

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Automatic Detection of High Frequency Oscillations in Humans with Epilepsy

Permalink

<https://escholarship.org/uc/item/1vq7p2ms>

Author

Charupanit, Krit

Publication Date

2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Automatic Detection of High Frequency Oscillations in Humans with Epilepsy

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Biomedical Engineering

by

Krit Charupanit

Dissertation Committee:
Assistant Professor Beth A. Lopour, Chair
Professor Zoran Nenadic
Professor Jack Lin

2019

Portions of Chapter II © 2017 Brain Topography
Portions of Chapter III © 2018 IEEE
All other materials © 2019 Krit Charupanit

DEDICATION

To

my beloved parents, Ittichai and Wanna Charupanit,
respected principle investigator, Dr. Beth Lopour,
my younger brothers, Gorn and Suppaluk Charupanit,
my family members,
my closest pal, Praopim Limsakul,
my colleagues, and friends,

who have always been my source of supports and inspirations,
given the greatest words of advises and encouragement,
stayed by my sides,
believed and never given up on me,
and understood and shared my hardship,

thank you all.

Especially, to my father who had guided me throughout his life.
I dedicate this work as a memorial.
Thank you and good luck for your new journey.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	iv
LIST OF TABLES	vi
ACKNOWLEDGMENTS	vii
CURRICULUM VITAE	viii
ABSTRACT OF THE DISSERTATION	ix
CHAPTER I: Introduction	1
CHAPTER II: A Simple Statistical Method for the Automatic Detection of Ripples in Human Intracranial EEG	16
CHAPTER III: Automated Detection of High Frequency Oscillations in Human Scalp EEG	53
CHAPTER IV: Detection of Anomalous High Frequency Events in Human Intracranial EEG	65
CHAPTER V: Amplitude of High Frequency Oscillations as a Biomarker of the Seizure Onset Zone	84
REFERENCES	106

LIST OF FIGURES

		Page
Figure 2.1	Detection procedure and an example of an HFO event.	43
Figure 2.2	Iterative procedure for threshold optimization.	44
Figure 2.3	Detector performance based on ROC curves when the threshold was individually optimized to each channel.	45
Figure 2.4	ROC curves when a single α value was applied to all channels within the same subject in simple detector with iterative threshold selection and the RMS detector.	46
Figure 2.5	A comparison of ROC curves when both automatic detectors were tested in three different ways.	47
Figure 2.6	A comparison of Sensitivity, FPR, and FDR when α was individually optimized for each channel, as a function of the number of consecutive peaks for HFO detection.	48
Figure 2.7	Histogram of α values at optimum thresholds for all 296 channels.	49
Figure 3.1	Flow chart of the automatic detection algorithm	56
Figure 3.2	Examples of HFOs and events rejected during post-processing.	58
Figure 3.3	The percentage of automatically detected candidates which were visually identified as HFOs.	60
Figure 3.4	Histograms of the HFO rate per minute and duration in ripple and fast ripple.	61
Figure 4.1	Data-flow diagram for ADA.	71
Figure 4.2	Histogram of total numbers of analyzed SOZ, nSOZ channels, and segments together with number of detected events for individual patients.	74
Figure 4.3	Percentages of the detected events in the ADA-only, RMS+ADA, and RMS-only groups for each subject for SOZ channels and nSOZ channels.	75
Figure 4.4	Amplitude envelopes of detected events separated into ADA-only, RMS+ADA, and RMS-only from three representative subjects.	76

	Page
Figure 4.5	Characteristics of detected events separated into SOZ and nSOZ channels. 78
Figure 4.6	Boxplots of CV for the rate and amplitude of detected events, separated into SOZ and nSOZ. 79
Figure 5.1	Flow chart for modified ADA. 90
Figure 5.2	The cHFO rate, cHFO amplitude, and aHFA rate are higher in the SOZ, but the measurements of amplitude are more stable over time. 94
Figure 5.3	Amplitude exhibits greater differences between SOZ and nSOZ channels than cHFO rate. 96
Figure 5.4	Classification of SOZ and nSOZ channels in individual subjects, amplitude provides better performance than cHFO rate. 97
Figure 5.5	Amplitude provides superior classification of SOZ and nSOZ channels with data pooled across all subjects. 98
Figure 5.6	The amplitude metric is robust to changes in the detector sensitivity. 100
Figure 5.7	The robustness of the metrics to the changes of the maximum number of clusters in clustering and classification process. 100

LIST OF TABLES

	Page
Table 2.1	Detection statistics of our detector with iterative threshold selection at the optimum threshold. 50
Table 2.2	Detection statistics for simple detector with non-iterative threshold selection at the optimum threshold. 51
Table 2.3	Detection statistics of RMS detector at the optimum threshold. 52
Table 3.1	Confusion matrix detailing the average percentage accuracy of visual review from seven patients for R and FR. 60

ACKNOWLEDGMENTS

During my journey, I have received numerous supports and encouragements from a great number of individuals. I would like to express the deepest appreciation to my advisor, Assistant Professor Beth A. Lopour, who inspired me throughout my graduate career. It is with her unwavering confidence in me that I was able to accomplish my goal. I would also like to thank my committee members, Professor Zoran Nenadic, and Professor Jack Lin for their valuable suggestions and supports.

It is very fortunate of me to be surrounded by nice friends and colleagues. Without them, it would be very difficult for me who was from different field of study to enjoy in bioengineering. I would like to thank Dr. Daniel Shrey for his supports, and suggestions. I would like to thank Dr. Rachel J. Smith, Derek Hu, and Dr. Michael D. Nunez for being a great friend and for their role model as a PhD student. I would like to thank all current and former lab members of Lopouratory for their help, advice, and friendship.

Lastly, I would like to thank Royal Thai government for their financially support throughout my Ph.D. journey.

CURRICULUM VITAE

Krit Charupanit

Education

- 2010 B.S. in Physics
Prince of Songkla University, Songkhla, Thailand
- 2013 M.S. in Material Science and Engineering
Mahidol University, Bangkok, Thailand
- 2016 M.S. in Biomedical Engineering
University of California, Irvine, USA
- 2019 Ph.D. in Biomedical Engineering
University of California, Irvine, USA

Publications

Charupanit K, Lopour BA. A simple statistical method for the automatic detection of ripples in human intracranial EEG. *Brain topography*. 2017 Nov 1;30(6):724-38.

Charupanit K, Nunez MD, Bernardo D, Bebin M, Krueger DA, Northrup H, Sahin M, Wu JY, Lopour BA. Automated Detection of High Frequency Oscillations in Human Scalp Electroencephalogram. In 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) 2018 Jul 18 (pp. 3116-3119). IEEE.

ABSTRACT OF THE DISSERTATION

Automatic detection of high frequency oscillations in humans with epilepsy

By

Krit Charupanit

Doctor of Philosophy in Biomedical Engineering

University of California, Irvine, 2019

Assistant Professor Beth A. Lopour, Chair

Epilepsy is a chronic neurological disorder characterized by seizures. Although most patients respond favorably to medications, some patients continue having seizures and require surgery or alternative treatments. Recently, high frequency oscillations (HFOs) have been proposed as a biomarker of epileptic tissue providing a seizure onset zone (SOZ) localization and relate to surgery outcomes.

Visual HFO identification, a gold standard HFO marking, has limitations, such as, subjective, and time consuming; therefore, automatic detection algorithms have been developed. However, the automatic detections suffer from complex optimization and specific to recording. We present an algorithm with amplitude threshold as single parameter that requires optimization tuned by an iterative procedure. Algorithm is used to study HFOs in intracranial (iEEG) and scalp EEG. In the iEEG, our detector achieved 99.6% sensitivity with 1.1% false positive rate (FPR), and 37.3% false detection rate. Furthermore, the algorithm was used to detect HFOs in scalp EEG. Of the marked candidate events, 40% and 60% were visually confirmed to be ripples and fast ripples by three reviewers.

As all HFO study rely on an empirical, derived from visual observation, rather than physiological definition, we introduce the anomaly HFO detection algorithm (ADA). The algorithm integrates machine learning techniques, including anomaly detection, pattern matching, and clustering and classification to identify anomalous patterns in high frequency signals without prior assumption of the shape, amplitude, or duration. The events detected by ADA are the different population to the conventional HFOs. The amplitude of detected events is a superior candidate as a SOZ biomarker with area under the receiver operation characteristic curve (AUC), sensitivity and FPR at 0.959, 93.6% and 5.6% when comparing to the rate of conventional HFO, which was exclusively used as a biomarker, (AUC:0.912, sensitivity:86.0% and FPR:13.3%). Moreover, the amplitude is more robust to the additional events, and stable across recording segments.

We believe ADA and simple detection algorithm will be powerful tools for the assessment and localization of epileptic activity providing unbiased estimation of HFO properties. Furthermore, the amplitude of HFO can become a superior candidate as the SOZ biomarker for epilepsy patients comparing to the rate of HFO.

CHAPTER I

Introduction

1.1 Epilepsy and Seizures

Epilepsy is a chronic neurological disease that affects 65 million people worldwide (Moshé, Perucca, Ryvlin, & Tomson, 2015). Epilepsy defined as “a disorder of the brain characterized by an enduring predisposition to generate epileptic seizures and by the neurobiological, cognitive, psychological, and social consequences of this condition” (Fisher et al., 2005). The terms seizure and epilepsy are distinct. A seizure is abnormal behavior resulting from irregular electrical discharges of neurons, while epilepsy refers to a chronic condition characterized by recurrent seizures (Benbadis, 2001). People that suffer from provoked seizures, or acute symptomatic seizures, which stem from an identified cause of transient or sudden events, such as injury to the head or brain, stroke, drug abuse, high fever, or low blood sugar, are not considered to have epilepsy (Misra & Kalita, 2011). To be diagnosed with epilepsy, a patient requires the occurrence of at least one unprovoked seizure, which is defined as a transient occurrence of signs and/or symptoms due to abnormal excessive or synchronous neuronal activity in the brain with no known external cause (Fisher et al., 2005).

Epileptic seizures can be divided into three broad categories including focal onset, generalized onset, and unknown onset (Berg et al., 2010). A focal epileptic seizure, or partial seizure, is a seizure that originates within a limited area of the brain. The effects of focal seizures vary depending on the brain regions that are involved. Focal epileptic seizures can be further subcategorized by the level of awareness of the subject. The International League Against Epilepsy (ILAE) defines the aware seizure as the seizure in which the subject is aware of

themselves and their environment during the seizure, even if the subject is immobile. In a focal impaired awareness seizure, the subject will have impaired awareness during some part of the seizure (Fisher et al., 2018). In some cases, a focal seizure begins in only one brain region and spreads out to both sides of the brain, developing into a generalized seizure, which is termed a secondary generalized seizure (Fisher et al., 2018).

A generalized epileptic seizure affects larger portions of the brain than a focal epileptic seizure. It occurs throughout the whole brain and cause symptoms over the entire body. ILAE subclassifies generalized seizures into motor and non-motor (absence) seizures. An absence seizure causes rapid blinking or staring blankly into space for a short time period, up to a few seconds. A motor seizure, which is the most common generalized seizure, can also be referred to as a tonic-clonic seizure. The basic symptoms for a motor seizure are a stiffening of muscles followed by violent muscle contractions, such as jerking of the limbs and face. Furthermore, other symptoms might occur during the seizure, such as biting of the tongue, jaw locking, reduction of breathing, and cyanosis of the lips (Benbadis, 2001).

When seizures cannot be classified as generalized or focal seizures, ILAE recommends keeping the seizure type unclassified until obtaining further information. This type of seizure is defined as an unknown onset seizure. However, in some cases, the seizure may be impossible to identify due to its unusual presentation, and it will be separately categorized as an unclassified seizure (Fisher et al., 2018).

1.1.1 Diagnosis and treatment

In order to make a precise diagnosis of epilepsy, the seizure type, semiology, electroencephalogram (EEG) findings, age, cause, family history, imaging studies, and additional

information are necessary and crucial pieces of information (Benbadis, 2001). To diagnose epilepsy, clinicians first review the symptoms and the medical history of the patient to identify the type of seizure or epilepsy syndrome that best explains the events. After that, seizure patterns and brain imaging from several different modalities, such as EEG, computerized tomography (CT), or magnetic resonance imaging (MRI), are used to help determine the cause.

The treatment of epilepsy varies depending on the patient. Common treatments used to prevent seizures are medication, a special diet, implantation of a brain stimulator, and brain resection surgery. Typically, the first treatments are noninvasive procedures, such as medication and diet. The first line of epilepsy treatment is anti-convulsant medications. The majority of medications act by blocking voltage-dependent sodium or calcium channels, and some medications operate by increasing GABA activity (Perruca, 2005). The type and dose of drug depends on the type of seizure, age, gender and other medical conditions of the patient such as pregnancy. While the majority of patients respond favorably to medication and become seizure-free, some will not respond to the drugs. The chance of the patient responding to a second medication, if the first medication did not work, is significantly reduced. In one third of cases, the patient will fail multiple medications and will be diagnosed with refractory epilepsy (Schuele & Lüders, 2008). These patients may require invasive treatments for their epilepsy, such as nerve stimulation or neurosurgery. Brain stimulation has been proposed as an alternative to drug therapy as it also has the advantages of reversibility and adjustability. The treatment by brain stimulation has been studied for more than 20 years (Theodore & Fisher, 2004). Seizure generation can be inhibited indirectly by stimulation of different brain structures (Schuele & Lüders, 2008). Stimulation targets used to treat epilepsy have included several region of the brain, such as, the cerebellum, the caudate nucleus, the thalamus, and the epileptic focus itself

(Theodore & Fisher, 2004). Although stimulation for uncontrolled epilepsy with electrodes implanted in the brain itself has been done for several years, the successful results have been confirmed by the very few controlled studies so far (Theodore & Fisher, 2004). Various factors related to the stimulation have not been confirmed, such as, the best structures to stimulate, the most effective stimuli, the high-low frequency stimulation having contrast effects in different seizure models and structures (Theodore & Fisher, 2004). On the contrary, while the surgery treatment is highly effective, neurosurgery, which can be either resective or disconnective surgery, is often considered to be a last resort after noninvasive treatments have failed because it permanently alters the structure or connectivity of the brain.

1.1.2 Neurosurgery for epilepsy

The epileptogenic zone reflects “the site of the beginning of the epileptic seizure and of their primary organization” (Jehi, 2018; Talairach & Bancaud, 1966). On the contrary, SOZ is the region of the brain in which the seizure activity first begins (Rosenow & Lüders, 2001). The SOZ is generally considered to be the target for resective surgery. The target regions also often include tumors, damaged neural tissue, or other abnormal brain regions. To locate the SOZ, various brain imaging techniques are employed, such as scalp EEG, iEEG, functional MRI (fMRI), and computerized tomography (CT) scans. In some cases, when the noninvasive imaging techniques are insufficient, the patient must be monitored using invasive electrical brain signal recordings, such as iEEG, to define the location and boundaries of the epileptogenic zone. Therefore, the presurgical procedure is crucial in order to pinpoint the targeted brain region for surgery. Consequently, the presurgical procedure sometimes requires several consecutive days of monitoring in the hospital, in order to capture data from multiple seizures.

Neurosurgical procedures for epilepsy typically involve resection (e.g. anteromedial temporal lobectomy, focal neocortical resection, lesional resection, and hemispherectomy) or disconnection procedures (e.g. corpus callosotomy and multiple subpial transections)(Spencer & Huh, 2008). The selected type of surgery depends largely on the location of the SOZ, the clinical condition of the patient, and additional complicating factors, including the cause of the epilepsy, the detrimental effects of the seizures and antiepileptic drugs on the developing brain, and the capacity for functional plasticity in younger patients (Spencer & Huh, 2008). Resection and disconnection surgeries are common in treating the epilepsy with the goal of completely removing or disconnecting the epileptogenic area without causing a permanent neurological deficit (Rosenow & Lüders, 2001). Sometimes this is done using laser interstitial thermal therapy, in which a laser is used to abolish the portion of brain tissue associated with SOZ.

The reported success rates vary across study centers. For example, Rosenow and Luders (2001) reported 30–85% of epilepsy patients who underwent surgery remain seizure-free (Rosenow & Lüders, 2001) while Engel (1993) found a ~ 60% successful rate (Engel, 1993). The success of the surgery depends on various factors including the epilepsy syndrome and the ability to localize and resect the epileptogenic zone.

1.2 Biomarkers for the seizure onset zone

As the localization of the SOZ is very important for epilepsy surgery, several attempts have been made to find a biomarker to identify this region in individual patients. The availability of a robust biomarker will lead to a better surgical outcome and help reduce the amount of time needed for intracranial monitoring in presurgical procedures.

1.2.1 Interictal epileptiform discharges

Interictal epileptiform discharges (IEDs) are pathological patterns in the EEG that can be observed in the period between seizures, called the interictal period, of patients with epilepsy. Interictal patterns are diverse (Curtis & Avanzini, 2001), but the best-known types of IED are spike and sharp wave. The original definition of IED was relatively broad, as “distinct waves or complexes distinguishable from the background activity, and with a sharp configuration”; consequently, the definition was refined to “a waveform that is characterized by an amplitude of 2.5 times that of the underlying background with a duration of either less than 70 milliseconds (spikes) or longer than 70 milliseconds (sharp waves),” at a symposium of the American Clinical Neurophysiological Society (International Federation of Societies for Electroencephalography, 1983). Several studies showed a correlation between IEDs and seizure location in intractable partial epilepsies (Avoli, 2001; Curtis & Avanzini, 2001; Curtis, Jefferys, & Avoli, 2012; Gotman, 1991); therefore, IEDs have been believed to be a potential biomarker for epileptogenic tissue. IEDs can be observed from both scalp EEG and iEEG recordings (Curtis et al., 2012; Gotman, 1991; Wang et al., 2013). The study also showed that patients who had spikes in multiple areas of the brain and later had surgery were less likely to become seizure-free post-surgery than the patients with well localized spikes (Bautista, Cobbs, Spencer, & Spencer, 1999). However, because IEDs can also be found in the irritative region and can spread to adjacent healthy brain structure outside the SOZ (Curtis et al., 2012), there is still no concrete conclusion regarding the relationship between IEDs and the SOZ

1.2.2 High frequency oscillations

As IEDs are not exclusive to the SOZ, other electrographic patterns have been studied, such as high frequency oscillations. Studies have shown that the analysis of oscillatory activity in a higher frequency band than is classically used for clinical purposes gives reliable information about the location of the SOZ (Bragin, Engel, & Staba, 2010; Dumpelmann, Jacobs, & Schulze-Bonhage, 2015; Engel et al., 2009; J. Jacobs et al., 2012; Julia Jacobs et al., 2009; M Zijlmans, Jacobs, Zelmann, Dubeau, & Gotman, 2009; Maeike Zijlmans et al., 2012). High frequency oscillations (HFOs) are defined as spontaneous electrographic patterns consisting of at least four cycles of an oscillation with a high amplitude that is distinguishable from the background in the frequency range of 80–500 Hz (G Buzsáki, Horváth, Urioste, Hetke, & Wise, 1992; Frauscher et al., 2017; J Jacobs et al., 2012). HFOs can be divided into two sub-frequency bands: ripple (80–250 Hz), and fast ripple (250–500 Hz); however, these frequencies were found empirically and are not strictly defined. Another type of HFO that has been reported recently is the very high frequency oscillation (600–2,000 Hz). These were found in mesiotemporal structures of patients with temporal lobe epilepsy (Brázdil et al., 2017); however, this high frequency activity is not well defined and understood at present.

