulsions: Efficacy of the MP5 monitor. *Seizure*, 18(3):225–227, Apr 2009. doi: 10.1016/j.seizure.2008.08.007. URL https://doi.org/10.1016/j.seizure.2008.08.007.

[94] S. F. et al. Prospective study of 2 bed alarms for detection of nocturnal seizures. *Journal of Child Neurology*, 28(11):1430–1433, Oct 2012. doi: 10.1177/0883073812462064. URL https://doi.org/10.1177/0883073812462064.

[95] W. J. van Elmpt et al. A model of heart rate changes to detect seizures in severe epilepsy. *Seizure*, 15(6):366–375, Sep 2006. doi: 10.1016/j.seizure.2006.03.005. URL https://doi.org/10.1016/j.seizure.2006.03.005.

[96] M. Z. et al. Heart rate changes and ECG abnormalities during epileptic seizures: Prevalence and definition of an objective clinical sign. *Epilepsia*, 43(8):847–854, Aug 2002. doi: 10.1046/j.1528-1157.2002.37801.x. URL https://doi.org/10.1046/j.1528-1157.2002.37801.x.

[97] S. B. et al. Pre-ictal heart rate variability assessment of epileptic seizures by means of linear and non-linear analyses. *Anadolu Kardiyol Derg.*, Oct 2013. doi: 10.5152/akd.2013.237. URL https://doi.org/10.5152/akd.2013.237.

[98] D. H. K. et al. Forecasting epilepsy from the heart rate signal. *Med. Biol. Eng. Comput.*, 43(2):230–239, Apr 2005. doi: 10.1007/bf02345960. URL https://doi.org/10.1007/bf02345960.

[99] I. Osorio. Automated seizure detection using ekg. *International Journal of Neural Systems*, 24(02):1450001, Jan 2014. doi: 10.1142/s0129065714500014. URL https://doi.org/10.1142/s0129065714500014.

[100] J. Allen. Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3):R1–R39, Feb 2007. doi: 10.1088/0967-3334/28/3/r01. URL https://doi.org/10.1088/0967-3334/28/3/r01.

[101] S. G. et al. Assessing autonomic function from electrodermal activity and heart rate variability during cold-pressor test and emotional challenge. *Scientific Reports*, 10(1), Mar 2020. doi: 10.1038/s41598-020-62225-2. URL https://doi.org/10.1038/s41598-020-62225-2.

[102] M.-Z. P. et al. Convulsive seizure detection using a wrist-worn electrodermal activity and accelerometry biosensor. *Epilepsia*, 53(5):e93–e97, Mar 2012. doi: 10.1111/j.1528-1167.2012.03444.x. URL https://doi.org/10.1111/j.1528-1167.2012.03444.x.

[103] G. R. et al. Multimodal wrist-worn devices for seizure detection and advancing research: Focus on the empatica wristbands. *Epilepsy Research*, 153:79–82, July 2019. doi: 10.1016/j.eplepsyres.2019.02.007. URL https://doi.org/10.1016/j.eplepsyres.2019.02.007.

[104] J. R. et al. Quantitative movement analysis differentiates focal seizures characterized by automatisms. *Epilepsy & Behaviour*, 20(4):642–647, Apr 2011. doi: 10.1016/j.yebeh.2011.01.020. URL https://doi.org/10.1016/j.yebeh.2011.01.020.

[105] H. L. et al. Quantifying limb movements in epileptic seizures through color-based video analysis. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 60(2):461–469, Feb 2013. doi: 10.1109/tbme.2012.2228649. URL https://doi.org/10.1109/tbme.2012.2228649.

[106] M. P. et al. Vision-based motion detection, analysis and recognition of epileptic seizures—a systematic review. *Computer Methods And Programs In Biomedicine*, 108(3):1133–1148, Dec 2012. doi: 10.1016/j.cmpb.2012.08.005. URL https://doi.org/10.1016/j.cmpb.2012.08.005.

[107] S. K. et al. Automatic segmentation of episodes containing epileptic clonic seizures in video sequences. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 59(12):3379–

3385, Dec 2012. doi: 10.1109/tbme.2012.2215609. URL https://doi.org/10.1109/tbme.2012.2215609.

[108] N. B. K. et al. Automated detection of videotaped neonatal seizures based on motion segmentation methods. *Clinical Neurophysiology*, 117(7):1585–1594, Jul 2006. doi: 10.1016/j.clinph.2005.12.030. URL https://doi.org/10.1016/j.clinph.2005.12.030.

[109] M. P. aet al. Vision-based absence seizure detection. IEEE, Aug 2012. doi: 10.1109/embc.2012.6345872. URL https://doi.org/10.1109/embc.2012.6345872.

[110] F. A. et al. Convolutional neural networks for real-time epileptic seizure detection. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging Visualization*, 6(3):264–269, July 2016. doi: 10.1080/21681163.2016.1141062. URL https://doi.org/10.1080/21681163.2016.1141062.

[111] Y. Y. et al. Video-based detection of generalized tonic-clonic seizures using deep learning. *IEEE Journal of Biomedical And Health Informatics*, 25(8):2997–3008, Aug 2021. doi: 10.1109/jbhi.2021.3049649. URL https://doi.org/10.1109/jbhi.2021.3049649.

[112] E. E. G. et al. Automated video-based detection of nocturnal convulsive seizures in a residential care setting. *Epilepsia*, 59:53–60, Apr 2018. doi: 10.1111/epi.14050. URL https://doi.org/10.1111/epi.14050.

[113] A. van Westrhenen et al. Erratum - automated video-based detection of nocturnal motor seizures in children. *Epilepsia*, Sep 2021. doi: 10.1111/epi.17056. URL https://doi.org/10.1111/epi.17056.