There is evidence of a relationship between HFOs and IEDs. For example, the majority of ripples and fast ripples were found to co-occur with IEDs in conventional EEG (Crépon et al., 2010; Julia Jacobs et al., 2008; Urrestarazu, Chander, Dubeau, & Gotman, 2007). Moreover, the pathological HFOs reflect hyperactive population spikes while IEDs reflect hypersynchronous postsynaptic potentials (Bragin, Wilson, & Engel, 2007; Grenier, Timofeev, & Steriade, 2001; Wang et al., 2013). HFOs were found to be more frequent in patients with a high number of IEDs (Wang et al., 2013). However, even though HFOs were found to be related to IEDs, they are more specific to the SOZ than IEDs (Fedele et al., 2016; Frauscher et al., 2017; Melani,

Zelmann, Dubeau, & Gotman, 2013) and are therefore thought to be a better metric to localize the SOZ.

HFOs have been observed in various brain regions. First, HFOs were found in the hippocampus in a study of a rodent model of epilepsy (Bragin, Engel, Wilson, Fried, & Mathern, 1999; G Buzsáki et al., 1992). Later, multiple studies also found HFOs in iEEG recordings from humans with epilepsy (Bragin et al., 2010; Bragin, Engel, Wilson, Fried, & Mathern, 1999; K. Charupanit & Lopour, 2017; Dumpelmann et al., 2015; Ferrari-Marinho et al., 2015; Gliske et al., 2018; Spring et al., 2018; M. Wu et al., 2018). More recently, HFOs have also been observed in human scalp EEG recordings (Bernardo et al., 2018; Krit Charupanit et al., 2018; J Jacobs et al., 2012; Kobayashi et al., 2010).

Early studies reported that pathological HFOs, specifically fast ripples (250–500 Hz), can be observed in epileptic rats and localized to the epileptic hippocampus that is generating the seizures (Bragin, Engel, Wilson, Fried, & Mathern, 1999; Bragin, Mody, Wilson, & Engel, 2002). The early studies of HFOs supported the hypothesis that fast ripples were distinctly pathological oscillations associated with epileptic brain (Bragin, Engel, Wilson, Fried, & Buzsáki, 1999; Bragin, Engel, Wilson, Fried, & Mathern, 1999; Matsumoto et al., 2013). However, later studies of HFOs contradicted this hypothesis, showing that both gamma and ripple frequency oscillations were also increased in human epileptogenic hippocampus (Crépon et al., 2010; Joshua Jacobs & Kahana, 2010; Greg A. Worrell et al., 2004) and neocortex (Blanco et al., 2011; Joshua Jacobs & Kahana, 2010; Schevon et al., 2009). Although most studies of high frequency activity focus on pathological HFOs that relate to epilepsy, physiological HFOs can also be found in healthy brain regions. Physiological HFOs have the same basic morphology as pathological HFOs and are typically found in the ripple frequency band. They are associated

with cognitive processing, information processing, and memory transference and consolidation (György Buzsáki & Silva, 2012; Kucewicz et al., 2014; Lachaux, Axmacher, Mormann, Halgren, & Crone, 2012; Matsumoto et al., 2013). Thus, the distinction between physiological and pathological HFO has remained blurred (Bragin et al., 2010; Engel et al., 2009; Matsumoto et al., 2013).

The rate of occurrence of HFOs has been almost exclusively used as an SOZ biomarker because in initial studies using visual detection, they were found more frequently in the SOZ compared to non-SOZ brain regions (Burnos et al., 2014; Cho et al., 2014; Dumpelmann et al., 2015; Fedele et al., 2016; Julia Jacobs et al., 2008; Liu et al., 2016; Nicolás von Ellenrieder, Andrade-Valença, Dubeau, & Gotman, 2012; M. Wu et al., 2018). In addition, the removal of brain tissue generating a higher rate of HFOs is strongly correlated with the outcome of epilepsy surgery (Haegelen et al., 2013; Julia Jacobs et al., 2010).

1.3 HFO detection

HFOs have typically been identified by either employing visual HFO marking by expert human reviewers or using automated HFO marking algorithms. In early studies, visual HFO detection was the only tool for marking HFOs and was widely applied to both scalp and intracranial EEG recordings (Bernardo et al., 2018; Frauscher et al., 2017; Pizzo et al., 2016; Spring et al., 2017; J. Y. Wu et al., 2010). Visual HFO detection is considered to be the gold standard and a benchmark for automated HFO detection (Chaibi, Lajnef, Sakka, Samet, & Kachouri, 2013; Frauscher et al., 2017; Nicolás von Ellenrieder et al., 2012; Zelmann et al., 2012). However, visual marking by reviewers is infeasible for longer studies. Therefore,

automated detectors have been increasingly developed to improve processing time, reliability, and repeatability of the results.

Automatic detection algorithms generally consist of two key steps: 1) initial detection for marking the candidate events and 2) rejection of artifacts and false positives from the group of candidate events. Most automated detectors focus on the first step without employing additional steps to reject false positives (Biro, Kachenoura, Albera, Bénar, & Wendling, 2013; Chaibi et al., 2013; K. Charupanit & Lopour, 2017; Dümpelmann, Jacobs, Kerber, & Schulze-Bonhage, 2012; Zelman et al., 2010), or they use automated post-processing steps to identify false positives (Blanco et al., 2010; Krit Charupanit et al., 2018; Fedele et al., 2016; Gliske et al., 2016; Matsumoto et al., 2013). In order to detect high frequency events that stand out from the background, the algorithms typically require the recording signal to be pre-processed and the events must fulfill various requirements, such as exceeding energy and duration thresholds. The energy of the signal is usually measured via various techniques, such as root mean square (RMS) amplitude (Blanco et al., 2010; Staba, Wilson, Bragin, Fried, & Engel, Jerome, 2002; Zelman et al., 2012), amplitude of filtered, rectified data (K. Charupanit & Lopour, 2017; Krit Charupanit et al., 2018), line length (Dümpelmann et al., 2012; Matsumoto et al., 2013), Hilbert envelope (Burnos et al., 2014; Crépon et al., 2010; Fedele et al., 2016), or power based on the time-frequency decomposition (Khadjevand, Cimbalkin, & Worrell, 2017). Minimum duration is another common condition of the detection, which is defined as a minimum number of oscillations or a minimum time above the energy threshold. The segments of the signal that meet all criteria of the algorithm will be considered as HFO candidates. Typically, the automatic HFO detection algorithms contain at least 3-5 interrelated parameters and must be optimized for individual subjects or electrodes in order to achieve the maximum detection accuracy (Krit

Charupanit et al., 2018; Maeike Zijlmans et al., 2012). Adding additional parameters to the algorithm can improve the reliability and accuracy of the detection, but this also increases the complexity of the optimization and makes implementation more difficult.

Semi-automated HFO detection combines automatic HFO detection and human supervision, with the goal of maintaining the advantages of both detection schemes. Although semi-automated procedures shorten the time required for visual detection, they are still generally infeasible for studies with long recorded signals. Some recent studies have analyzed data recorded over 24 hours or more (Bernardo et al., 2018; Joyce Y. Wu, Koh, Sankar, & Mathern, 2008). The general procedure of semi-automated HFO detection starts with applying a hyper-sensitive automatic HFO detector to capture HFO candidates. Then the detected candidates are visually reviewed by trained human reviewers to reject falsely detected events (Crépon et al., 2010; Gardner, Worrell, Marsh, Dlugos, & Litt, 2007; Staba et al., 2002).

1.4 Barriers to applying HFOs to clinical practice

HFOs have been a focus of research for two decades; however, only a few centers are attempting to employ HFOs in clinical practice due to several barriers. One of the biggest obstacles is the insufficient understanding of HFO physiology. Despite the extensive number of studies of HFOs, both visual and automated HFO detection schemes still rely on the original empirical definition of an HFO derived from visual observation (G Buzsáki et al., 1992). It requires HFO candidates to exceed energy and duration thresholds, as a means of separating events from the background. Because the shape and amplitude of the HFO waveform can vary depending on the distance between the recording electrode and the neural generator (György Buzsáki, Anastassiou, & Koch, 2012), a rigid template of shape and amplitude is likely to be

insufficient for HFO detection. Together with the lack of a well-defined physiological definition and its detailed knowledge of the mechanism underlying HFOs, HFOs used in different studies are diverse and non-standardized which limits our understanding of these events and will likely be deemed too risky to use them to guide brain surgery.

The second limitation is the HFO detection. Although visual detection is considered to be the gold standard, with the capability to reject artifacts and adapt to different baseline activities, its use is limited because it is highly time-consuming and subjective. In Zelman et al. 2012, there was only 61% agreement between two trained reviewers (Zelman et al., 2012), while Spring et al. 2017 also showed that the interrater reliability was poor and visual markings were significantly different for most pairs of the six tested reviewers (Spring et al., 2017). Since most studies employed only two reviewers to mark HFOs, they are prone to bias that arises from poor agreement between reviewers. Moreover, visual detection is highly time-consuming. Visual HFO marking for 2 minutes of iEEG recordings with eight channels can take more than an hour (K. Charupanit & Lopour, 2017; Krit Charupanit et al., 2018). To overcome this limitation, automated detectors have been developed and have gradually become more reliable and versatile. However, automatic HFO detectors often suffer from high false detection rates (Amiri, Lina, Pizzo, & Gotman, 2016; Chaibi et al., 2013; Dümpelmann et al., 2012; Gliske et al., 2018; Matsumoto et al., 2013; Zelman et al., 2012) and complex optimization processes. In order to achieve accurate detection, automated detectors require a training process for parameter optimization, and this is usually repeated for individual patients or channels. Furthermore, automated HFO detectors are usually designed and optimized for a specific research center and data type. As a result, these algorithms do not easily generalize to new datasets. The time and effort needed for validation and optimization become a huge barrier for clinical implementation.

1.5 Scope of the dissertation

In this doctoral dissertation, we develop new tools to study HFOs by proposing two new automated HFO detection algorithms with different goals. The first HFO detector, simple HFO detector, aims to perform with the minimum requirement of the parameter optimization and training. The second detector, anomalous HFO detection algorithm, intends to identify HFOs without the prior definition of conventional HFO. Then we employ both HFO detectors to study the characteristics of HFOs in iEEG and scalp EEG recording with the ultimate goal of using HFOs to accurately localize the SOZ in refractory epilepsy patients that require brain surgery.

To provide an efficient and powerful tool to identify HFOs, in Chapter II we propose a novel algorithm for automatic detection of HFOs in human iEEG data (K. Charupanit & Lopour, 2017). The algorithm design focuses on the simplicity and ease of implementation. This method requires the optimization of a single parameter, the amplitude threshold, which is directly related to the sensitivity and specificity of the detection. The algorithm employs an iterative approach that stems from the idea of distinguishing transient burst-like events from background noise that is similar to the single-molecule fluorescence experiments (Grange et al., 2008). Here, we adapt this technique to use the rectified filtered iEEG signal instead. The algorithm estimates the amplitude probability distribution of the background activity based on each oscillation in the filtered rectified data, and the distribution is used to calculate an optimal amplitude threshold. We directly compared the performance of our algorithm to the RMS detector (Staba et al., 2002) by testing them on the same dataset of visually identified HFO events (Zelmann et al., 2012). We evaluate the performance of HFO detectors by using the HFOs detected by visual detection as a benchmark. As our algorithm requires the optimization of only a single parameter, it offers

advantages over alternative methods for HFO detection that involve several inter-related parameters.

The high frequencies associated with HFOs (80-500 Hz) were not traditionally analyzed in scalp EEG; however, improvements in EEG recording technology now enable the recording of scalp EEG with very high sampling rate. Subsequently, studies have reported that HFOs can also be observed in scalp EEG recordings of epilepsy patients (Andrade-Valenca, Dubeau, Mari, Zelmann, & Gotman, 2011a; Frauscher et al., 2017; Kobayashi et al., 2010; Pizzo et al., 2016; Nicolás von Ellenrieder et al., 2012; Joyce Y. Wu et al., 2008; Zelmann, Lina, Schulze-Bonhage, Gotman, & Jacobs, 2014). In Chapter III, we adapted our simple automatic HFO detection algorithm from Chapter II (K. Charupanit & Lopour, 2017) to detect HFOs in scalp EEG recordings. The target frequency bands of detection were both ripples, (80-250 Hz) and fast ripples (250- 500 Hz). Because artifacts are prevalent in scalp EEG recordings, it is essential to implement several additional false positive rejection criteria as post-processing steps to the primary detection.

Although HFOs are a promising biomarker for the epileptogenic zone, there no physiological definition of an HFO has been established. HFO detection typically relies on the empirical definition of the HFO derived from visual observation. Therefore, HFO detection can lead to biased estimates of HFO characteristics, thereby decreasing their performance as metrics to localize the SOZ. In Chapter IV, we present a novel anomaly detection algorithm (ADA) to detect any high frequency events ($> 80\text{Hz}$) in the iEEG that stand out from the background, regardless of amplitude or morphology. ADA integrates several machine learning techniques and does not require human supervision, parameter optimization, or prior assumptions about the shape, amplitude, or duration of the events. We then measure the characteristics of detected

anomalous events, including rate, amplitude, and duration, and we compare them to conventional HFOs detected by the RMS detector, which is a standard approach that relies on an energy-based threshold.

Previous studies showed that the amplitude of pathological HFOs is higher than physiological HFOs across different frequency bands(Alkawadri et al., 2014; Matsumoto et al., 2013; Pail et al., 2017; van Klink, Frauscher, Zijlmans, & Gotman, 2016; Nicolás von Ellenrieder et al., 2012; Wang et al., 2013) suggesting that HFO amplitude could be used as a biomarker of the SOZ. However, the reported differences in amplitude between SOZ and non-SOZ electrodes were typically small and sometimes not statistically significant. We hypothesize that this may be a byproduct of amplitude-based HFO detection and the selection of parameters for the algorithm. In Chapter V, we measured the amplitude and rate of HFOs in human iEEG using two automatic detection algorithms. First, we measured the amplitude and rate of anomalous high frequency events identified using the ADA from Chapter IV. Second, we repeated the analysis using a conventional HFO detector based on RMS amplitude. For both detectors, we compared the rate and amplitude of detected events in SOZ and non-SOZ electrodes and evaluated the classification accuracy and robustness of these two metrics.

Both novel automated HFO detectors presented here, the simple HFO detector and the ADA, require only minimal optimization procedures and human involvement. Therefore, we believe these techniques can aid researchers in studying HFOs from various perspectives and can be consistently applied across data from different research centers with different recording modalities. Moreover, our detectors will be a powerful tool to assist clinicians in the diagnosis and surgical treatment of patients with refractory epilepsy.

CHAPTER II

A Simple Statistical Method for the Automatic Detection of Ripples in Human Intracranial EEG

2.1 Introduction

High frequency oscillations (HFOs) are a promising biomarker for the delineation of the seizure-onset zone in epileptic patients (Bragin et al., 2010; Maeike Zijlmans et al., 2012), and the removal of brain regions exhibiting HFOs has been correlated with a greater likelihood of seizure-free surgical outcome (Fujiwara et al., 2012; Julia Jacobs et al., 2010). Pathological HFOs are broadly defined as spontaneous electrographic events of at least three to four full oscillations that are clearly distinguishable from the background signal, with frequencies ranging from 80-500 Hz (G Buzsáki et al., 1992; J Jacobs et al., 2012). Visual identification by expert reviewers is a widely used procedure for detecting HFOs in both scalp and intracranial EEG recordings (Ferrari-Marinho et al., 2015; Julia Jacobs et al., 2014), and it is currently considered the gold standard (G A Worrell et al., 2012). This type of manual detection affords great adaptability to recordings with different baseline levels, and rejection of artifacts can be accomplished simultaneously. However, the results are subjective; Zelman et al. (2012) reported 60.98% agreement between two trained, human reviewers. The process is also highly time-consuming, as visual identification of HFOs in 2 minutes of EEG recordings with 8 channels can take about two hours. For these reasons, visual analysis is infeasible for longer studies. To overcome these disadvantages, there has been a push to develop automatic detection algorithms to improve the speed and accuracy of detection across different recordings and clinical settings, with the long-term goal of eventually replacing traditional visual identification.

HFO detection typically consists of two steps: (1) initial detection of candidate events and (2) rejection of artifacts and other false positives. A majority of published detectors focus primarily on the first step of this process and rely on human visual validation for the rejection of artifacts (Crépon et al., 2010; Gardner et al., 2007; Staba et al., 2002) or accept the results of detection without post-processing (Chaibi et al., 2013; Dümpelmann et al., 2012; Zelman et al., 2010). To detect high frequency events that stand out from the background, these algorithms filter the raw data in the frequency band of interest and then estimate the energy of the signal using various techniques, including the RMS amplitude (Staba et al., 2002), line length (Gardner et al., 2007), Hilbert transform (Crépon et al., 2010), Hilbert Huang transform (Chaibi et al., 2013), or a combination of techniques (Dümpelmann et al., 2012; Zelman et al., 2010). Then a threshold is applied, along with other parameters such as minimum duration, minimum number of oscillations, or minimum time between successive events. Increasing the complexity of the algorithm can improve accuracy, but it also makes implementation more difficult. Most algorithms require a complex optimization of three to four parameters, and this optimization step is critical for achieving accurate performance (Zelman et al., 2012).

More recently, several algorithms for automated artifact rejection have been developed (Amiri et al., 2016; Burnos et al., 2014; Cho et al., 2014; Gliske et al., 2016). These algorithms are independent from the initial detection, which is accomplished through either human visual analysis (Amiri et al., 2016) or an existing detector (Burnos et al., 2014; Cho et al., 2014; Gliske et al., 2016). These methods reduce the number of false detections, but again increase the complexity and number of parameters, making implementation more challenging.

Overall, both detection and artifact rejection algorithms tend to be optimized for the recordings of a specific research group, as they are highly sensitive to characteristics of the data,

including the source and prevalence of artifacts. This makes it difficult to directly compare results of different studies and remains a barrier to the translation of automated HFO detection to the clinical setting. We hypothesize that simplifying the detection procedure will aid in overcoming some of these challenges.

Therefore, we propose a new algorithm for automatic detection of HFOs in human intracranial EEG (iEEG) data focusing on simplicity and ease of implementation. This method requires optimization of only one parameter, the amplitude threshold, which is directly related to the sensitivity and specificity of the detection. We propose two approaches for selection of the threshold: an iterative method and a non-iterative method. The iterative algorithm stems from the idea of distinguishing transient burst-like events from background noise, as is done in single-molecule fluorescence experiments (Grange et al., 2008). In those experiments, the fluorescence intensity is analyzed, which is always a positive quantity; therefore, we adapt the rectified filtered iEEG signal to the algorithm. Based on the amplitude of each oscillation in the filtered data, an estimate of the probability distribution of the background activity is obtained via an iterative process and is used to calculate an optimal amplitude threshold. In the non-iterative method, the threshold is calculated directly from the mean and standard deviation of the filtered, rectified signal. We directly compared the performance of our algorithm to the Root Mean Square (RMS) detector (Staba et al., 2002) by testing them on the same dataset of visually identified HFO events from Zelman et al. 2012. Because our algorithm achieves accurate performance and requires the optimization of only a single parameter, it offers advantages over alternative methods for the initial detection of candidate HFO events.

2.2 Methods

2.2.1 Electrophysiological recordings

All recordings used in this analysis were previously published by Zelman et al. (2012). The iEEG recordings were collected between September 2004 and April 2008 from 45 patients with medically refractory epilepsy who underwent depth macro-electrode implantation at the Montreal Neurological Hospital. Recordings were sampled at 2 kHz and low-pass filtered at 500 Hz. The dataset contained 1-minute segments of slow wave sleep from 19 patients, each with 10-39 channels. As in the original study, channels with continuous artifacts, channels containing zero visually-detected HFOs, and channels without distinguishable baseline were excluded, leaving 373 channels for analysis.

2.2.2 Channel selection and event identification

HFOs were visually identified by two experienced reviewers at the Montreal Neurological Institute. In addition to HFOs, reviewers marked baseline segments of data with no oscillatory activity (no possible HFOs) with a maximum length of 400 ms (NegBASE). If baseline segments exceeded the maximum duration, they were split into multiple 200 ms-long events. For this study, automatic detection was performed on the ripple band (80-250 Hz) because there were very few visually marked fast ripples (250-500 Hz).

Two validation schemes introduced by Zelman et al. (2012), referred to as “strict” and “open” validation were used to evaluate the performance of the automatic detectors. The strict validation was based on HFOs (PosAND) and baseline segments (NegBASE) jointly marked by both reviewers, representing positives and negatives respectively. In the open validation scheme, events marked by either one or both reviewers (PosANY) were considered, since events marked by only one reviewer are possible HFOs, but with lower probability. In this scheme, the negative

baseline was defined as any iEEG segment with a minimum of 25 ms separation from any PosANY event. Sensitivity and FPR were evaluated based on strict validation; on the contrary, false detection rate (FDR) employed open validation.

We excluded four channels because they contained only fast ripples (250-500 Hz) with no ripples, leaving us with 369 channels for analysis. Within these channels, 20% (73 channels) were randomly selected to serve as training channels to optimize the parameters of the RMS detector. The remaining 296 channels were used to evaluate the performance of both automatic detectors, which contained 7561 PosAND HFOs, 12316 PosANY HFOs, and 43101 NegBASE segments that were visually identified.

2.2.3 Automatic detection algorithm

The detection procedure starts with a band-pass filter (80-250 Hz; finite impulse response filter, $f_{stop1} = 70$ Hz; $f_{pass1} = 80$ Hz; $f_{pass2} = 250$ Hz; $f_{stop2} = 260$ Hz; stopband attenuation = -60 dB) applied to the iEEG recordings. The signals are filtered forward and backward to obtain zero-phase distortion. The filtered signals are then rectified, and the amplitude of each oscillatory cycle is measured by identifying each “peak” (local maximum) in the rectified data (Figure 2.1). Our algorithm then consists of choosing a threshold for the local maxima and identifying events that have a minimum number of oscillations above that threshold. More specifically, we look for a number of rectified peaks that exceed the threshold within a window of time, e.g. for 6 consecutive peaks, at least 5 are above the threshold.

Here, we propose two methods for choosing the threshold: (1) calculated based on the mean and standard deviation of the peak amplitudes and (2) determined via an iterative process to estimate the amplitude distribution of the background activity. We will refer to these as the

“non-iterative” and “iterative” methods, respectively. In the non-iterative method, the threshold is calculated directly from the peak amplitudes (local maxima) of the rectified filtered signal, defined as the mean plus a number of standard deviations. In the iterative process for threshold selection, we first create a histogram of the peak amplitudes (Figure 2.1E). We model the background activity of the channel with a gamma distribution, $f(x)$, which can be estimated as

$$f(x) = \frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-\frac{x}{\theta}}, \quad x \in (0, \infty)$$

where k and θ are the shape and scale parameters of the gamma probability distribution, respectively. In channels with HFOs, the presence of the high amplitude activity will cause an observable long tail at the upper end of this distribution, superimposed on the background activity. In the first iterative step, k and θ are estimated using the height of all peaks (via MATLAB function “gamfit”), which will include both background activity and HFOs. The shape and scale parameters are used to construct an estimated probability distribution $f(x)$ to represent the amplitude of background activity (Figure 2.2A). We then define a cutoff of $F(x) > 1 - \alpha$, where $F(x)$ is the cumulative probability distribution function of $f(x)$ (Figure 2.2B). Peaks with amplitude above the cutoff are excluded from the distribution and then k and θ are recalculated for the new improved estimate of the background activity, $f(x)$. Note that the cutoff value is tuned by a single parameter, α , related directly to the number of peaks that are removed in each step of the process; for example, if $\alpha = 0.01$, peaks that fall within the top 1% of the estimated $f(x)$ will be removed during each iteration. The threshold based on alpha is determined relative to the amplitude of the signal, rather than being an absolute threshold in terms of microvolts. For example, with $\alpha = 0.01$, a signal with a mean amplitude of 100 μV will have a higher threshold than a signal with a mean amplitude of 50 μV . The process is then repeated with the new estimate of $f(x)$ until the values of k and θ converge and no additional peaks are removed with

each iteration. At this point, $f(x)$ is the estimate of the underlying background distribution given a tolerance for false positives (α). We then define the threshold for detection based on the cutoff value of the final iteration (Figure 2.2B). The threshold will depend on the magnitude of α , with higher values of α leading to a lower threshold. In both approaches, the user must only provide one parameter to set the threshold, either the number of standard deviations above the mean or the tolerance for false positive detections, respectively. After threshold selection, all peaks with amplitudes above the threshold are marked, and events with at least 5 out of 6 consecutive peaks above the threshold are defined as HFOs.

We tested the detector performance as a function of the threshold (Section 2.2.5), and we also measured the robustness of the algorithm to changes in the shape of the amplitude probability distribution, the number of iterations for threshold convergence, and the required number of consecutive peaks above threshold (see Section 2.3.3).

2.2.4 RMS detector

The RMS detector (Staba et al., 2002) is based on the energy of the moving average of the root mean square amplitude of the filtered signal. In the original publication, the signal was bandpass filtered from 100-500 Hz, and segments of data in which the RMS value exceeded a threshold for at least 6 ms were marked as pre-qualified events. The threshold was defined as five standard deviations above the mean RMS value. Events less than 10 ms apart were joined together and considered as a single event. Finally, to be qualified as HFOs, the events were required to have at least six rectified peaks above a second threshold, which was three standard deviations above the mean of the rectified filtered signal. Originally, the detector was designed to identify HFOs in hippocampus and entorhinal cortex microwire recordings in humans (Staba et

al., 2002), and it was used in the studies of microelectrode recordings in temporal regions (Staba et al., 2007, 2004). The RMS detector is one of the most widely used automatic HFO detectors in published studies (Blanco et al., 2010; Gardner et al., 2007; Gliske et al., 2016; Zelmann et al., 2012).

Before applying the RMS detector to our dataset, we used the training data to optimize all parameters, in order to ensure the best possible performance. We implemented the same filter for both detectors (see Section 2.2.3), while the following parameters required optimization: duration of moving window to calculate RMS amplitude, minimum event duration, RMS threshold (1st threshold), rectified threshold (2nd threshold), and minimum time gap between consecutive candidate events. All parameters except the RMS threshold, which was selected from the receiver operation characteristic (ROC) curve, were optimized via the training channels (20% of all channels). We tested 4-10 ms for RMS window size; 4-24 ms for minimum event duration; 1-3.5 times the standard deviation above the mean for the rectified threshold; and 4-13 ms for minimum time gap between consecutive events. The HFO detection was performed in the training channels, and a single ROC curve was constructed for each parameter combination by varying the number of standard deviations for the RMS threshold. The parameter set that gave the best performance at the optimum RMS threshold (the point on the ROC curve that was closest to the upper left corner of the plot) was chosen as default: 10 ms for RMS window size, 6 ms for minimum event duration, one standard deviation for rectified peak threshold, and 13 ms minimum time gap between consecutive events. Then the RMS detector with the optimized parameters was applied to the 296 channels reserved for testing and the RMS threshold was optimized using a ROC curve in the same way as α for our detector (see Section 2.2.5). Note that these parameters were different from the original study of Staba et al. 2002.

2.2.5 Evaluation of detector performance

2.2.5.1 Performance measured from ROC curve

The performances of both automatic detectors were evaluated via a ROC curve, which is plotted using sensitivity and FPR under the strict validation scheme. To calculate sensitivity and FPR, true positives (TP) were defined as detected events that matched PosAND. False positives (FP) were automatically detected events that corresponded to a visually marked baseline. True negatives (TN) referred to the NegBASE segments containing no automatically-detected HFOs. False negatives (FN) were undetected PosAND. Then sensitivity was defined as $TP/(TP+FN)$, and FPR was defined as $1 - TN/(FP+TN)$. The ROC curves were created by varying α for our detector and varying the first RMS threshold (specifically the number of standard deviations) for the RMS detector. The optimum values of α and the threshold were chosen by identifying the best balance between sensitivity and FPR, the point on the ROC curve that was closest to the upper left corner of the plot. The performance at this optimum value illustrates the best possible performance of the detector. In addition, the area under the ROC curve (AUC), which ranges from zero to one (one indicating perfect performance), is used to represent the detector's overall performance.

Note that FPR reflects only the falsely detected events within the visually-marked baseline segments; therefore, this measure does not account for the events that are neither HFOs (PosANY) nor baseline (NegBASE), which are likely to be moderate amplitude events. Accordingly, the false detection rate (FDR) was introduced as the number of detected events that did not overlap with PosANY events, including false positives, divided by the total number of detected events.

Optimization of α or the threshold was performed in three different ways: (1) across channels, with α or the threshold optimized for each individual channel, representing the best possible outcome of the detector, (2) across patients with a single value for all channels within each patient, illustrating the robustness of the detector when applied to different patients, and (3) grouping all events together, using a single value for all channels.

2.2.5.2 Performance measured from cross-validation

Cross-validation was used as a more stringent evaluation of the detection performance. Ideally, the detector would be optimized for each channel based on visual markings from a short amount of data within that particular channel, at least one minute free of artifact, and those settings would be used for the remaining data. However, this was not possible here because the dataset contained only one minute of data for each channel. We therefore measure the performance of the detector using independent subsets of channels for training (optimizing) and testing. This implementation method would only be used if visual detection is not available for all channels. It illustrates the performance of the detector when a single parameter setting is applied to new data, without taking the characteristics of the new signals into account.

We apply an eight-fold cross-validation technique, where 296 channels are randomly assigned to eight subsample groups with 37 channels each. For each iteration, one subsample group is selected for testing and the seven remaining groups are used for optimizing. The average optimum α value from all seven training groups (derived from ROC curves) is used to perform automatic detection in the remaining testing group. This procedure is repeated eight times, so each subsample group is tested once.

2.3. Results

2.3.1 Detector performance

2.3.1.1 Comparing the performance across channels

The ROC curves for each channel were created using a range of α values or the number of standard deviations to determine the amplitude threshold. The optimum threshold for each channel (the threshold on the ROC curve closest to the upper left-hand corner, representing the best performance) was selected, and then the associated performance was vertically averaged across all 296 channels (Figure 2.3). Note that this performance, in which the threshold was optimized for each channel, represents the best possible detection performance for the detector. For our detector with iterative threshold selection, the averaged AUC was 0.995. The average sensitivity was 99.6%, and the FPR and FDR were 1.1% and 37.3%, respectively (Table 2.1). For the non-iterative method, the sensitivity, FPR, FDR and AUC were 99.5%, 1.1%, 36.9%, and 0.995, respectively (Table 2.2). The RMS detector had 98.9% sensitivity, 0.7% FPR, and 34.6% FDR, with an area under the ROC curve of 0.995 (Table 2.3). The difference in detection performance was tested with a Wilcoxon signed rank test. The test failed to reject the null hypothesis when comparing the two different thresholding schemes for our detector. This was expected, as the only difference was how the threshold was selected. However, the sensitivity and AUC of our detector, with the iterative method, were statistically significantly higher than the RMS detector, at $p < 0.01$ and $p < 0.05$, respectively, but the FDR was significantly higher, $p < 0.05$. Our detector, with the iterative method, exhibited error-free performance in many channels. In 257 of 296 total channels (86.8%), all visually marked PosAND events were identified at the optimum threshold (100% sensitivity), and 187 of them (63.2% of all channels)

had all PosAND events detected with zero FPR, representing perfect performance (AUC =1). Only 9 of the 296 channels (3.0%) had sensitivity below 95%.

The optimum RMS threshold ranged from -1.7 to 8.5 (mean 2.15) standard deviations above the mean RMS value. Similarly, in our detector with the non-iterative approach, the standard deviation of the rectified peak amplitude ranged -0.5 to 10.0 (mean 2.41). The negative values of standard deviation for the optimum threshold need to be interpreted with caution because this indicates that the optimum threshold was lower than the mean value for the signal.

2.3.1.2 Comparing the performance across patients

We then assessed the performance of each detector when a single value of α or number of standard deviations was used for all channels within a patient's dataset. This illustrates the robustness of the detector when it is applied to channels with different amplitudes and background characteristics. All ROC curves within each patient were averaged together with equal weighting, and the average AUC, sensitivity, FPR, and FDR were obtained for each of the 19 subjects (Figure 2.4). Here, a single α value or number of standard deviations cannot optimally represent all channels due to the variation in background activity in each channel. In channels with an active background, the detector requires a higher α (lower threshold) to capture more peaks because there are random peaks from background activity mixed in with the detected peaks of HFOs, while HFOs can be detected with a lower α (higher threshold) in quiet channels. Therefore, the performance of our detector with iterative threshold selection decreased (0.980 AUC, 93.6% sensitivity, 5.2% FPR, and 66.9% FDR; see Table 2.1) compared to the case where it was optimized for every single channel. A similar trend can be observed for both the non-

iterative method (0.976 AUC, 92.7% sensitivity, 6.0% FPR, and 66.6% FDR; see Table 2.2) and the RMS detector (0.973 AUC, 92.2% sensitivity, 5.4% FPR, and 64.0 % FDR; see Table 2.3).

2.3.1.3 Comparing the performance across all events

Both detectors were tested when a single α value or number of standard deviations was used across all 296 channels. Again, all channels were weighted equally. As in the previous section, we expected a decrease in detection performance because the level of background activity varies across channels and subjects. For our detector with iterative threshold selection, the optimum value of α was 0.065. At this parameter setting, it identified 75.3% of PosAND events (5696 out of 7561) and 65.4% of PosANY events (8051 out of 12316); further, 95.1% of baseline segments (NegBASE, 40985 out of 43101) were free of detected HFO events. The performance indicates that reliable detection was achieved by using one single value of α to iteratively select a threshold in all channels, with an average sensitivity of 93.1%, FPR of 6.4%, FDR of 71.6%, and AUC of 0.981. Similar outcomes were found for our detector with the non-iterative approach (92.3%, FPR of 6.1%, FDR of 68.0%, and AUC of 0.979 at the optimum threshold, which was 0.80 standard deviations above the mean of rectified peak amplitude). The RMS detector also illustrated comparable accuracy, with an average sensitivity of 92.0%, FPR of 7.0%, FDR of 66.2% and AUC of 0.976 at a threshold equal to the mean plus 0.63 standard deviations of RMS amplitude. The ROC curves for both detectors are shown in Figure 2.5.

2.3.1.4 Cross-validation of the automated detection algorithm

In the previous sections, optimization of the threshold was based on the prior knowledge of all visually detected events as a “true” reference via ROC curve. However, in practice, it may

not be realistic to complete visual detection on all channels, especially in the case of intracranial recordings, which often have over one hundred channels. Cross-validation techniques serve as a measure of this worst-case scenario, in which the detector is optimized on one set of channels and the parameters are applied to a different set of channels. This is the most stringent test of performance, as each channel's background characteristics and HFO amplitude are independent from those of the other channels. We implemented eight-fold cross-validation, in which seven subgroups of channels were used to optimize the threshold and subsequently tested on the remaining channels. Note that this is not typically how the detector would be implemented. In general, it is best to use a portion of the data within each particular channel to optimize the parameters for that channel, e.g. use the first minute of a ten-minute recording.

For our detector with iterative threshold selection, the optimum α value of the eight cross-validation iterations ranged from 0.036 to 0.038. The average sensitivity, FPR, and FDR were 80.0%, 1.3%, and 49.4%, respectively. In the case of our detector paired with non-iterative thresholding, the performance decreased to 64.4%, 0.4%, and 31.2% for sensitivity, FPR, and FDR respectively. The optimum number of standard deviations for the threshold ranged from 2.35 to 2.5. Similarly, the RMS detector had 65.8% sensitivity, 0.5% FPR, and 33.4% FDR. The optimal thresholds for the RMS detector ranged from 2.10 to 2.20 standard deviations above the mean RMS. There was large decrease in sensitivity for both detectors when the threshold was determined using the amplitude of the signal (mean plus number of standard deviations). These results indicate that the iterative method for threshold selection is more robust than directly using the statistics of the amplitude because it adapts to the characteristics of each channel.

2.3.2 Parameter optimization to reduce the false detection rate

When optimizing α individually for each channel, our detector demonstrated very high sensitivity, but the FDR was 37.3%. False detection of events remains a barrier to the practical implementation of automatic algorithms. There are several means to reduce the FDR, e.g. applying post-processing steps (Amiri et al., 2016; Burnos et al., 2014; Cho et al., 2014; Gliske et al., 2016) or using human validation (Crépon et al., 2010; Gardner et al., 2007; Staba et al., 2002). Here we propose another approach, in which α is optimized based on FDR instead of FPR.