[114] H. A. et al. Epileptic seizure detection based on video and EEG recordings. IEEE, Oct 2017. doi: 10.1109/biocas.2017.8325156. URL https://doi.org/10.1109/biocas.2017.8325156.

[115] R. Marks. *Introduction to Shannon sampling and interpolation theory*. Springer-Verlag, New York, 1991. ISBN 9780387973913.

[116] D. Barrett. *SSH, the secure shell : the definitive guide*. O'Reilly, Sebastopol, CA, 2005. ISBN 9780596008956.

[117] H. Al-Ghaib and R. Adhami. On the digital image additive white gaussian noise estimation. IEEE, Aug. 2014. doi: 10.1109/iaict.2014.6922089. URL https://doi.org/10.1109/iaict.2014.6922089.

[118] P. H. Westfall. Kurtosis as peakedness, 1905–2014.r.i.p. 68(3):191–195, Jul 2014. doi: 10.1080/00031305.2014.917055. URL https://doi.org/10.1080/00031305.2014.917055.

[119] J. M. Eargle. Peak and RMS values of waveforms. In *Electroacoustical Reference Data*, pages 236–237. Springer, 1994. ISBN 978-1-4613-5839-8.

[120] R. Tibshirani. The lasso method for variable selection in the cox model. *Statist. Medicine*, pages 385–395, 1997.

Appendix A

Non-Video Seizure Detection Methods And Devices

Table A.1: The current state of the art in seizure detection devices; their names, price, the seizure types they claim to detect, the reported detection accuracy and latency, and the published paper describing these results.

Modality	Device Name	Price	Seizure Type	Detection Accuracy	Latency	Article
Scalp EEG	epihunter	34-39€ /month	Absence	Sensitivity: 99.6% FAR: 4.7/24h PPV: 90.3%	Not stated	Loeckx et al., 2019 [79]
Scalp EEG (implanted)	24/7 EEG SubQ	Not stated	All types	Sensitivity: Unknown, reported as high, FAR: 1/h	Not stated	Viana et al, 2020 [80]
ACM	SmartMonitor Inspyre	$20 to activate, + an additional $15-$50/month	GTCS	Sensitivity: 31-92.3% PPV: 12.9%	Not stated	Patterson et al., 2015 Lockman et al., 2011 Velez et al., 2016
ACM	Epi-Care	£1,399.00	GTCS	Sensitivity: 89.7%, FAR: 0.2/24h PPV: 47%	55s	Beniczky et al., 2013 [85]
sEMG	SPEAC system	Not stated	GTCS	Sensitivity: 76-95% FAR: 0.02-2.54/24h PPV: 3-95%	7.7s	Szabó et al., 2015 Halford et al., 2017
Under-Mattress (piezoelectric)	EMFIT MM	$594.00	GTCS	Sensitivity: 21-89%, FAR: 0.03-0.13/24h PPV: 43%	9s	Narechania et al., 2013 [90] Van Poppel et al., 2013 [91] Arends et al., 2018 [92]
Under-Mattress (microphone)	MedPage MP5	£246.00	GTCS	Sensitivity: 7-63%, FAR: 4.4/24h PPV: 3.3%, NPV: 99.8%	3-20s	Carlson et al., 2009 [93] Fulton et al., 2012 [94]
PPG + ACM	Nightwatch	£1,249.00	GTCS, tonic hyperkinetic, clusters	Sensitivity: 85%, FAR: 0.23/night PPV: 49%	Not stated	Arends et al., 2018 [92]
EDA + ACM + PPG + Temperature	Empatica E4	$1,690.00	GTCS	Sensitivity: 92-100%, FAR: 0.2-1/24h (Both use same algorithm, according to published literature)	Not stated	Regalia et al, 2019 [103]
EDA + ACM + Gyroscope + Temperature	Empatica Embrace2	£249.00 + £10-45/month	GTCS		Not stated	

Table A.2: The current state of the art in seizure detection methods; the seizure types they claim to detect, the general functioning of each algorithm, the reported detection accuracy and latency, and the published paper describing each method.

Modality	Seizure Type	Algorithm	Detection Accuracy	Latency	Article
Scalp EEG	Focal Seizures	Wavelet Decomposition + 1-Nearest-Neighbours classification	Accuracy: 97.9%	Not stated	Xie et al., 2012 [75]
Scalp EEG	Not stated	Higher Order Spectra (HOS) features + SVM classification	Accuracy: 98.5%	Not stated	Acharya et al., 2012 [76]
Intracranial EEG	Not stated	Probabilistic Neural Network	Sensitivity: 100%, FAR: 1.1/h for one patient, failed for the other patient	10 min before seizure onset (prediction)	D'Alessandro et al., 2005 [77]
Intracranial EEG	GTCS, Focal Seizures	Feature extraction + SVM	Sensitivity: 94.5%, Specificity: 95.3%, FAR: 0.58/h	11.1s	Liu et al., 2012 [105]
ACM	Nocturnal Frontal Lobe Seizures	Detection of movement epochs, moving avg., thresholding	Sensitivity: 91.7%, Specificity: 83.9%	Not stated	Cuppens et al., 2009 [81]
ACM	GTCS, Focal, Psychogenic Nonepileptic	Time-domain and Poincaré plot features + SVM	Sens: 87.0%, FAR: 2.25/24h. GTCS only: Sens: 95.8%, FAR: 0.74/24h	Not stated	Kusmakar et al., 2018 [82]
sEMG	GTCS	High-pass filter with 150 Hz cutoff, zero-crossing count	Sensitivity: 100%, FAR: 1/24h	13.7s	Conradsen et al., 2012 [87]
sEMG	GTCS	Hotelling's T^2 analysis of compound muscle action potentials	Sensitivity: 95%, 1 false alarm in 1399h of recording	20s	Szabó et al., 2015 [88]
ECG	Focal Seizures with impaired awareness	Fuzzy Clustering Algorithm to find R-R interval variability	Forecasting Sensitivity: 86%	10min to 30s before onset (prediction)	Kerem et al., 2005 [98]
ECG	Not stated	Finding R-R interval to analyse HR variability	For different settings: Sensitivity: 98%. FAR: 9.5/h / / Sensitivity: 86%, FAR: 1.1/h	0.8s before onset /13.8s after onset	Osorio et al., 2013 [99]