More specifically, we optimized the threshold using the precision and recall curve, which is the plot of sensitivity and 1-FDR, rather than the ROC curve. Here, the optimum threshold was defined as the point on the curve that was closest to the upper right corner of the plot. Our detector with iterative threshold selection achieved a sensitivity of 89.1% (originally 99.6%), but the FDR was reduced by more than half to 16.5% (from 37.3%; see also Table 2.1). There were only 20 channels (6.8%) in which the FDR exceeded 50% and only 2 channels (0.7%) with sensitivity less than 50%. Both the non-iterative method and RMS detector performed similarly; see also Table 2.2 and 2.3. Note that, as with all other results presented here, the only parameter involved in the optimization of our detector was α or the number of standard deviations. Therefore, adjusting this single parameter enables us to choose whether we prioritize high sensitivity or high specificity.

2.3.3 Algorithm design

In addition to the parameter α , the design of the automatic detection algorithm involved several important elements. For example, we also chose a probability distribution to model the amplitude of the background activity, the maximum number of iterations for convergence of the

threshold, and the number of consecutive oscillations required to define an HFO event. However, we found that changes to these elements did not have a significant effect on the detection performance, as demonstrated below.

2.3.3.1 Probability distribution for the background activity

A key step in the detection algorithm is the estimation of the amplitude probability distribution for the background activity. To estimate the true distribution with minimal influence from HFOs, the 33 channels with the fewest visually detected HFOs (i.e. PosANY was less than three events) were used as representative of background activity. We compared gamma distributions to the amplitude histograms of the rectified band-pass filtered iEEG recordings. A Kolmogorov-Smirnov test showed that 31 channels were consistent with a gamma distribution ($p < 0.05$). Furthermore, a Q-Q plot showed that the peak amplitude distribution was a good match to the reference distribution at low quantiles (data not shown). At high peak amplitudes, however, the sample distributions deviated from the reference. Therefore, both the Kolmogorov-Smirnov test and the Q-Q plot confirmed our hypothesis that without HFOs, the gamma distribution would be a good model for the peak amplitude distribution of the background activity. We performed the same tests with a zero-truncated normal distribution to determine whether it could be used as a model of the probability distribution of the peak amplitude. Interestingly, the two distributions gave comparable performance, which suggests that the detection outcome is robust to a change of the model distribution. For all results presented here, we selected a gamma distribution as the default model for the background activity.

2.3.3.2 Number of consecutive peaks above threshold

For the purposes of detection, the definition of an HFO may vary slightly depending on the automatic detector algorithm and the study. Generally, an HFO is required to have 3-5 consecutive cycles above a threshold (Julia Jacobs et al., 2008; Staba et al., 2002). Here we used the number of local maxima (“peaks”) in the rectified filtered iEEG recordings to count the oscillations, with one oscillation consisting of two consecutive rectified peaks. We tested performance while varying this parameter from 4 to 8 consecutive peaks. When a high number of consecutive peaks was required, the optimum α value was increased (indicating a lower threshold value), enabling events with more consecutive peaks to be detected. Our results indicated that detector performance was robust to changes in the required number of oscillations; because the threshold was recalculated each time, α increased as a function of the number of consecutive peaks. As shown in Figure 2.6, the difference between each condition was primarily in the FDR value (more than 3% difference between the top three conditions with highest sensitivity, 42.5%, 39.3%, and 40.3%), while the sensitivity difference among the best conditions was less than 0.5% (99.7%, 99.4%, and 99.2% in the 4 out of 5, 5 out of 6, and 6 out of 7 conditions, respectively). We chose 5 out of 6 consecutive peaks within the test window, i.e. for a group of 6 consecutive peaks, at least 5 needed to exceed the threshold, as a default setting for our detector because it offered the best balance of sensitivity and FDR at the optimum threshold.

2.3.3.3 Number of iterations

Our algorithm employed an iterative process together with the estimate of the peak amplitude distribution, which enabled us to optimally adjust the threshold based on the selected α . After several iterations, e.g. up to 10-11 iterations in channels with very high α value, the estimate of the peak amplitude distribution converged, and the stable cutoff value was used as

the detection threshold. Once this occurred, additional iterations of the algorithm did not affect the final estimate of the distribution. The rate of convergence varied and depended on the magnitude of α and the peak amplitude probability distribution. In our study, 15 iterations were sufficient for all channels while still minimizing the computational cost. With our desktop system (CPU: intel i7-4790k, 16 GB of RAM), one minute of iEEG with 20 channels recorded at a 2 kHz sampling rate required approximately 1-2 seconds for filtering and rectification; after that, the detection procedure took an additional 2-3 seconds per channel for a single value of α . To calculate the ROC curves, 40-60 different values of α were tested, which took 4-6 minutes for each channel.

2.3.4 Optimum alpha values

Alpha (α) was the only parameter in our algorithm that had a substantial impact on the accuracy of HFO detection. Since α is related to the threshold, it affects the number of detected events. Lowering α , which raises the threshold, reduces the detection sensitivity together with the FPR and FDR. In this case, the detected events have a higher probability of being real HFOs, but overall fewer HFOs are detected. On the contrary, increasing α , which decreases the threshold, allows more falsely detected events which results in an increase in sensitivity; however, the FDR and FPR also increase as more peaks are detected.

In section 2.3.1, we showed that the detection performance was best when α was individually optimized for each channel, rather than using a single value of α for multiple channels. Recall that alpha determines the percentage of peaks in each iteration that are marked by the detector as being possibly associated with an HFO. Therefore, in a channel with a quiet background and prominent HFOs, we can select very few peaks (high threshold, low alpha) and

still detect all HFOs. On the other hand, when there is an active background, the peaks associated with HFOs are mixed with “noisy” peaks, which may be approximately the same height. In this case, the detector must mark many more peaks as possibly associated with HFOs (high alpha) in order to detect the events. Across all channels, the optimum alpha values ranged from 0.0001 to 0.135. Most channels had α values on the order of 10^{-3} when we optimized the detector with an ROC curve (Figure 2.7A). When α was optimized using a precision and recall curve, its value tended to be lower, which increased the threshold and reduced the number of falsely detected events, (Figure 2.7B); however, there were some channels in which α was significantly higher. Overall, detection performance in those channels was poor, but these channels tended to have a very few or only a single visually marked HFO.

2.4. Discussion

Here we have presented a simple automatic detection algorithm for HFOs that consists of filtering and rectifying the data, then counting consecutive peaks above a threshold. An optimum threshold can be determined using an iterative process to estimate the statistical properties of the background activity, or it can be based on the mean and standard deviation of the peak amplitudes. The algorithm requires optimization of only a single parameter, greatly simplifying the implementation of automatic detection. Importantly, we have shown that the performance of this detector is comparable to or better than the most commonly used published algorithm, which makes it a suitable technique for the initial detection of candidate HFOs. It could also be paired with an algorithm for the automatic rejection of artifacts and false positive detections; because these post-processing steps increase the number of parameters and overall complexity, it is advantageous to use the simplest possible approach for the initial detection of candidate events.

The iterative and non-iterative methods used in our detection algorithm gave comparable results in most cases, but sensitivity using iterative threshold selection was superior in a cross-validation test. When the threshold was optimized for each individual channel, both approaches performed similarly: average sensitivity, FPR, and FDR were approximately 99%, 1%, and 37%, respectively. This result was expected, as the only difference was the method of threshold selection; therefore, the ROC curves of both cases were theoretically the same. A slight difference in the observed performance between the two approaches, which was not statistically significant, resulted from differences in the smoothness of the ROC curve. Using a fixed threshold for all channels caused a decrease in detection performance, due to the variation between channels. However, our detector still provided satisfactory results with approximately 93% sensitivity, 6% FPR, and 70% FDR for both the iterative and non-iterative approaches. The biggest difference between the two approaches occurred for the cross-validation test. The iterative method yielded better sensitivity (80.0% compared to 64.4% for our detector with the non-iterative method). While the non-iterative method resulted in lower FPR and FDR, the sensitivity is the most crucial factor, as false detections can be reduced by post-processing procedures. Therefore, when optimizing channels independently, the iterative method is more robust to variation across channels compared to the non-iterative method in which the threshold is directly calculated from the mean and standard deviation of the amplitude.

The two thresholding schemes each have strengths and weaknesses. The non-iterative method is the simplest method we tested, and it is also the least computationally expensive. However, our results showed that there is no intuitive way to choose the number of standard deviations for the threshold. When the threshold was individually optimized for each channel, the optimum number of standard deviations ranged from -0.5 to 10.0 (mean 2.41). Thirteen of

296 channels had a negative number of standard deviations as the optimum threshold, resulting in a threshold that was below the mean. In those channels, the overall amplitude of the signal was significantly higher than other channels, causing the mean and standard deviation to be much larger than usual. The amplitude distribution was non-normal and highly skewed; therefore, the mean and standard deviation were not appropriate metrics for the selection of a threshold. Overall, the wide range and possibility of needing negative values presents an obstacle for appropriate selection of a threshold. On the other hand, the iterative method is more complex and computationally intensive, but it can adapt to the characteristics of the data in each channel. It therefore performs better when the detector is optimized on one dataset and then applied to an independent set of data. Moreover, the optimal α values ranged from 0.0005 to 0.145 (mean 0.037); selection of this parameter is easier, as its range is narrower, and its value is always positive.

The detection performance of our algorithm was either comparable or superior to the RMS detector when tested using the same iEEG data, visually-marked HFOs, and evaluation scheme. It was comparable when testing the best possible performance using a ROC curve, but the sensitivity for the iterative method was superior in the cross-validation test (80.0% for iterative method, 65.8% for RMS). This is somewhat expected because the RMS detector was initially designed for use with microwire data, and it has been shown that the characteristics of high frequency activity depend on electrode size (G. a. Worrell et al., 2008). However, the RMS detector has been successfully implemented in various recording schemes, including microwire electrodes (Staba et al., 2002, 2004) and macro- and microelectrodes (Gliske et al., 2016; Zelmann et al., 2012). Importantly, the results of our tests demonstrate that good performance is dependent on proper optimization of all parameters, not just the threshold. Optimization of four

parameters using the training channels was necessary to achieve detection performance similar to our detector. Here, most of the optimum parameters that were used for the RMS detector were different from the original study, especially the RMS threshold. The average RMS threshold across the 296 test channels was only 0.80 standard deviations above the mean, while the original study recommended 5 standard deviations. Similar to the non-iterative method, the optimal RMS threshold had a wide range of values, including negative values of the standard deviation. If we had limited the standard deviation to positive values to ensure that the HFO amplitude was greater than the mean, the performance of the RMS detector would have been significantly worse.

Taking all these results into consideration, we recommend the following guidelines for implementation of an automated detector. The best performance will be achieved if visual detection is first performed on a small segment of data from each channel, approximately one minute. These visual detections can be used to optimize α or the number of standard deviations for each channel, which can then be applied to detection in the remainder of the dataset. In this situation, we anticipate that our detector and the RMS detector will perform similarly, but the four parameters of the RMS detector must be optimized on an independent dataset prior to this procedure. Our detector does not require this initial optimization step. On the other hand, if visual detection is only available for a subset of channels, the parameter α can be optimized using that group of channels and then applied to the rest of the dataset. This is analogous to the cross-validation test we performed. Therefore, in this case, the best performance will be achieved with our detector, using the iterative method of threshold selection.

The high number of falsely detected events remains a challenge in the implementation of automated detection. While threshold optimization based on the ROC curve yields high

sensitivity and extremely low FPR, this low FPR must be interpreted with care because it does not reflect all falsely detected events. The FPR includes only the detected events that lay within visually marked baseline segments. Detected events that occurred within unmarked sections of the signal (not clearly baseline and not an HFO) were not included in this number, but instead contributed to the FDR. The FDR for both detectors was as high as 36.1%, which means that approximately one third of detected HFO events occurred outside of the PosANY visual markings. Several techniques can be implemented to reduce this value. Here, we optimized the threshold based on a precision and recall curve, rather than using a ROC curve. When we did this for the iterative method, the average sensitivity across all channels decreased from 99% to 89%, but the FDR showed a much greater decrease, from 36% to 17%. This trade-off between sensitivity and specificity may be desirable in some applications, such as the analysis of long datasets in which visual rejection of falsely detected events is not feasible. Note that the detection performance in several channels suffered due to very few HFOs (there were 17 channels with only one visually-marked HFO), a very active background with continuous high frequency activity, and the detection of artefactual waveforms (the sensitivity was below 50% in two channels, and the FDR exceeded 50% in 27 out of 296 channels). Another option to reduce the FDR and detection of artifacts is to apply a post-processing step to eliminate falsely detected events and leave only “true” HFOs. This can be done either automatically, using an artifact rejection algorithm (Amiri et al., 2016; Burnos et al., 2014; Cho et al., 2014; Gliske et al., 2016) or data classification via clustering (Blanco et al., 2010; Malinowska, Bergey, Harezlak, & Jouny, 2015), or manually with supervision by experts.

There are two possible limitations to the automatic algorithm presented here. First, for the iterative threshold selection, we assumed that HFOs were rare and that the high amplitudes

associated with these events would be superimposed on a stable distribution of the background amplitude. If HFOs occur frequently, the estimation of the background activity will be inaccurate and will result in a drop of performance. However, channels with this characteristic are typically quite difficult to interpret, even for human reviewers, so it is perhaps not surprising that it also provides a challenge for automatic detectors. Moreover, the dataset used for testing contained such channels, and this method still demonstrated good overall performance. Second, because the detection relies on identification of local maxima in the rectified filtered EEG signals, the accuracy may be sensitive to high amplitude events, including artifacts. On the other hand, the requirement that each event have a number of consecutive peaks above the threshold will help reduce the detection of short, sharp transients, as they usually cause only a few full oscillations in which the amplitude is high enough to exceed the threshold in the filtered data. It is important to note that the dataset we used for performance testing contained epileptiform discharges and other artifacts, and we compared the automated detection results to a gold standard of visual detection.

The iterative procedure presented here has some similarities to the one utilized in the MNI detector (Zelmann et al., 2012), however there are several important differences. First, the MNI detector incorporated two different algorithms for HFO detection, depending on whether the iEEG channel contained continuous high frequency activity. The iterative procedure for determining the threshold was used only in cases where an insufficient amount of baseline activity was detected. Here, we use a single algorithm for all channels. Second, the iterative procedure in the MNI detector is based on the moving average of the RMS within a window of data. In our algorithm, the iterative choice of threshold is based on detection of local maxima in the filtered, rectified data. This enabled us to detect high amplitude oscillations directly, rather than indirectly through energy measures. Third, the MNI detector requires the optimization of

many parameters. The baseline detection requires selection of a time window, amount of overlap, and a threshold for the wavelet entropy; the HFO detection relies on selection of time windows for calculation of the RMS and the moving average of the RMS, a threshold for the cumulative distribution function, an energy threshold, and a minimum amount of time between events. Here, our simplified detection algorithm used only four parameters (α , the choice of background probability distribution, the max number of iterations, and the number of required consecutive oscillations), and we were able to show that the choice of α was the only one that affected the performance of the detector. While our algorithm provides levels of accuracy that are similar to other published detectors, we feel that the value of our detector lies in the simplicity of the optimization, which is based on only one parameter. This will enable it to be applied broadly to a variety of data types.

Roehri et al. (2016) also reported a technique to estimate the background activity for the purposes of HFO detection. A continuous wavelet transform was applied to the data and distributions of the real and imaginary coefficients were created at each frequency. It was then assumed that the background noise could be represented by a Gaussian fit to the central portion of those distributions. This enabled estimation of the background activity without requiring manual selection of a baseline segment of data. The distributions for each frequency were normalized using this technique, thus whitening the data. Whitening of electrophysiological data can prevent HFO detection from being dominated by low frequencies, which naturally have higher amplitude. This procedure is related to the non-iterative approach presented here, but we create amplitude distributions using local maxima (“peaks”) in the filtered data, rather than using all data points. Moreover, our results suggest that iterative estimation of the background provides a distinct advantage in choosing a threshold; therefore, combining the whitening methods

suggested by Roehri et al. with iterative background estimation may be a powerful technique for HFO detection. On the other hand, if the goal is to match visually marked HFOs, whitening may not be desirable, as humans viewing bandpass filtered data will naturally be biased toward low frequency events.

There are several remaining barriers to achieving automated detection accuracy which is sufficient for clinical use. The physiological mechanism underlying HFOs is not fully understood, and therefore detection is based on empirical definitions (Engel & da Silva, 2012; Jefferys et al., 2012). To complicate matters further, the shape of the HFO waveform can vary depending on the position of the electrode and its distance from the generating tissue (György Buzsáki et al., 2012). This implies that a rigid template of HFO shape may be insufficient for detection. While these ideas can be loosely factored into visual detection, a precise physiological definition of HFOs will be crucial for the next generation of automatic detection. Until that definition is understood, however, visually-detected HFO events will act as the gold standard for evaluating the performance of automated algorithms.

Furthermore, the distinction between physiological and pathological HFOs is not understood. There is evidence that a high rate of ripples and fast ripples is indicative of the seizure onset zone, even without attempting to distinguish between pathological and physiological events (Bragin et al., 2010; Maeike Zijlmans et al., 2012). Both ripple and fast ripple rates are higher in seizure onset zones, confirmed by surgical outcome (Julia Jacobs et al., 2010). Other studies have suggested that HFOs should be divided into two categories based on frequency (Bragin, Engel, Wilson, Fried, & Buzsáki, 1999; Staba et al., 2002), where ripples are considered to be physiological events and fast ripples are associated with epileptogenicity (Staba et al., 2007). However, to differentiate between pathological and physiological HFOs, frequency

alone is insufficient (Matsumoto et al., 2013). In the present study, the default parameters were configured for the detection of ripples, but we expect the algorithm to work for other frequency bands as well, simply by changing the bandpass filter. By iteratively estimating the background amplitude distribution and choosing the threshold, the procedure adapts to the characteristics of the data, regardless of the frequency band or recording modality.

In conclusion, this simple algorithm can be used to automatically detect HFOs with a high degree of accuracy, confirmed by comparison to visually marked events. When compared directly using the same dataset, our detector's performance equaled or exceeded the most commonly used HFO detection method based on the RMS amplitude. The statistical detection of HFOs at different confidence levels was done by estimating the peak amplitude probability distribution of the background activity and then identifying oscillatory iEEG events with amplitudes that were statistically higher than the background. Because the algorithm requires the optimization of only one parameter, related to the percentage of allowable false positive events, this technique can be applied consistently across data from different research centers and different recording modalities. Due to its high detection sensitivity and simple optimization procedure, our detector provides advantages over other techniques for the initial detection of candidate HFO events. It is suitable to be paired with an algorithm for the automatic rejection of artifacts and false positives, or it may be used with human validation. Overall, this type of automated algorithm is less subjective and much faster than visual HFO detection, and it has the potential to be a powerful tool for the assessment and localization of epileptic activity.