Appendix B

Images of the Prototype

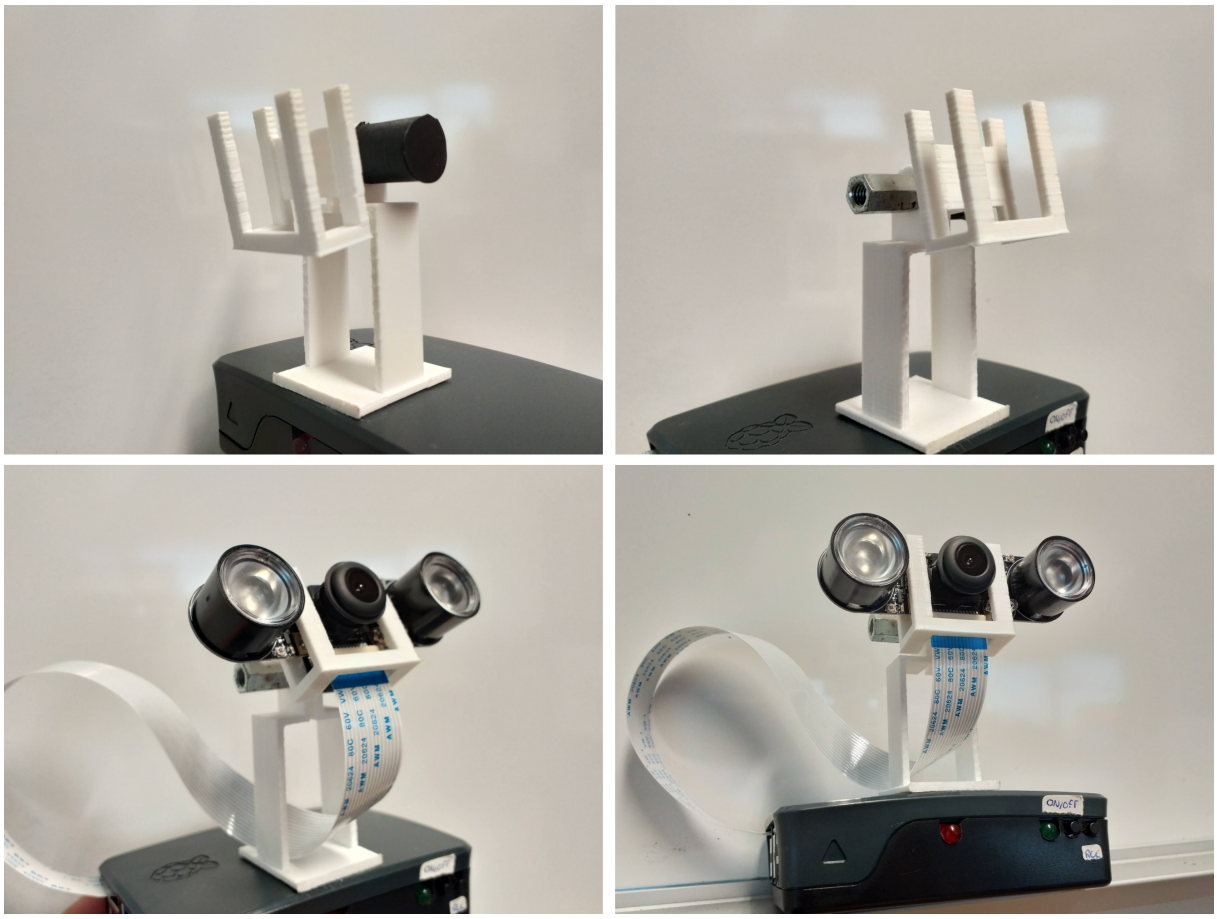

Figure B.1: The 3D model for the camera support, designed with the 3D modelling software OpenSCAD, and 3D printed using the BEETHEFIRST+ 3D printer located at IT, glued on top of the RPI case.

(a) (b)

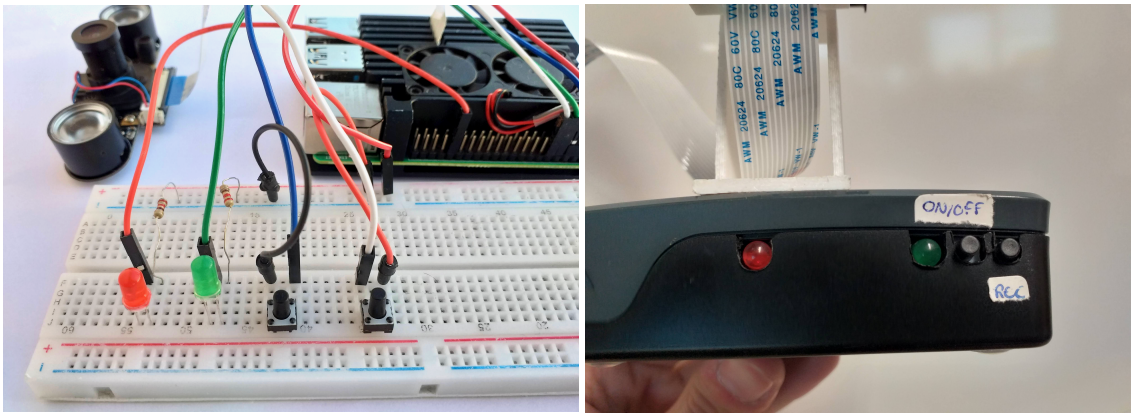

Figure B.2: (a) – The circuit shown in *Fig. 4.2* was implemented in a breadboard at first. On the left, the two LEDs are connected with their respective 220 Ω resistors and to GPIO12 and GPIO14, as well as to GND, on the RPI GPIO pins. On the right, the two buttons can be seen. The REC button is connected to 3V3 and to GPIO15, while the ON/OFF button is connected to GPIO3 and GND. (b) – The circuit was then assembled within the official case for the RPI 4 Model B. The red LED is on the left and the green LED is in the middle, while the ON/OFF and REC buttons are on the right, with temporary markings in paper to indicate which is which.

(a) (b)

Figure B.3: (a) – The support built for the 4G router, attached to the lamp pole, with the router placed on it. (b) – The 3D model for the display encasing, designed with the 3D modelling software OpenSCAD and 3D printed using the BEETHEFIRST+ 3D printer located at IT.

Appendix C

Seizure Detection Algorithm

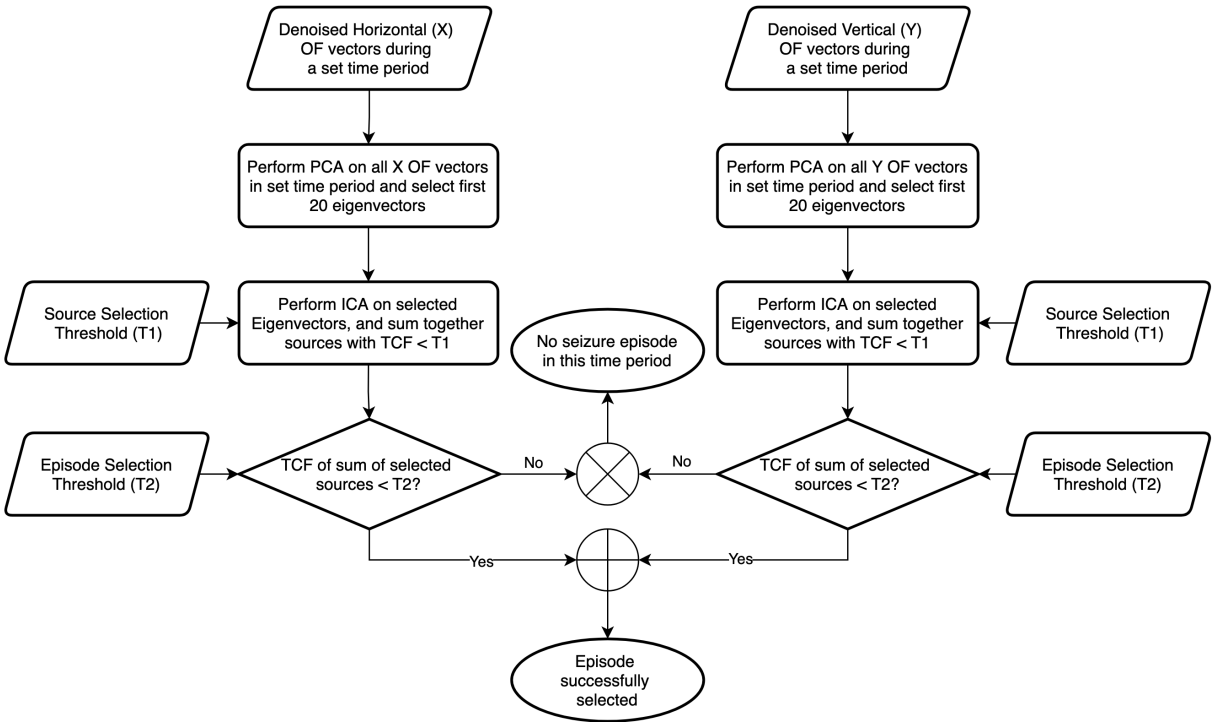

Figure C.1: Flowchart detailing the structure of the segmentation step. This step entails the successive calculation of PCA and ICA, as well as source selection, such as it was described in section 5.5.2, for both X and Y OF vectors. If TCF resulting from the sum of the selected sources is lower than the Episode Selection Threshold (T2) for either the X or the Y components, this episode is considered to be successfully selected. Otherwise, no seizure episode is detected during this period.