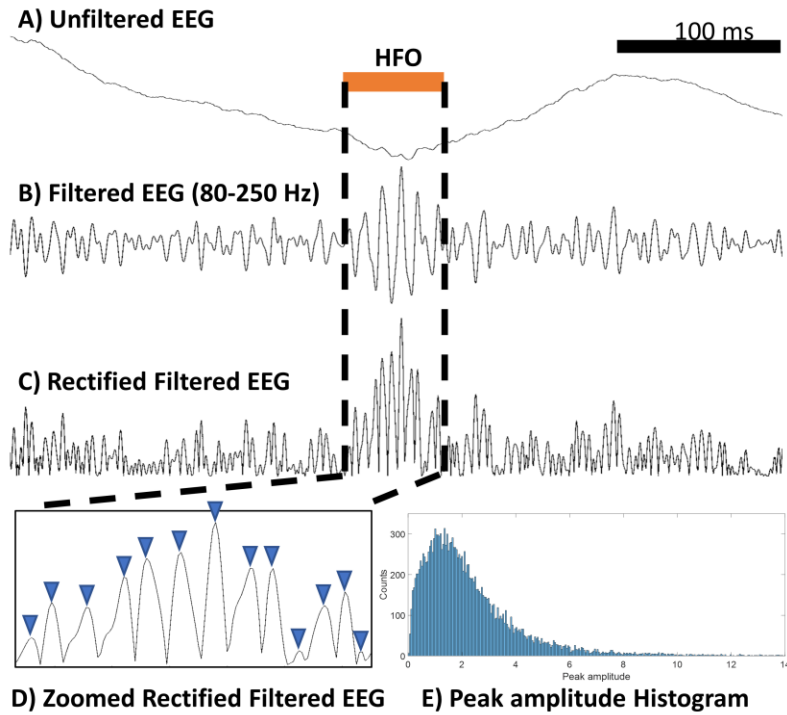


Figure 2.1 Detection procedure and an example of an HFO event (orange bar).

(A) The algorithm starts with unfiltered EEG signal. (B) Then the EEG signal is bandpass filtered from 80-250 Hz. (C) The filtered signal is then rectified. (D) We define “peaks” as local maxima in the rectified, band-pass filtered data (blue triangles). (E) A histogram of peak amplitudes resembles a gamma distribution. HFOs will be superimposed upon this distribution, causing an extraordinarily long tail on the right side.

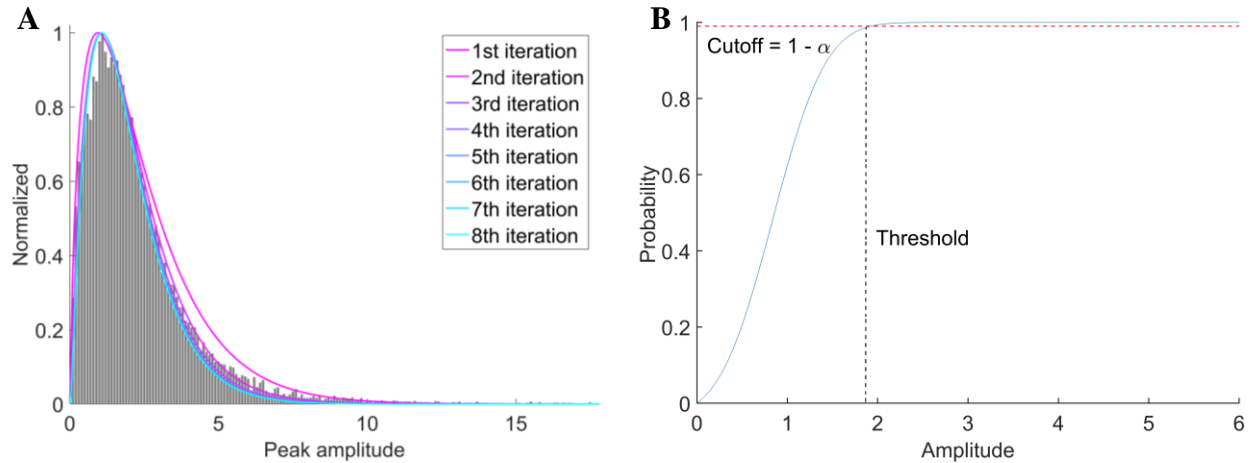


Figure 2.2 Iterative procedure for threshold optimization.

(A) The amplitude probability distribution for a single channel, including both HFOs and background activity (gray bars). Lines represent the estimated background distribution $f(x)$ over 8 iterations. For each iteration, the amplitude values that are higher than the cutoff $F(x) > 1 - \alpha$ are removed before re-estimation of the new distribution. After several iterations, the estimate of the background distribution converges, providing an optimum threshold for the detection of HFOs. (B) The cutoff value is based on the cumulative probability distribution function $F(x)$ (blue line), and it is directly related to the amplitude threshold in the iEEG data.

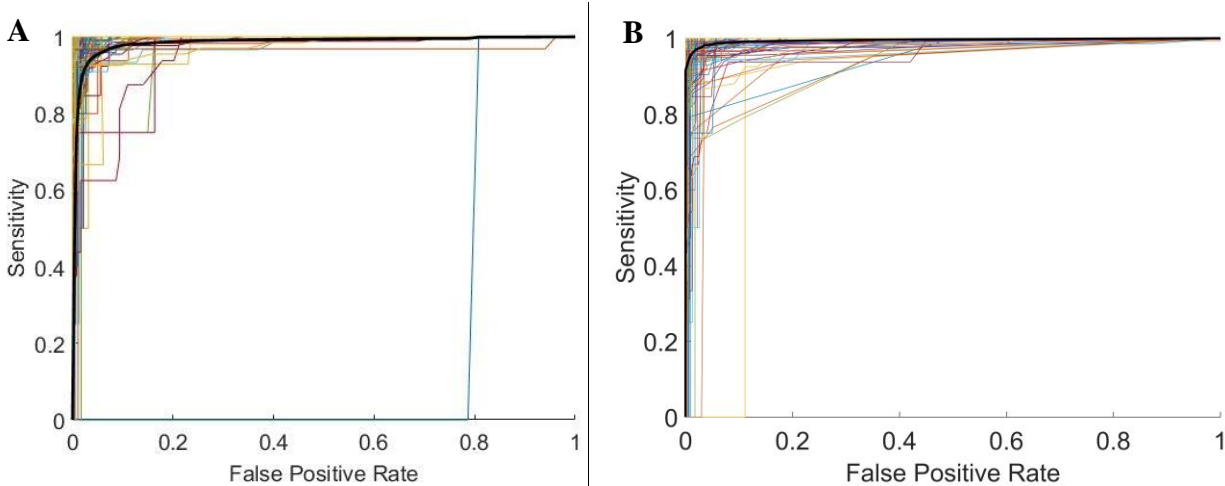


Figure 2.3 Detector performance based on ROC curves when the threshold was individually optimized to each channel.

The subfigures show ROC curves for all 296 channels (thin lines) and the vertical average ROC curve (thick black line) of (A) our detector with iterative threshold selection and (B) RMS detector. The reported optimum threshold was the point on each ROC curve that was closest to the upper left corner.

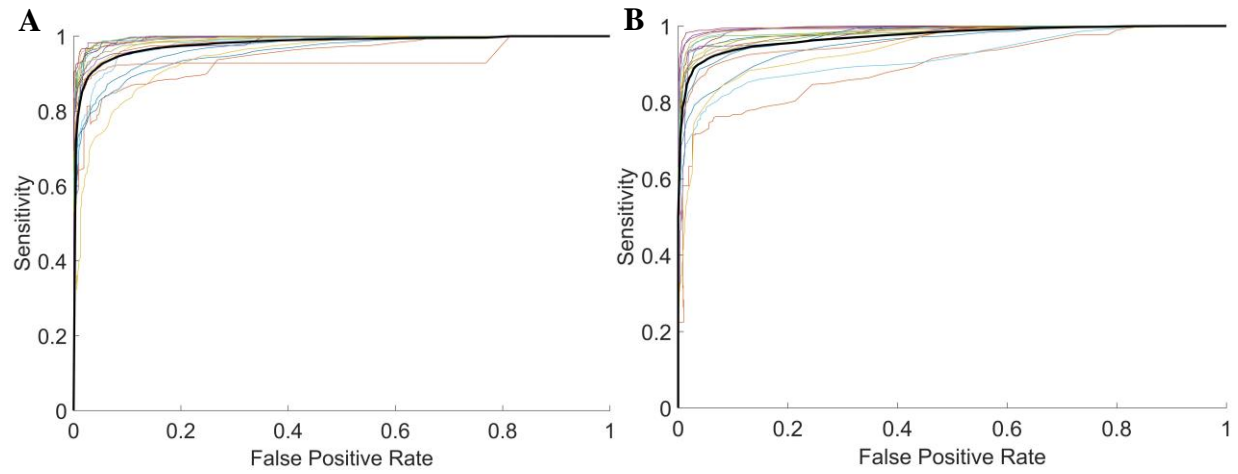


Figure 2.4 ROC curves when a single α value was applied to all channels within the same subject in (A) our detector with iterative threshold selection and (B) the RMS detector.

The thin lines represent the average of all channels within a subject, and the thick black line is the vertical average of all equally weighted 19 patients.

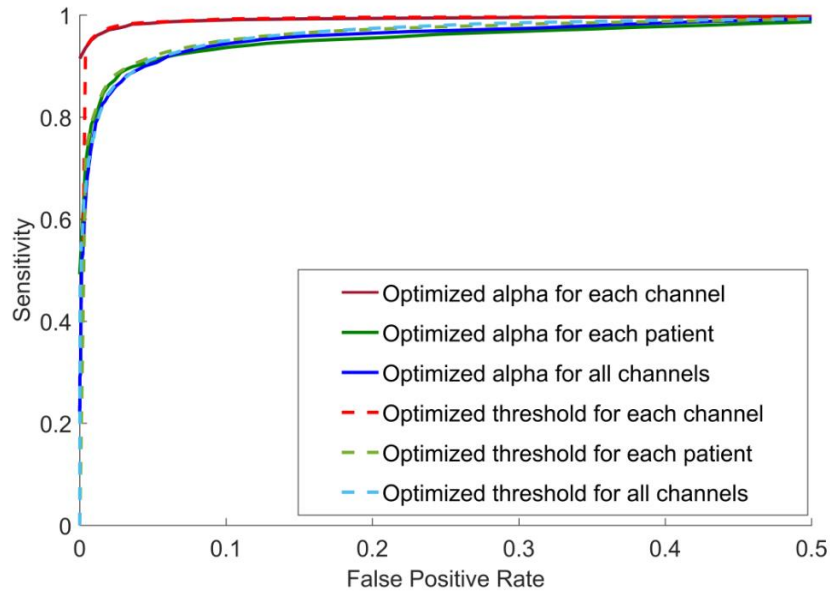


Figure 2.5 A comparison of ROC curves when both automatic detectors are tested in three different ways.

The solid lines represent the ROC curves for our detector, while dashed lines show the results for the RMS detector. First, the threshold was optimized for each individual channel and averaged across all channels (red lines). Second, averaged across patients with a constant threshold for all channels within each patient (green lines). Third, a single threshold value for all channels (blue lines).




Figure 2.6 A comparison of Sensitivity, FPR, and FDR when α was individually optimized for each channel, as a function of the number of consecutive peaks for HFO detection.

Results are shown for nine different criteria for the number of peaks required for each HFO, including 4 consecutive peaks (4/4), 4 out of 5 consecutive peaks (4/5), and up to 8 consecutive peaks (8/8). Recall that two “peaks” (local maxima) in the rectified, filtered data constitute one full oscillation. As a default configuration, we chose 5 out of 6 consecutive peaks above threshold (5/6) as it resulted in marginally better performance than the other choices. Note that better detection performance corresponds to high sensitivity, but low FPR and FDR.




Figure 2.7 Histogram of α values at optimum thresholds for all 296 channels.

(A) α values when the threshold was optimized using sensitivity and FPR. (B) α values when the threshold was optimized using sensitivity and FDR.

Table 2.1 Detection statistics of our detector with iterative threshold selection at the optimum threshold.

Sensitivity, FPR, and AUC were computed based on strict validation (PosAND and NegBASE); FDR was calculated with open validation (PosANY).

	Sensitivity (%)			FPR (%)			FDR (%)			AUC
	Mean \pm SD	Median	Min, Max	Mean \pm SD	Median	Min, Max	Mean \pm SD	Median	Min, Max	
Average across 296 channels (Single α for each channel)	99.6 \pm 1.4	100.0	87.5, 100.0	1.1 \pm 5.1	0.0	0.0, 80.8	37.3 \pm 30.6	35.7	0.0, 96.5	0.995
Average across 19 patients (Single α for each patient)	93.6 \pm 3.7	94.7	86.0, 98.2	5.2 \pm 2.9	4.4	1.7, 1189	66.9 \pm 10.3	68.7	49.1, 79.7	0.980
All events pooled together (Single α for all channels)	93.1	-	-	6.4	-	-	71.6	-	-	0.981
Average across 296 channels (Optimized threshold by FDR)	89.1 \pm 13.8	94.4	33.3, 100.0	6.1 \pm 23.6	0.0	0.0, 100.0	16.5 \pm 20.6	9.0	0.0, 94.4	-
Eight-fold Cross-validation	80.0 \pm 3.4	80.3	74.9, 84.8	1.3 \pm 0.6	1.1	0.6, 2.6	49.4 \pm 5.1	48.4	41.8, 55.6	-

Table 2.2 Detection statistics for our detector with non-iterative threshold selection at the optimum threshold.

Sensitivity, FPR, and AUC were computed based on strict validation (PosAND and NegBASE); FDR was calculated with open validation (PosANY).

	Sensitivity (%)			FPR (%)			FDR (%)			AUC
	Mean \pm SD	Median	Min, Max	Mean \pm SD	Median	Min, Max	Mean \pm SD	Median	Min, Max	
Average across 296 channels (Threshold optimized for each channel)	99.5 \pm 1.4	100.0	87.5, 100.0	1.1 \pm 5.0	0.0	0.0, 79.5	36.9 \pm 29.2	36.1	0.0, 96.5	0.995
Average across 19 patients (Single number of SD used for threshold for each patient)	92.7 \pm 4.3	94.7	85.0, 98.8	6.0 \pm 3.8	4.7	1.5, 16.0	66.6 \pm 9.7	67.5	46.9, 83.2	0.976
All events pooled together (Single number of SD used for threshold for all channels)	92.3	-	-	6.1	-	-	68.0	-	-	0.979
Average across 296 channels (Optimized threshold by FDR)	88.7 \pm 13.6	92.7	33.3, 100.0	1.0 \pm 9.4	0.0	0.0, 79.5	16.9 \pm 20.5	9.7	0.0, 94.6	-
Eight-fold Cross-validation	64.4 \pm 4.3	64.4	58.9, 71.2	0.4 \pm 0.3	0.3	0.2, 1.2	31.2 \pm 6.4	30.4	22.9, 42.0	-

Table 2.3 Detection statistics of RMS detector at the optimum threshold (Staba et al., 2002).

Sensitivity, FPR, and AUC were computed based on strict validation (PosAND and NegBASE);

FDR was calculated with open validation (PosANY).

	Sensitivity (%)			FPR (%)			FDR (%)			AUC
	Mean \pm SD	Median	Min, Max	Mean \pm SD	Median	Min, Max	Mean \pm SD	Median	Min, Max	
Average across 296 channels (Threshold optimized for each channel)	98.9 \pm 4.3	100.0	48.3, 100.0	0.7 \pm 2.0	0.0	0.0, 18.8	34.6 \pm 29.0	33.3	0.0, 97.5	0.995
Average across 19 patients (Single number of SD used for threshold for each patient)	92.2 \pm 5.1	93.7	76.3, 98.2	5.4 \pm 3.4	4.7	1.5, 12.8	64.0 \pm 10.1	66.0	36.3, 78.5	0.973
All events pooled together Single number of SD used for threshold for all channels)	92.0	-	-	7.0	-	-	66.2	-	-	0.976
Average across 296 channels (Optimized threshold by FDR)	88.6 \pm 14.4	95.5	33.3, 100.0	7.3 \pm 25.6	0.0	0.0, 100.0	16.9 \pm 21.1	9.8	0.0, 97.5	-
Eight-fold Cross-validation	65.8 \pm 4.7	65.0	58.1, 71.7	0.5 \pm 0.4	0.3	0.2, 1.5	33.4 \pm 7.2	30.4	24.0, 44.4	-

CHAPTER III

Automated Detection of High Frequency Oscillations in Human Scalp EEG

3.1 Introduction

HFOs, defined as spontaneous electrographic events in the brain with frequencies ranging from 80-500 Hz, are a promising biomarker of the seizure onset zone in epileptic patients. While this frequency range is not traditionally analyzed in electroencephalography (EEG), there is emerging evidence that these pathological events are visible on scalp recordings (Andrade-Valenca, Dubeau, Mari, Zelman, & Gotman, 2011b; Frauscher et al., 2017; Kobayashi et al., 2010; Pizzo et al., 2016; Joyce Y. Wu et al., 2008; Zelman et al., 2014). Interictal HFOs (70–200 Hz) were observed in scalp EEG associated with continuous spike-waves during slow-wave sleep (Kobayashi et al., 2010), and interictal fast oscillations (< 100 Hz) were specific in identifying children with epilepsy and localized ictal onsets (Joyce Y. Wu et al., 2008). While several automated algorithms have been developed to identify HFOs on scalp EEG recordings (Chu et al., 2017; Nicolás von Ellenrieder et al., 2012), a majority of prior studies relied on time-consuming human visual identification. Thus, developing a reliable automatic HFO detection algorithm for scalp EEG is crucial. Moreover, the ability to noninvasively measure and automatically detect HFOs would have an impact on the clinical diagnosis and treatment of epilepsy.

Automatic HFO detection typically consists of two steps: (1) initial detection of HFO candidate events and (2) rejection of artifacts and false positives. A recently published detector of fast oscillations (40-200 Hz) in scalp EEG recordings employed narrow frequency bands and root mean square amplitude threshold, and classification based on narrowband and wideband

amplitude and absolute narrowband amplitude was applied to reduce false positives (Nicolás von Ellenrieder et al., 2012). A semi-automated method to detect the co-occurrence of spike-ripple events (100-300 Hz) in scalp EEG employed a threshold for amplitude envelope as an initial detection metric and then categorized events using six features, e.g. maximum amplitude, number of oscillations, and Fano factor, together with human visual validation (Chu et al., 2017). Recently, several automated artifact rejection algorithms have been developed to be independently applied after initial HFO detection (Amiri et al., 2016; Burnos et al., 2014; Gliske et al., 2016). Although post-processing can reduce the number of false detections, it also increases the complexity and number of parameters, increasing the difficulty of implementation for new datasets.

As a result, we apply our recently published automatic HFO detection algorithm, which was used to study HFOs in human iEEG recordings (K. Charupanit & Lopour, 2017), to detect HFOs in scalp EEGs. The target frequency bands were ripple, (R, 80-250 Hz) and fast ripple (FR, 250-500 Hz). Furthermore, several false positive rejection algorithms are implemented as post-processing steps.

3.2 Method

3.2.1 Electrophysiological recordings

Scalp EEG data were recorded between 2013-2015 at UCLA Mattel Children's Hospital as a part of the Tuberous Sclerosis Complex Autism Center of Excellence Network (TACERN) database. The recordings were collected under standard 10-20 configurations at 2 kHz sampling rate from seven children with epilepsy and tuberous sclerosis complex. The average subject age was 362 days (range 87-531 days). For each patient, a 12-minute segment of non-REM sleep

EEG was selected. Channels with continuous artifacts and data segments with excessive movement artifacts were excluded, leaving 9-12 minutes of data for each channel. In total, we analyzed 1280 minutes of data from 118 channels. EEG data were re-referenced to a longitudinal bipolar montage to minimize the presence of artifacts.

3.2.2 Automatic HFO detection

In this work, the preliminary HFO detection was accomplished using the automated algorithm from Charupanit et al. 2017 (K. Charupanit & Lopour, 2017), which had previously only been validated for iEEG data. The algorithm required one parameter to be optimized, α . Based on this parameter, the detection threshold was determined iteratively using an estimate of the amplitude probability distribution of the background activity. To apply this algorithm to scalp EEG, the only deviation from the original algorithm was the number of consecutive peaks above the threshold; we used a more relaxed condition of six out of eight consecutive peaks (five out of six was used in the iEEG study). HFO detection was performed separately for two frequency bands, R and FR. For each patient, a single α value was applied to all channels, which ranged from 0.001-0.01 and 0.00001-0.0005 for R and FR, respectively. After the candidate events were identified, several post-processing steps were applied to reduce false detections (see Section 3.2.3). The flow chart of the entire detection process is shown in Figure 3.1.

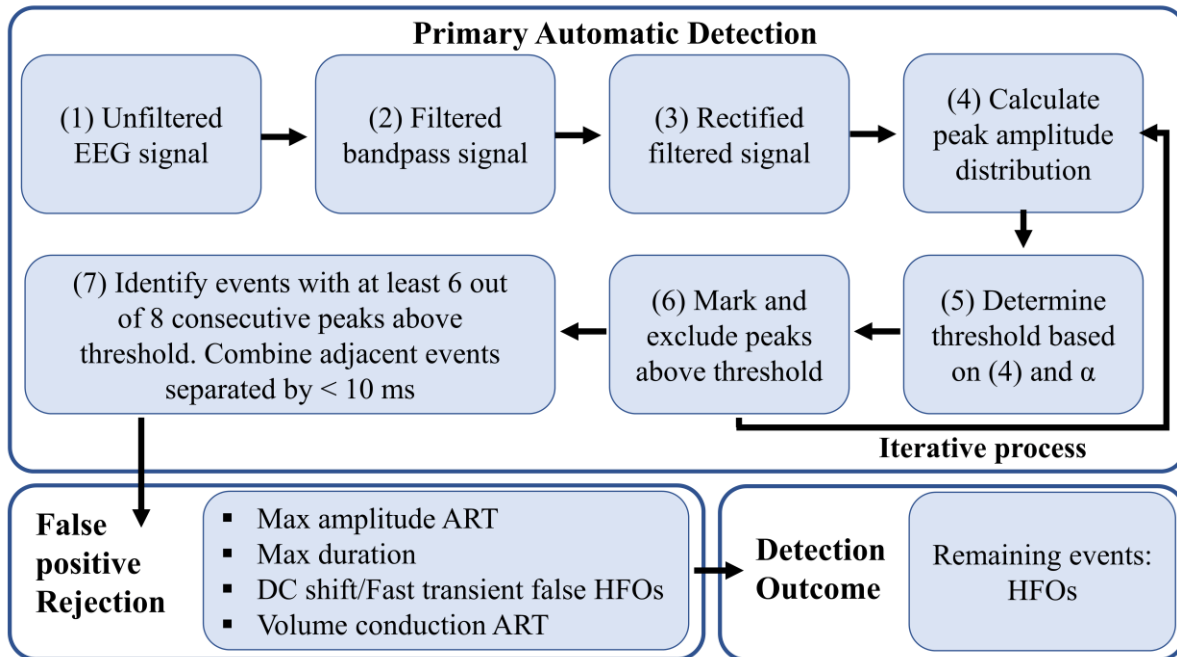


Figure 3.1 Flow chart of the automatic detection algorithm.

3.2.3 Automatic rejection of falsely detected events

We implemented four automatic algorithms to reject false positive detections.

Max Amplitude. Events with extremely high amplitude, greater than expected for signals of neural origin, were regarded as artifact. Any events with a peak amplitude (local maxima of the rectified filtered signal) greater than 15 standard deviations above the mean were classified as artifact. This threshold was calculated in non-overlapping one-min windows.

Max Duration. The typical duration of an HFO is in the range of 40-100 ms (Amiri et al., 2016; Nicolás von Ellenrieder et al., 2012). Longer events tended to be the segments with active background activity, possibly due to muscle noise, which was difficult to distinguish from HFOs. Candidate events longer than 100 ms were excluded.

DC Shift/Fast Transient Generated False HFO. This criterion was adapted from Gliske et al. 2016 (Gliske et al., 2016). It is based on the assumption that electrographic activity at very

high frequencies (above 700-800 Hz) is caused by the filtering of a DC shift or fast transient in the raw signal. Here, an event was marked as artifact when its line-length in the 750-850 Hz bandpass signal was higher than a threshold of mean plus five standard deviations of the average line-length, calculated in 500-ms non-overlapping windows.

Volume Conduction Artifact. The neural generators associated with HFOs are thought to be small (Zelmann et al., 2014), so we expect to see HFOs occurring simultaneously in only one or two channels. However, candidate events were sometimes detected in multiple channels concurrently, due to artifacts such as muscle noise. Thus, events that were detected at the same start and end time, within the range of ± 5 ms, in five channels or more were rejected as artifacts.

3.2.4 Detector validation

The automatic detector was validated by human review. Visual marking of HFOs by experienced reviewers has been considered as a gold standard; however, even using two independent reviewers to visually identify HFOs is subject to bias resulting from poor interrater reliability (Spring et al., 2017). Therefore, we employed three reviewers and used the pre-marked candidate HFOs from the automatic detection to improve the reliability of traditional visual marking by two persons. Visual verification was performed via a custom MATLAB graphical user interface to display 400 ms of data from a single channel, including (1) broadband EEG (1-500 Hz), (2) R bandpass, and (3) FR bandpass signal, and the start and end time of the event were marked [13]. Reviewers KC, BL, and MN were each presented with a random order of detected events consisting of (1) Candidate HFOs: all remaining events after removing artifacts and false positive detections, (2) Rejected artifacts: events removed during the false detection rejection step, and (3) Background: randomly selected segments that did not overlap with HFOs

or artifacts. Artifacts and background events were each approximately 20% of the total number of events; this number varied randomly between datasets. Reviewers were blind to clinical information, number of HFOs, and percentage of false HFOs and background segments to reduce bias in the reviewing process. Reviewers were asked to mark each event as “HFO”, “False HFO/Artifact”, or “Background”. The events with contradicting markings were settled with majority vote.

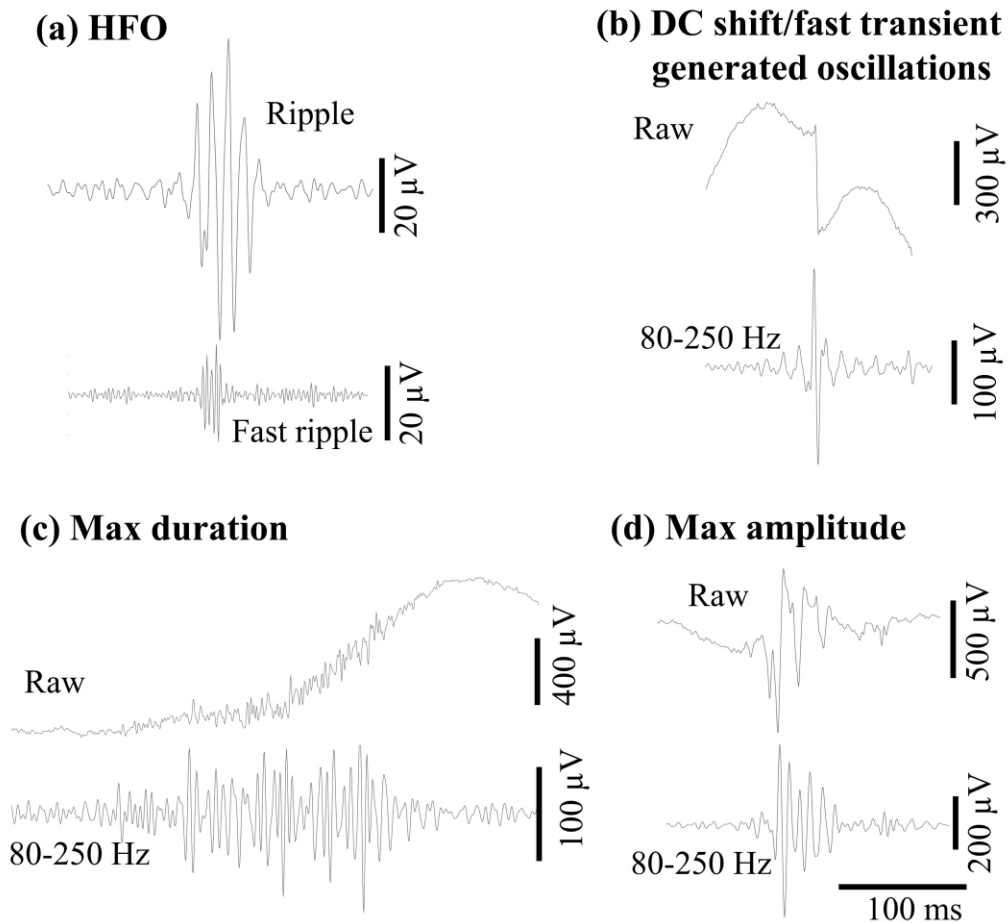


Figure 3.2 Examples of HFOs and events rejected during post-processing

(a) HFOs: R and FR, (b) False HFO due to DC shift/fast transient, (c) Event exceeding max duration, and (d) Event exceeding max amplitude rejection.

3.3 Result

3.3.1 Automatic detected candidate

The detection results were similar for R and FR bands in terms of the ratio between false positive events and total detected candidates. There were 715 and 444 candidates for R and FR, respectively. In the case of R, 231 (32.3%) events were removed in the rejection process leaving 484 (67.7%) candidate HFO events. For FR, 298 out of 444 events (67.1%) passed the artifact rejection steps. The most common detected artifacts were false oscillations due to DC shift/fast transients: 113 (15.8%) and 117 (26.4%) events from R and FR, respectively. Several rejected events met the rejection criteria in more than one category.

3.3.2 Human visual validation

All candidates that passed the rejection steps were visually marked by three reviewers. The accuracy was averaged equally across all seven patients. Based on the visual marking in R band, 43.2% were identified as true HFOs, 1.7% were artifacts, and 55.1% were classified as background activity. In FR, reviewers marked 63.2% of the candidates as HFOs, 5% as artifacts and 31.8% as background activity. All added background segments were correctly identified. However, only about 50% of the artifacts were correctly classified. The result of the validation via visual review is summarized by a confusion matrix (Table 3.1). The percentage of true HFOs varied significantly between patients (Figure 3.3). Two patients had multiple channels with active high-frequency background which obstructed the identification of HFOs (Zelmann et al., 2012). In those subjects, reviewers marked <30% of candidates as HFOs. Overall, the rates of true HFOs were higher in FR than R in all but one patient.

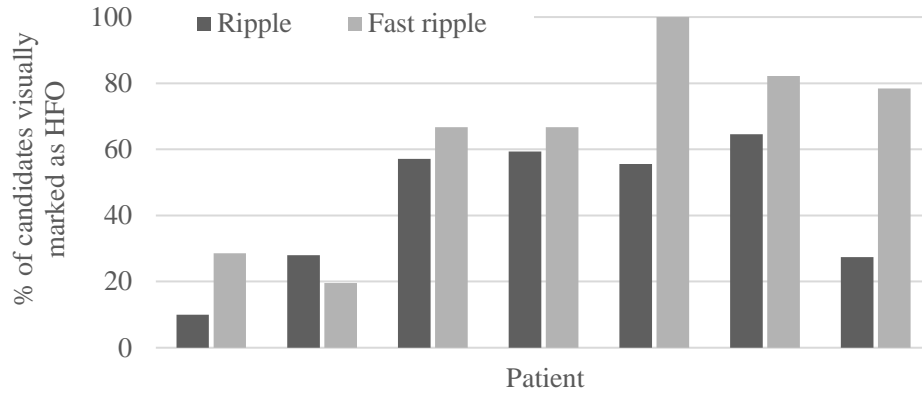


Figure 3.3 The percentage of automatically detected candidates which were visually identified as HFOs.

Table 3.1 Confusion matrix detailing the average percentage accuracy of visual review from seven patients for R and FR (FR is in parenthesis).

		Visual review		
		HFO	Artifact	Background
Automatic Detection	HFO	43.2 (63.2)	1.7 (5.0)	55.1 (31.8)
	Artifact	34.6 (40.7)	54.7 (51.5)	10.7 (7.8)
	Background	0.0 (0.0)	0.0 (0.0)	100.0 (100)

For the visual validation, the agreement between reviewers was 68.7%, 83.0%, and 62.6% in R, and 82.3%, 72.4%, and 76.7% in FR. Surprisingly, the agreement between reviewers and the percentage of HFO candidates confirmed by visual review were both higher in the FR band.

3.3.3 Characteristics of HFOs

We analyzed characteristics of the visually validated HFOs including duration and rate. HFOs were found in 72 (61.0%) and 43 (36.5%) out of 118 channels for R and FR, respectively. The average rates of HFOs in R and FR, based only on the channels where HFOs were found, were 0.24 ± 0.22 and 0.29 ± 0.22 per min, respectively (Figure 3.4). The average duration of R was 47.0 ± 16.3 ms, while FR were half as long at 23.7 ± 12.6 ms. The HFOs were found mostly in the temporal, occipital, and frontal channels located farthest from the vertex.

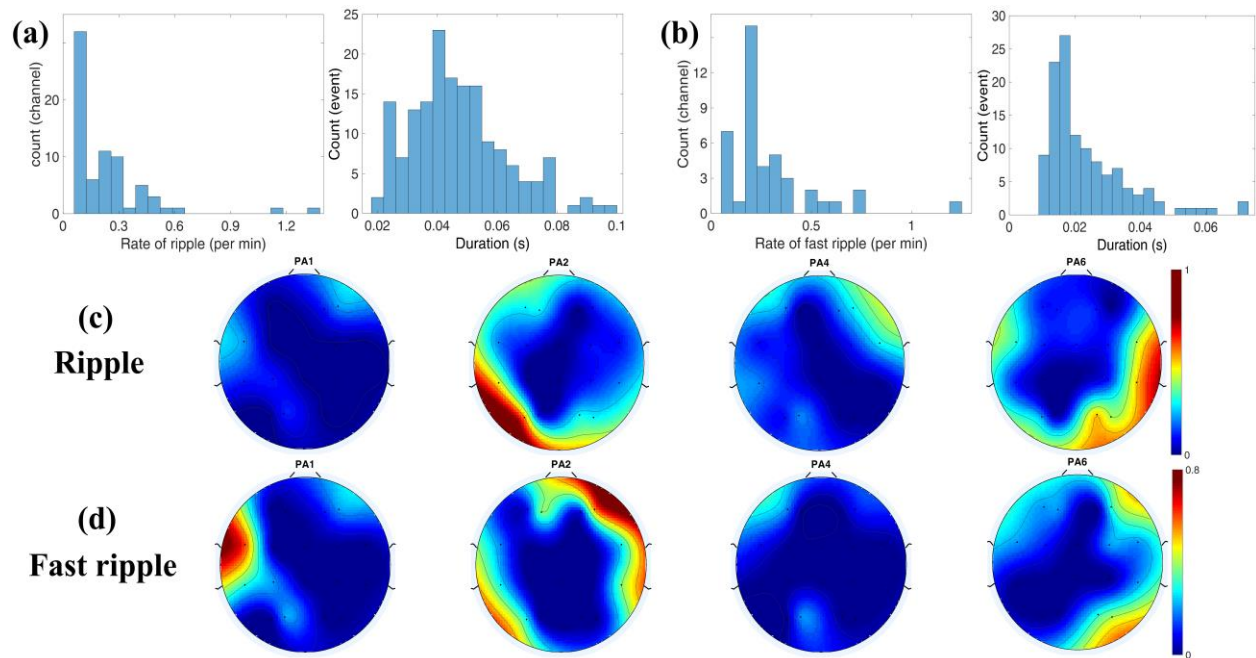


Figure 3.4 Histograms of the HFO rate per minute and duration in (a) ripple and (b) fast ripple.

Examples of topographies of the rate of HFO of two patients with the highest (PA4 and PA5) and lowest (PA1 and PA2) percentage of true HFOs in (c) R and (d) FR.

3.4 Discussion and conclusion

We have shown that our simple automatic detector originally designed for iEEG recordings, can be adapted for the detection of HFOs in scalp EEG recordings. The only parameter in the detection algorithm we have changed is the number of consecutive peaks above the threshold. The range of α is similar to the iEEG study because the threshold depends on the peak amplitude probability distribution of the signal and is self-adapting via iterative process. Note that the detection performance could be improved if the threshold was optimized for each channel, rather than using a single value for each subject (K. Charupanit & Lopour, 2017). To reduce false positives and identify true HFOs, we applied four post-processing rejection steps and human validation. These processes excluded many false positive events in both R and FR.

Approximately 35% of the candidate HFOs were removed by four rejection algorithms; of the remaining events, only about 5% were visually identified as artifacts. This confirmed the effectiveness of the post-processing steps in removing artifacts. On the contrary, 30% of removed artifacts were classified as HFOs in the visual review process. The most commonly rejected candidates that were marked as true HFOs were events from the max amplitude rejection. Since the rejected events resembled HFOs, but with higher amplitude, it caused the reviewers to mistakenly mark them as HFOs. In various cases, it is impossible to conclude whether they are true HFOs or artifact-induced oscillations. Overall, the post-processing steps greatly reduced the false detections, but they may have sacrificed a small portion of HFOs in the process.

Visual identification was used to verify the detected events. A recent study showed a poor degree of inter-rater reliability when HFOs were visually identified by two reviewers (Spring et al., 2017). However, in our study, the agreement between three reviewers was 63-83%, and the percentages were unexpectedly higher in FR than R. In several cases of detected R

candidates, the presence of high amplitude active background caused difficulty in classification of the candidates and resulted in incongruity between reviewers' decisions. Conversely, the background in FR band was less active, so the candidates were more distinct and easier to be classified. Here, approximately half of HFO candidates were visually identified as background activity because of the generous detection threshold, allowing several active background segments to be selected. This was to maximize the sensitivity of the detector and, simultaneously, to reduce bias in the visual review process by introducing several non-HFO events.