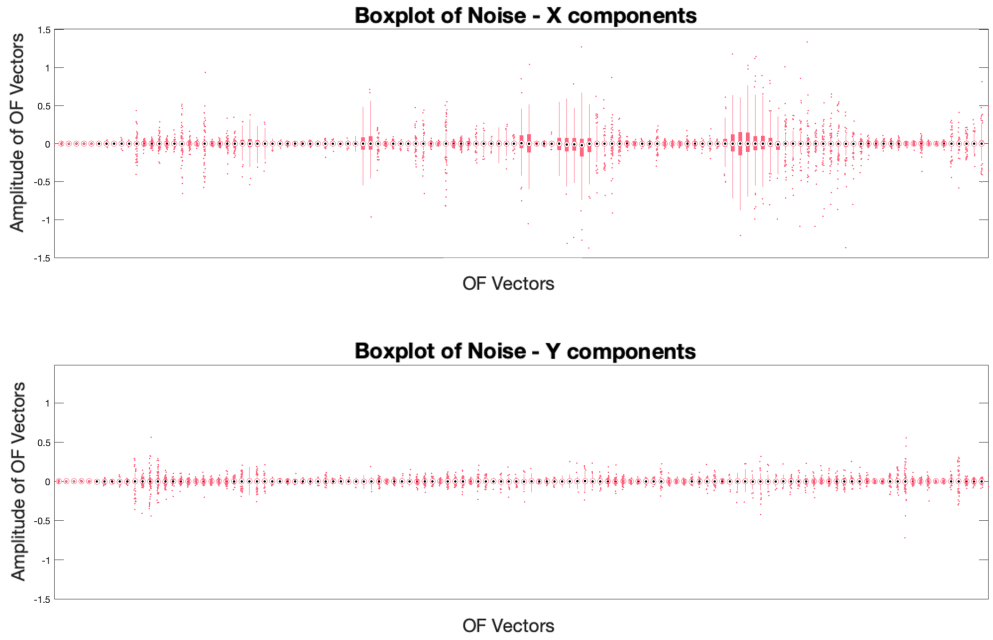

Figure C.2: A zoom-in of the boxplot in *Fig. 5.5*. It shows the distribution of the X and Y components of some OF vectors individually. Black dots denote the median, boxes denote the interquartile range and whiskers represent the variability outside of it. Whisker length is specified as 3 times the interquartile range. Outliers are plotted as dots if they fall outside this range.

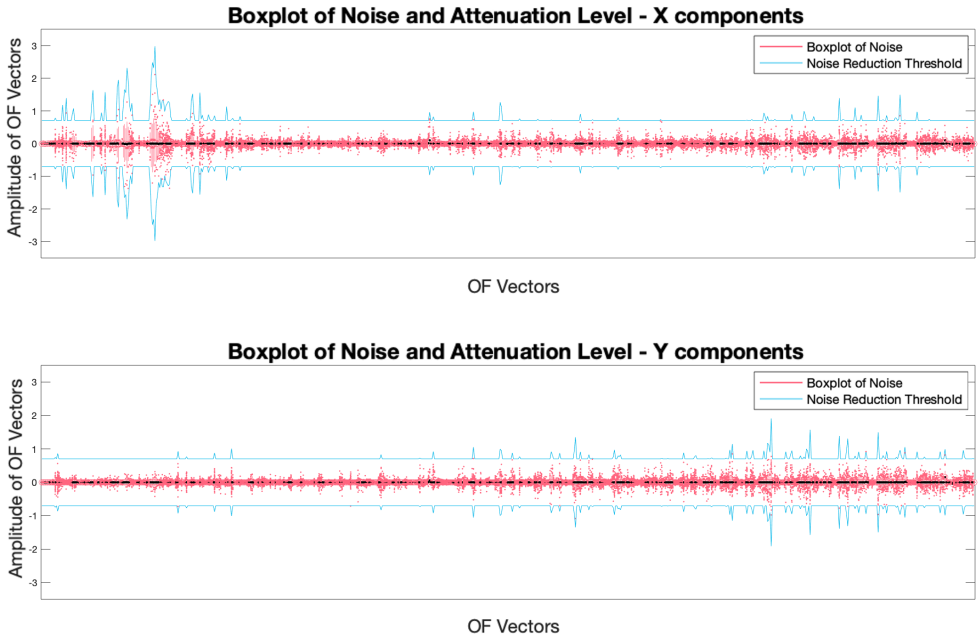

Figure C.3: Boxplot (in pink) showing the distribution of the X and Y components of all OF vectors individually, and combination of fixed and adaptive threshold values for each OF vector (in blue).

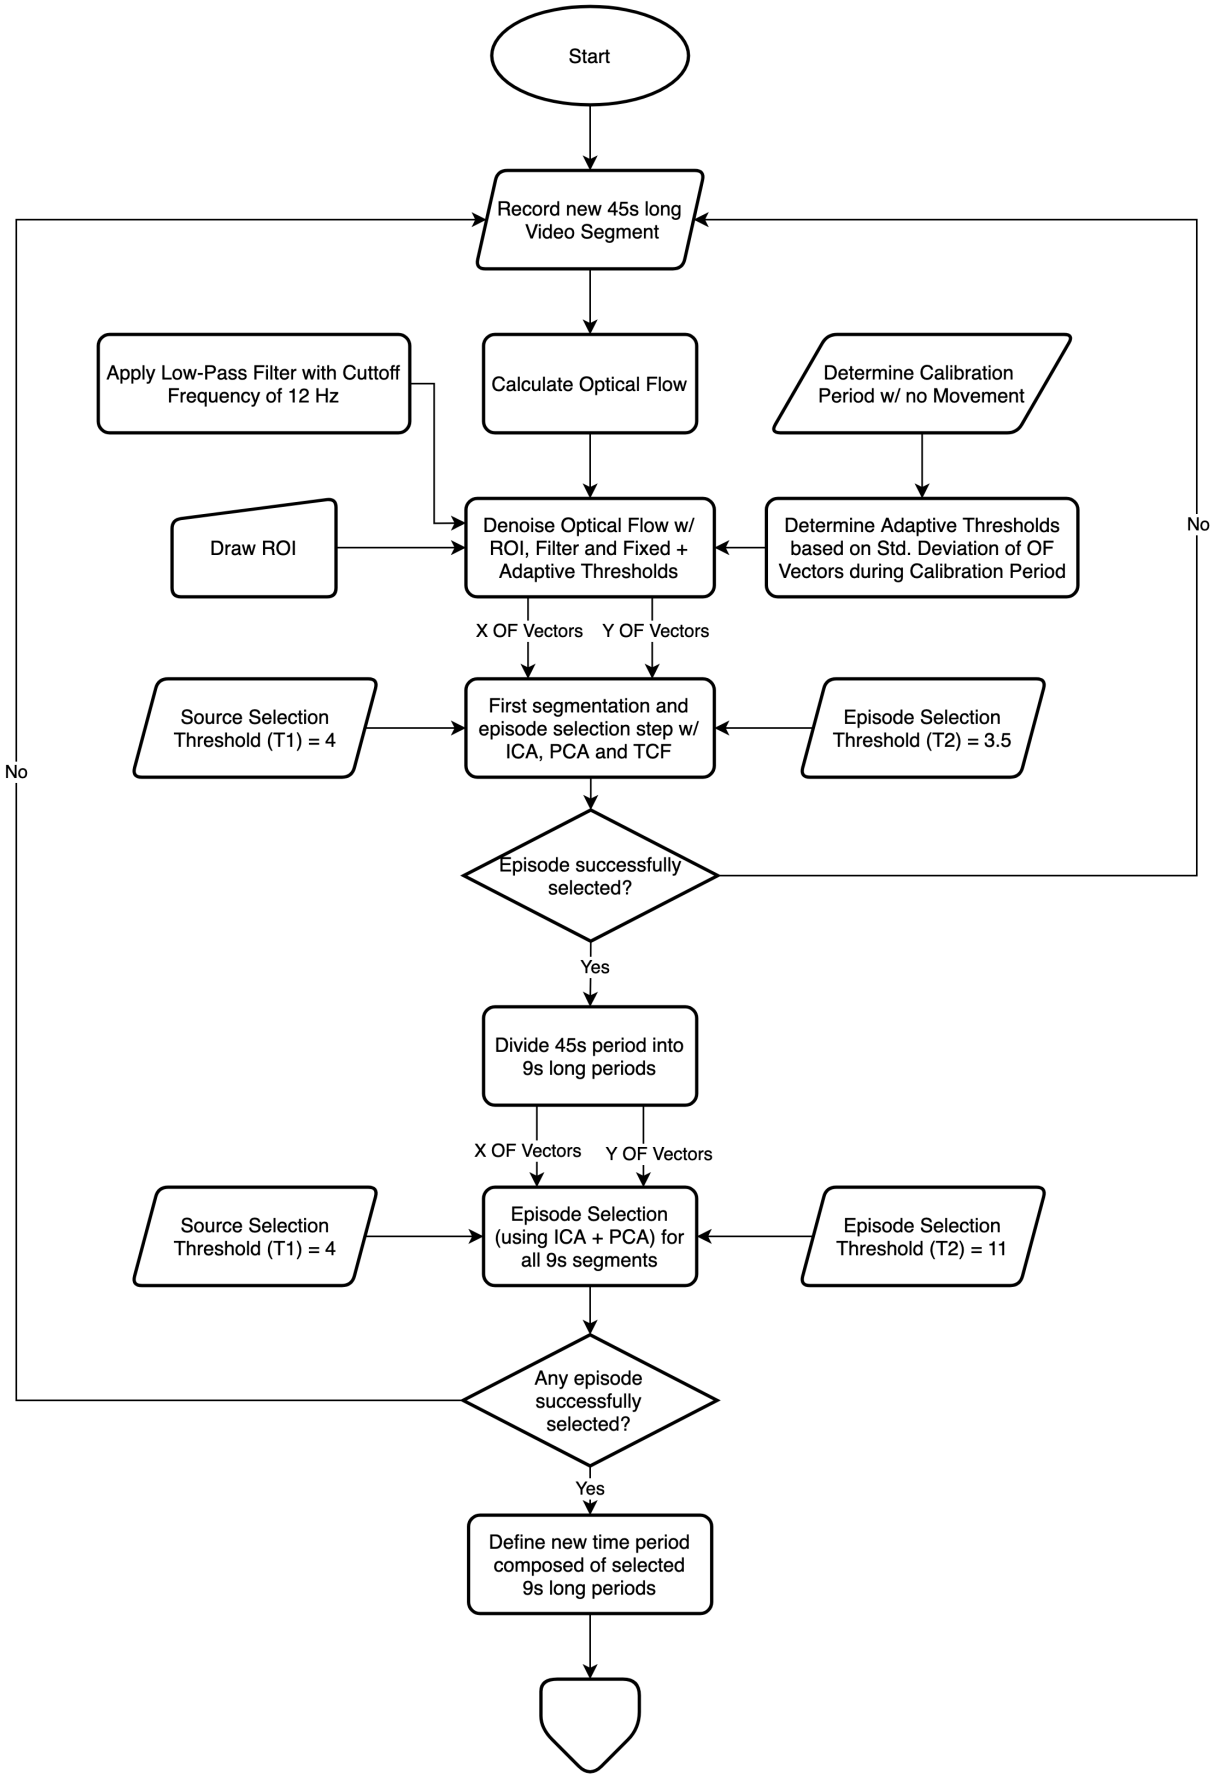

Figure C.4: First part of the flowchart detailing the structure of the seizure detection algorithm. After the OF vectors are calculated and denoised according to the process defined in Section 5.4.2, two successive segmentation steps are performed.