The characteristics of the scalp HFOs were different than those seen in iEEG. Scalp HFOs were more difficult to be distinguished and identified; therefore, the rate of HFOs was drastically lower than in iEEG studies (Dümpelmann et al., 2015; Malinowska et al., 2015). However, the duration was similar because the HFO definition used in the detection algorithm was identical. It required a number of oscillations standing out from background, which limited the minimum duration of the detected events. Moreover, the maximum duration artifact rejection also capped the maximum duration. The HFOs were found mostly in temporal and occipital channels. In general, the channels with high appearances of R tended to have a higher rate of FR; however, in some cases, e.g. patients 2 and 6, the regions with high FR rate showed low occurrences of R and vice versa.

The automated detection was reasonably accurate and greatly reduced the time needed to mark HFOs; average time was about 10-15 minutes for each patient for 10 minutes of data and 19 channels. However, this process could be further improved by minimizing the number of non-HFOs to go through visual validation. For example, we could add a background candidate rejection step by setting a minimum amplitude threshold, or we could increase the threshold of

the preliminary HFO detection to reduce the number of candidates. To achieve the potential of HFO as a biomarker for epileptogenic tissue, it is crucial to develop reliable automatic detection for noninvasive recordings to complement invasive HFO studies. This will further our knowledge of these unique events and support clinical translation.

CHAPTER IV

Detection of Anomalous High Frequency Events in Human Intracranial EEG

4.1 Introduction

Twenty to forty percent of patients with epilepsy will not achieve seizure freedom using medication, and this may lead them to consider surgery as a treatment option(French, 2007). The surgical procedure typically relies on localization of the seizure onset zone (SOZ) using intracranial electroencephalography (iEEG) to guide resection of neural tissue. Recent studies have shown that high frequency oscillations (HFOs) occur more frequently in the SOZ(Burnos et al., 2014; Cho et al., 2014; Dumpelmann et al., 2015; Fedele et al., 2016; Julia Jacobs et al., 2008; Liu et al., 2016; Nicolás von Ellenrieder et al., 2012; M. Wu et al., 2018), and the surgical removal of brain regions with high incidences of HFOs has been correlated to a higher likelihood of seizure freedom after surgery(Cho et al., 2014; Fedele et al., 2016; Fujiwara et al., 2012; Guragain et al., 2018; Julia Jacobs et al., 2010; J. Y. Wu et al., 2010). These results suggest that HFOs may be a valuable marker for localization of epileptogenic tissue during surgical planning. Moreover, HFOs occur in interictal periods, so their use may enable clinicians to shorten the duration of invasive monitoring.

HFOs are empirically defined as spontaneous electrographic patterns consisting of at least four cycles of an oscillation in the frequency range of 80–500 Hz, with a high amplitude that is distinguishable from the background(G Buzsáki et al., 1992; Frauscher et al., 2017; J Jacobs et al., 2012). Because these are transient events, detection of HFOs is a critical step in the localization procedure. The gold standard for detection is visual identification(Chaibi et al.,

2013; Frauscher et al., 2017; Nicolás von Ellenrieder et al., 2012; Zelman et al., 2012), but automated detectors are increasingly being implemented to save time and improve reliability and reproducibility(Spring et al., 2017, 2018). Automatic detection algorithms generally consist of the same basic steps: they identify a period of increased high frequency energy (measured with root-mean-square (RMS) amplitude(Blanco et al., 2010; Staba et al., 2002; Zelman et al., 2012), amplitude of filtered, rectified data(K. Charupanit & Lopour, 2017; Krit Charupanit et al., 2018), line length(Dümpelmann et al., 2012; Matsumoto et al., 2013), Hilbert envelope(Burnos et al., 2014; Crépon et al., 2010; Fedele et al., 2016), or as a peak in the time-frequency decomposition(Khadjevand et al., 2017)), then verify that the event exceeds a minimum duration or a minimum number of oscillations. Many detection algorithms include additional steps to merge consecutive events and reject false positives.

Despite the large number of automated algorithms that are currently available, there are two challenges of HFO detection that have not yet been addressed. First, both visual and automated detection rely on the empirical definition of an HFO derived from visual observation(G Buzsáki et al., 1992). There is currently no clear physiological definition that can be used to guide the selection of detection parameters such as amplitude, duration, and number of cycles, as studies have shown significant overlap between pathological and physiological HFO(Alkawadri et al., 2014; Khadjevand et al., 2017; Matsumoto et al., 2013; Pail et al., 2017; Wang et al., 2013). However, the optimization of such parameters is critical to the accuracy of the detector(K. Charupanit & Lopour, 2017; Krit Charupanit et al., 2018; Zelman et al., 2012). This is directly related to the second challenge: existing detection methods require complex optimization procedures. These algorithms typically contain at least 3-5 interrelated parameters, and the detection accuracy is highest when the parameters are optimized for individual

subjects (Krit Charupanit et al., 2018; Maeike Zijlmans et al., 2012). As a result, these algorithms do not easily generalize to new datasets, and the time and effort needed for validation and optimization is a barrier to implementation in a clinical setting.

Therefore, we propose a new algorithm for detection of transient high frequency events in iEEG data that addresses these two challenges. Rather than identifying events with specific features, our anomaly detection algorithm (ADA) detects any events that stand out from the background, regardless of amplitude or morphology. This could include conventional HFOs, oscillations similar to HFOs but with lower amplitude, oscillations with irregular amplitude profiles, artifacts (if they are present in the data), and other unique patterns. ADA integrates several machine-learning techniques, including anomaly detection, time series pattern matching, clustering, and classification. It is fully automated and does not require any human supervision, parameter optimization, or prior assumptions about the shape, amplitude, or duration of the events. Here, we present our algorithm, demonstrate its use on human iEEG data, and compare the detection results to those of a standard HFO detection algorithm. We hypothesize that ADA will enable unbiased estimation of HFO properties, which has the potential to lead to the development of more accurate biomarkers of the epileptogenic zone.

4.2 Methods

4.2.1 Patients and recordings

Intracranial EEG recordings were collected from 36 adult patients between April 2015 and December 2017 at University of California, Irvine, Medical Center. We selected 11 patients (five female, 38.2 ± 16.9 years old) with medically refractory epilepsy who underwent electrode implantation to localize the SOZ for possible surgical resection. The selected recordings had to

fulfill the following inclusion criteria: 1) the SOZ was clearly localized to one or more iEEG channels by experienced neurophysiologists based on seizures that occurred during the monitoring period; 2) electrode locations were confirmed using co-registered pre-implantation and post-implantation structural T1-weighted magnetic resonance imaging scans; 3) minimum recording duration of six hours with no seizures, collected overnight while the patient was resting; and 4) a minimum sampling frequency of 2kHz. The recordings were 20 to 90 hours in duration with a 2 kHz (one patient) or 5 kHz (ten patients) sampling rate and contained data from a total of 1148 electrodes (104.3 ± 33.1 electrodes per patient). Electrodes that could not be clearly localized to the gray matter, electrodes with continuous electrographic artifact, and electrodes within the regions of immediate seizure spread outside of the SOZ were excluded from the analysis. We analyzed the remaining 55 SOZ electrodes, an average of 5.0 ± 3.0 channels per patient, and 122 channels outside the SOZ (which we term nSOZ), an average of 11.1 ± 6.3 channels per patient. We then selected multiple three-minute segments of iEEG for each patient using the following rules. The segments were clipped from overnight iEEG records between 11 PM to 6 AM. Each segment was at least one hour away from a seizure onset time, and we ensured that segments from the same patient were recorded at least 15 minutes apart. This resulted in the selection of two to three segments per hour from each subject. In total, 118 segments were analyzed (mean of 10.7 ± 2.8 segments per patient or 32.1 ± 8.4 minutes per patient, with a range of 7 to 17 segments per patient).

4.2.2 Anomaly Detection Algorithm

The novel algorithm described here aims to separate anomalous high frequency events from the baseline background signal without human supervision or assumptions about the

appearance of the events. The procedure consists of three parts: 1) preprocessing, 2) constructing a distance matrix, and 3) clustering and classification (Figure 4.1). All data analysis procedures were implemented in MATLAB 2018b (MathWorks, MA) using custom-written code.

Preprocessing: The iEEG was re-referenced to a bipolar montage via subtraction of adjacent electrodes, and the resulting signals were modified using the Simple Diff method (Gardner et al., 2007; Roehri et al., 2016). This method flattened the power spectrum in the frequency domain by enhancing the energy of the high frequency activity and suppressing the low frequency components. The flattening process started with application of the fast Fourier transform (FFT) to the data. Then the FFT power spectrum was multiplied by a constant scalar factor of $1 - \cos(2\pi f/fs)$, where f was the frequency and fs was the sampling rate. The final modified iEEG signal in the time domain was obtained by applying an inverse FFT to the modified flattened power spectrum. Each modified iEEG segment was then filtered using an 80 Hz high-pass finite impulse response digital filter. The data were filtered in the forward and reverse directions to avoid phase distortion using the `filtfilt` function in MATLAB.

Constructing the distance matrix: To construct the distance matrix for one segment of data from one channel, we first selected two window sizes: a 1.5 ms small window and 50 ms large sliding window. The small window was used to downsample the iEEG, which was necessary to reduce processing time; within each small window, the representative amplitude of the iEEG was calculated as the average amplitude of all data points. The large 50 ms sliding window was used with 50% overlap for event detection. Each large window consisted of 33 small windows, which was long enough to contain a typical HFO event.

We then measured the distance between the iEEG time series in all pairs of large windows, using the MATLAB function `dtw` to calculate dynamic time warping (DTW)(Chiba, 1978; Paliwal, Agarwal, & Sinha, 1982). DTW measures the similarity between two temporal sequences while being robust to phase differences. Then the distance matrix (DM) was created by assembling these calculated distances into a nonnegative, square, two-dimensional symmetric matrix with elements corresponding to the pairwise distances.

Clustering and classification: The upper triangular elements of each DM were then converted into vector form for the linkage function used to create a hierarchical cluster (HC) tree. The unweighted average distance linkage method was applied as an algorithm for computing the distance between clusters. For clustering, a maximum of seven clusters was used, although comparable results were achieved with a maximum number of clusters ranging from seven to thirteen. The cluster containing the highest number of members (where each member was one 50 ms window of iEEG data from the associated electrode) was designated as the background cluster, and the remaining six clusters were merged together and defined to be the anomaly group. Finally, for each iEEG electrode, any overlapping 50 ms windows within the anomaly group were merged into single events.

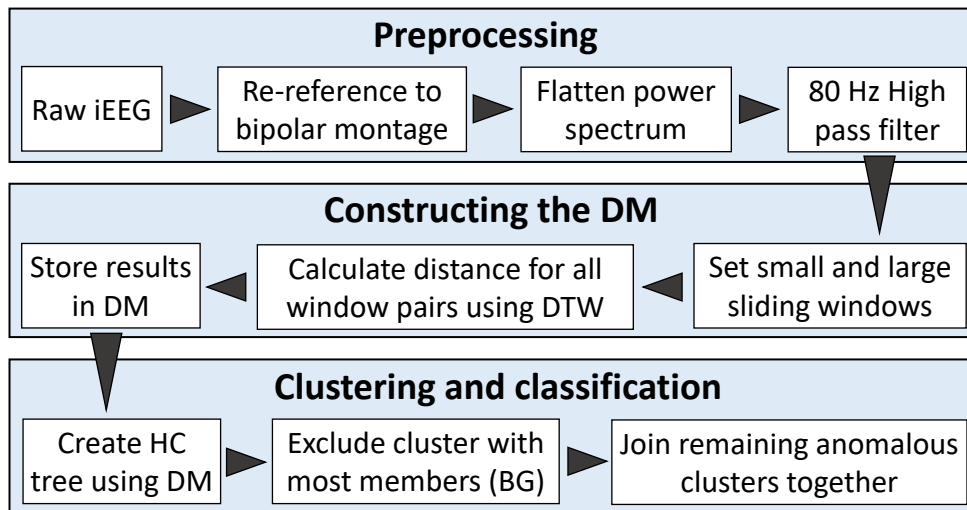


Figure 4.1 Data-flow diagram for ADA.

The algorithm consists of three main parts (large shaded boxes): preprocessing, constructing the DM, and clustering and classification. Each small white box represents a major processing step within each main part, and the arrowheads show the flow of the algorithm. Abbreviations are DM: distance matrix, DTW: dynamic time warping, HC: hierarchical cluster, BG: background cluster.

4.2.3 RMS detector

The RMS detector was originally designed to identify HFOs in hippocampus and entorhinal cortex microwire recordings in humans (Staba et al., 2002), and it was later used to identify HFOs from microelectrode recordings in temporal regions (Staba et al., 2007, 2004). Since then, the RMS detector has become one of the most widely used automated detectors in publications related to HFOs (Cho et al., 2014; Gardner et al., 2007; Gliske et al., 2016; Murphy, Paternos, & Santaniello, 2017; M. Wu et al., 2018). Here, we use the RMS detector to identify traditional HFO events for comparison with the results of the ADA.

The RMS detector is based on the moving average RMS amplitude of the 100 to 500 Hz bandpass filtered signal (finite impulse response filter, roll off -33 dB/octave). The parameters for the RMS detector used in our study matched the original publication (Staba et al., 2002). We calculated the RMS amplitude using a three ms sliding window and the RMS threshold was defined as five standard deviations (SD) above the mean RMS value of the entire length of the signal. Segments of iEEG were considered to be HFO candidates when the RMS amplitude exceeded the threshold for at least six ms. Consecutive candidate events less than ten ms apart were joined together as a single event. Finally, the candidate events were accepted as HFOs if at least six rectified peaks exceeded a second threshold, defined as three SDs above the mean of the rectified filtered signal.

4.2.4 Characteristics of Detected Events

Hereafter, we will refer to events identified by the RMS detector as conventional HFOs (cHFO). Events identified via ADA, which do not have specific thresholds for amplitude or number of oscillations, will be referred to as anomalous high frequency activity (aHFA).

We compared the results of ADA and RMS detection by analyzing the shapes and characteristics of three groups of events: events detected only by ADA, only by RMS, and both by ADA and RMS, which we will refer to as ADA-only, RMS-only, and RMS+ADA, respectively. The shape of the events in each group was compared using the amplitude envelope from the Hilbert transform (median, 95th percentile, and maximum). We also measured four characteristics of the aHFA and cHFO: 1) rate, defined as the average number of events per minute per channel, 2) amplitude, defined as the average value of the upper amplitude envelope over the duration of the event, 3) duration, and 4) coefficient of variation (CV). The CV is a

statistical measure of dispersion of data around the mean, represented by the ratio of the SD to the mean. We used the CV to measure the consistency of the event characteristics across all segments of data. All characteristics of aHFA and cHFO were compared between SOZ and nSOZ channels. We employed a Wilcoxon rank sum test to determine whether the event characteristics were significantly different, and the significance for all analyses was set at $p < 0.05$.

4.3 Results

4.3.1 Incidence and morphology of aHFA and cHFO

Across all iEEG electrodes, a total of 598 SOZ and 1360 nSOZ three-minute epochs from 11 patients were analyzed (Figure 4.2A). Overall, 21294 cHFOs (14008 in SOZ and 7286 in nSOZ) were detected using the RMS detector, and 22311 aHFA (6335 in SOZ and 15976 in nSOZ) were detected using the ADA (Figure 4.2B).

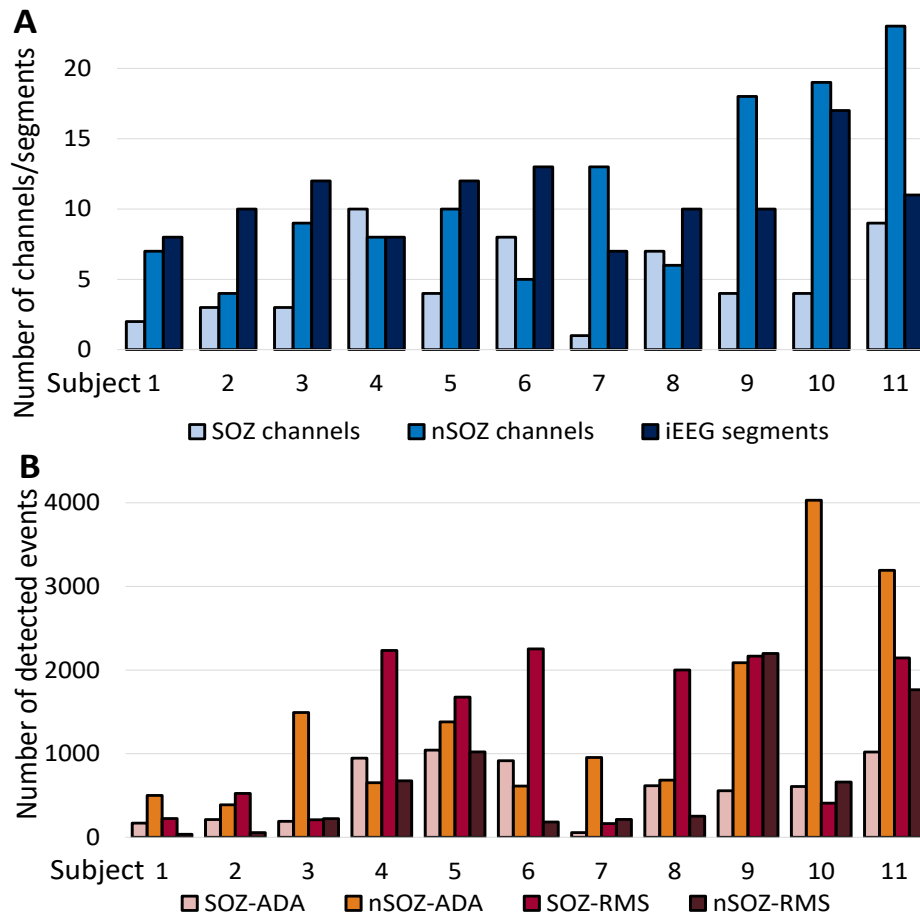


Figure 4.2 Histogram of total numbers of analyzed SOZ, nSOZ channels, and segments together with number of detected events for individual patients.

A) Total numbers of SOZ and nSOZ channels and number of three-minute segments of iEEG data analyzed for each individual patient. B) Total number of detected events from each individual patient divided into SOZ and nSOZ channels.

The number of identified events in RMS+ADA, ADA-only, and RMS-only varied greatly between SOZ and nSOZ (Figure 4.3). In every subject, a subset of events was detected by both the ADA and RMS detectors. In the SOZ, a majority of events were RMS-only (55.5%), while 39.1% of events were RMS+ADA and 5.4% were ADA-only. In five subjects, less than 5% of

events in the SOZ were ADA-only, indicating that the events detected by ADA were a subset of the cHFOs identified by the RMS detector. In contrast, a majority of events in the nSOZ channels were ADA-only (65.9%), while 10.9% were RMS+ADA and 23.2% were RMS-only.