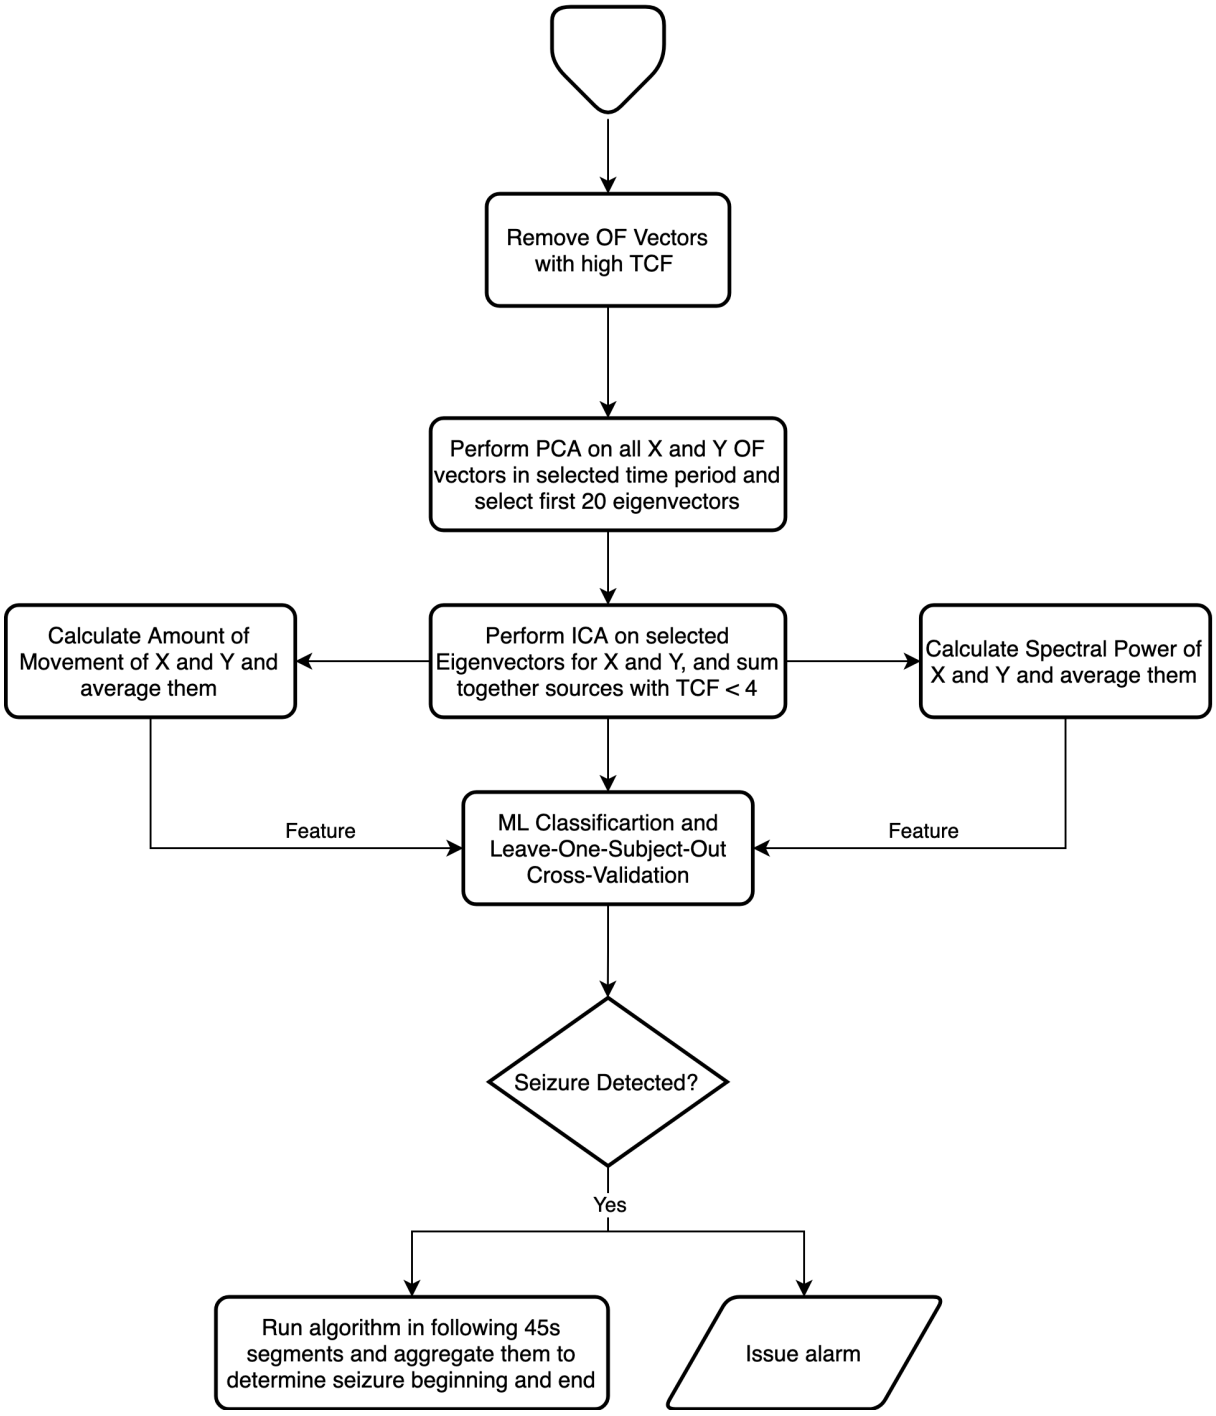

Figure C.5: Second part of the flowchart detailing the structure of the seizure detection algorithm. If, after the two segmentation steps, a possible seizure episode is segmented, ICA source selection is performed again to obtain a time-series, from which features are calculated. These features are then used to perform classification. If a seizure is detected from this classification, an alarm is issued, and then this algorithm is run again for the following 45s periods to determine the beginning and end of the seizure.

Appendix D

Results

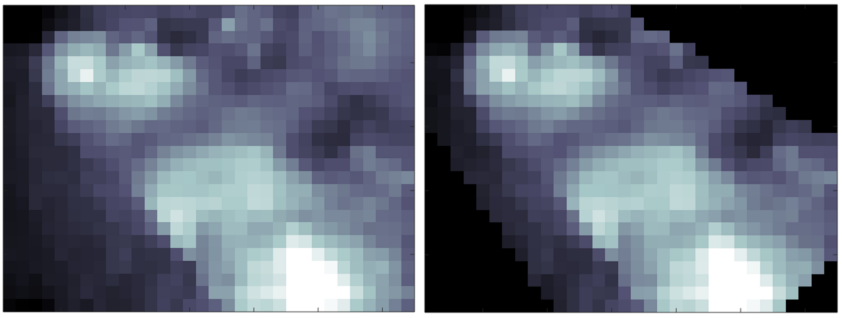

Figure D.1: Standard deviation of all X OF vectors (left) and denoised X OF vectors (right) during the seizure period, in their spatial position in the video frame. Note the higher standard deviation vectors in the center, which correspond to the OF vectors with patient movements. The largest impact comes from the ROI, which enables a better isolation of the patient's movements.

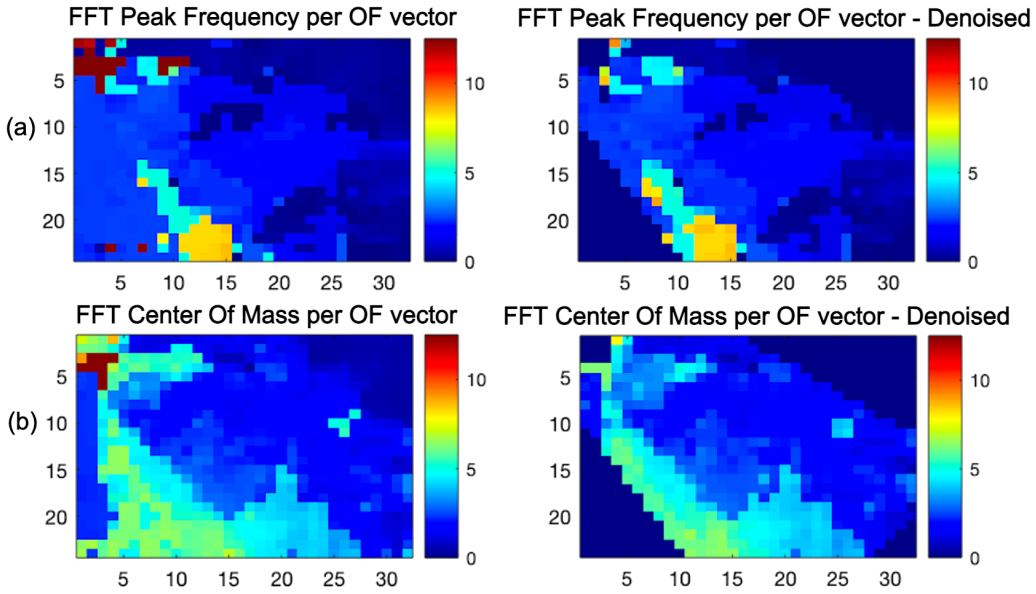

Figure D.2: (a) – Frequency of the maximum value (or peak) of the frequency spectra of the denoise X OF vectors. (b) – The center-of-mass (COM) of the frequency spectra of the denoised X OF vectors.

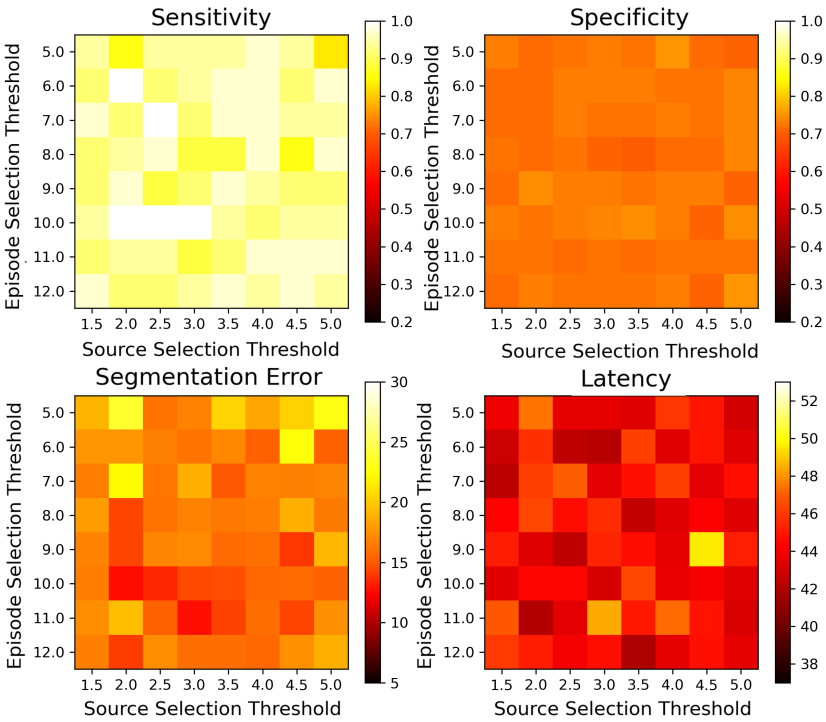

Figure D.3: The performance of the algorithm for different source and episode selection thresholds for the second (9s) segmentation step, measured by its sensitivity, specificity, latency and RMSE. The model was run 5 times for each threshold value pair, and the mean of these 5 runs was taken.