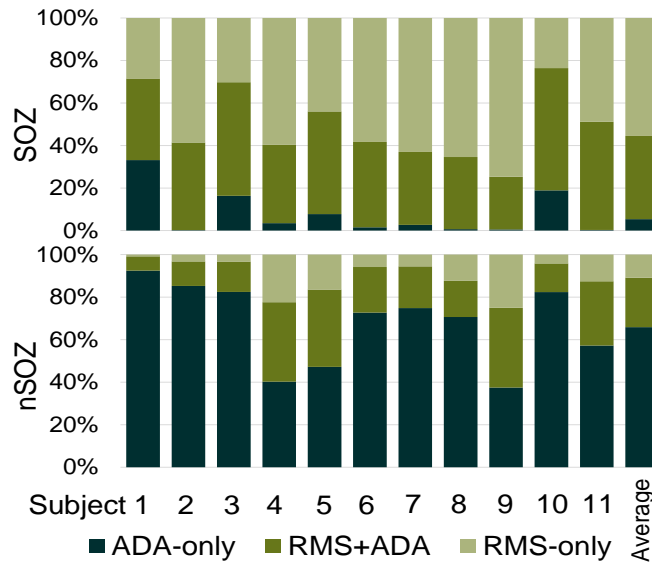


Figure 4.3 Percentages of the detected events in the ADA-only, RMS+ADA, and RMS-only groups for each subject for SOZ channels (top) and nSOZ channels (bottom).

To compare the morphology of events, we plotted the median, 95th percentile, and max amplitude envelopes for each group of events (Figure 4.4). All groups of events exhibited amplitude profiles that reached a peak at the center of the event and tapered off at the edges, consistent with the traditional definition of an HFO. The amplitude of events in the RMS-only group was higher than those in the ADA-only group, suggesting that ADA may have considered a number of high amplitude events to be non-unique. The low-amplitude events in ADA-only were likely missed by the RMS detection scheme due to the application of a strict amplitude detection threshold. This result is consistent with the data in Figure 4.3; because the ADA detects

low-amplitude events, aHFA represent a larger percentage of events in nSOZ channels, where high amplitude events are infrequent.

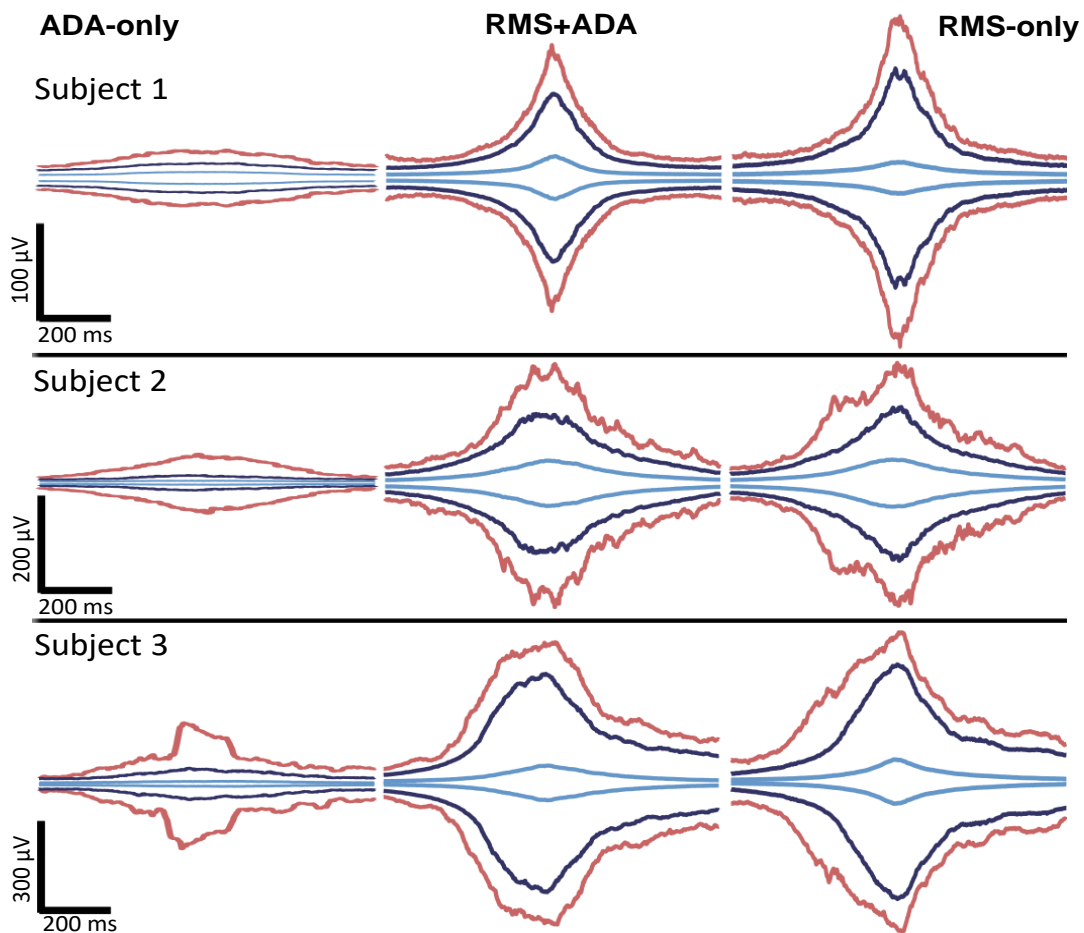


Figure 4.4 Amplitude envelopes of detected events separated into ADA-only (left panel), RMS+ADA (center panel), and RMS-only (right panel) from three representative subjects. The maximum, 95th percentile, and median amplitude envelope are represented by red, blue, and light blue lines, respectively.

4.3.2 Characteristics of aHFA and cHFO

Because the rate (average number of HFOs per minute) has been almost exclusively used as a biomarker of the SOZ, we first measured the aHFA rate and cHFO rate in individual

subjects, as well as the total rates when all events were pooled together (Figure 4.5A). Our results for rate using the RMS detector were consistent with previous studies: the rate of cHFO in the SOZ was 7.8 ± 5.0 per minute, which was significantly higher than the rate in the nSOZ, 1.8 ± 2.9 per minute. The individual results for all 11 patients also exhibited significant differences in cHFO rate between SOZ and nSOZ. The results for rate using ADA were less consistent. The average rate of aHFA was significantly different between the SOZ (3.5 ± 2.9 per minute) and nSOZ (3.9 ± 2.2 per minute), with the nSOZ having a higher rate. However, only six out of eleven patients had significantly different rates of aHFA in SOZ and nSOZ. In five out of the six patients, the nSOZ had a higher rate. Therefore, the rate of aHFA does not provide reliable separation between SOZ and nSOZ channels.

In contrast to the rate, the event amplitude showed robust differences between SOZ and nSOZ channels using both ADA and the RMS detector (Figure 4.5B). The mean amplitudes of cHFO ($38.0 \pm 33.7 \mu\text{V}$) and aHFA ($40.1 \pm 32.5 \mu\text{V}$) in SOZ were significantly higher than in nSOZ (8.8 ± 9.1 and $7.1 \pm 9.1 \mu\text{V}$, respectively). These differences were statistically significant for all 11 individual subjects using both detection methods. Moreover, the amplitudes were consistent across subjects, such that a single, common threshold of approximately $15 \mu\text{V}$ can approximately separate SOZ and nSOZ across all subjects. This is not true for the rate, which requires selection of a patient-specific threshold to separate SOZ and nSOZ channels. The robustness and consistency of these results suggests that the amplitude of aHFA and cHFO may be potential biomarkers of the SOZ.

The average duration of detected events within SOZ channels was significantly longer than the duration in nSOZ channels for both detection schemes (Figure 4.5C). This difference was statistically significant for 8 subjects using ADA and 9 subjects using the RMS detector,

with events in the SOZ generally having a longer duration. However, for individual subjects, the duration of cHFOs provided smaller separation between SOZ and nSOZ channels than event amplitude or rate.

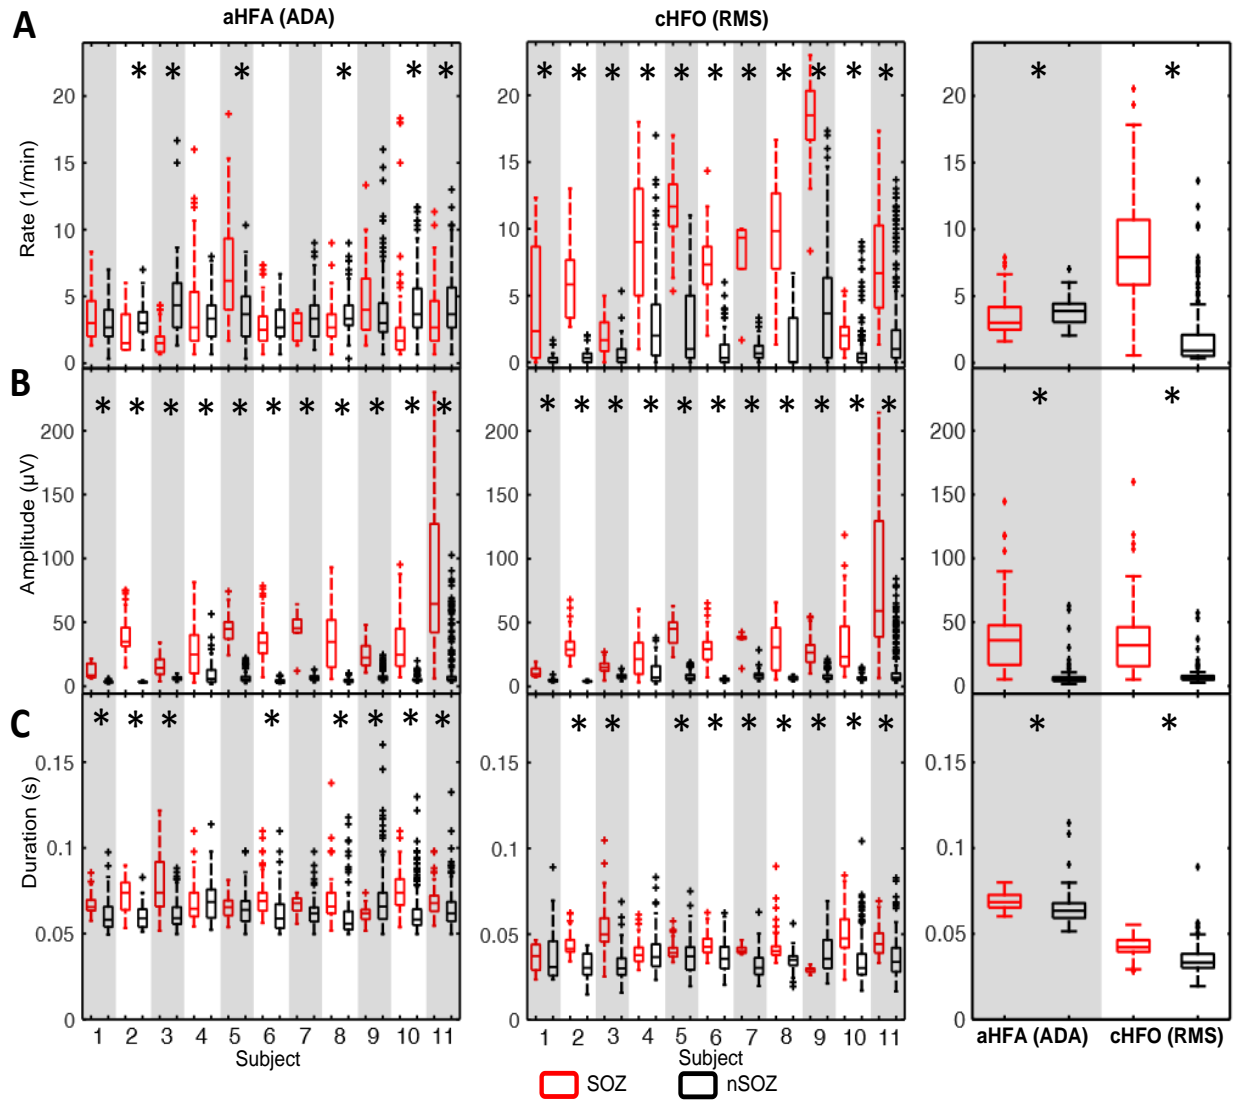


Figure 4.5 Characteristics of detected events separated into SOZ (red) and nSOZ channels (black).

Boxplots show the (A) rate, (B) amplitude, and (C) duration of detected events for individual subjects (aHFA in left column and cHFO in middle column) and overall characteristics of

detected events when all segments are pooled together (right column). * p-value < 0.05, Wilcoxon rank sum test.

Finally, we calculated the CV to assess the consistency of the measurements across segments of data (Figure 4.6). A low CV value indicates a consistent measurement, where the SD is lower than the mean. The amplitude CV was significantly lower than the rate CV when calculated using the same detector and channels, suggesting that the estimates of amplitude are more stable over time for both detection schemes. The CV for amplitude in the SOZ was higher than in nSOZ, but it was still generally lower than the CV values for rate.

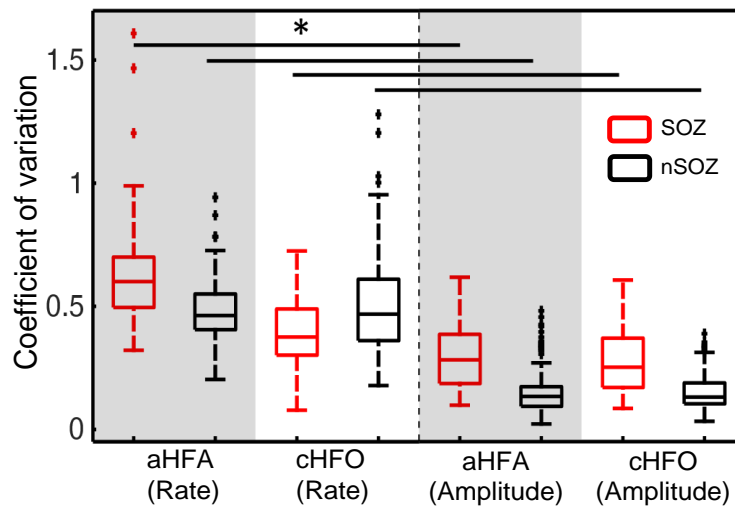


Figure 4.6 Boxplots of CV for the rate and amplitude of detected events, separated into SOZ (red) and nSOZ (black).

* p-value < 0.05, Wilcoxon rank sum test.

4.4 Discussion

Here, we have presented a novel algorithm for detection of anomalous high frequency events in electrophysiological data and applied it to human iEEG. The ADA is unsupervised and does not require complex optimization procedures or prior assumptions about the shape or

amplitude of the events. When we compared the events detected by ADA to those of a standard energy-based detector, we found significant overlap, indicating that ADA is sensitive to traditionally-defined HFOs. However, we also found that ADA identifies additional events that would not exceed the energy threshold of standard algorithms. For example, many aHFA events had a low amplitude, particularly in the nSOZ. The rate of aHFA was not consistently different between the SOZ and nSOZ; however, the amplitude of aHFA provided reliable and robust separation between SOZ and nSOZ channels, suggesting it as a possible biomarker of the SOZ.

Robust identification of clinically-relevant HFOs has remained a challenge due to the lack of a physiological definition of these events and the limitations of current detection methods. Because the mechanism underlying HFOs is not fully understood, detection is typically guided by the empirical definition which requires selection of an optimum energy threshold to separate events from the background. This is not a trivial task because the shape and amplitude of the HFO waveform can vary depending on the distance between the electrode and the neural generator (György Buzsáki et al., 2012) and the characteristics of the background activity vary over time. Therefore, a rigid template of shape and amplitude is likely to be insufficient for HFO detection. In addition, current automated algorithms require complex optimization procedures to maximize accuracy (Zelmann et al., 2012), and it is common for them to suffer from a high number of falsely detected events (Chaibi et al., 2013; K. Charupanit & Lopour, 2017; Zelmann et al., 2012). Visual identification by expert reviewers has been widely used for detecting HFOs in both scalp and iEEG recordings (Ferrari-Marinho et al., 2015; Julia Jacobs et al., 2014), as human reviewers can simultaneously adapt to changes in the background activity and perform artifact rejection. However, it is highly time-consuming and has poor agreement between reviewers (Spring et al., 2017), which reduces the generalizability of the results. ADA addresses

these challenges, as it does not require human input and can therefore perform unbiased detection and estimation of event characteristics.

In the present study, the default parameters for ADA were configured to detect anomalous events at frequencies greater than 80Hz. The small window was used to downsample the data to reduce the calculation time(Keogh, Lin, & Fu, 2005); the approximate detection time for three minutes of data from a single channel of iEEG was ~4-5 minutes using a desktop PC. With parallel computing or optimization of the algorithm, it may be possible for the size of this window to be reduced or to exclude this step from the procedure. The size of the 1.5 ms small window is equivalent to downsampling of the signal to 666Hz, which means that our algorithm is primarily detecting events in the ripple frequency band (80-250 Hz). However, the algorithm could be configured for other frequency bands as well, simply by changing the sizes of the small and large windows. The large sliding window was chosen to match the approximate duration of HFOs reported in prior studies(Bagshaw, Jacobs, Levan, Dubeau, & Gotman, 2009; Matsumoto et al., 2013; Staba et al., 2002). Because the algorithm contains a step to join overlapping large windows in the anomaly group into single events, a range of large window sizes can be used without affecting the results. In the clustering and classification process, the maximum number of clusters was seven; however, this number can be altered, if it is higher than the expected number of anomalous patterns. We tested the algorithm with a range of seven to thirteen clusters, and there was no noticeable difference in the results in this range because the background cluster is always several orders of magnitude larger than any other cluster. Therefore, we chose to use seven clusters, as it reduced the processing time.

We found that aHFA had a longer duration than cHFO, but there were significant differences between the two detection algorithms that contributed to this result. The aHFA had a

minimum duration of 50 ms, corresponding to the size of the large sliding window, while the cHFO duration was measured as the length of time that the RMS amplitude exceeded the energy threshold. This impacted the measurement of amplitude, as well. The same event detected with ADA may have a lower amplitude than when it is detected with the RMS detector because the average amplitude will include some background activity at the edges of the window.

Conventionally, the rate of cHFO has been used as a biomarker of the SOZ in studies of high frequency activity related to epilepsy (Burnos et al., 2014; Cho et al., 2014; Dumpelmann et al., 2015; Fedele et al., 2016; Julia Jacobs et al., 2008; Liu et al., 2016; Nicolás von Ellenrieder et al., 2012; M. Wu et al., 2018). Here, we found that the rate of aHFA was less predictive of the SOZ than the rate of cHFO detected using a strict energy threshold, due to the detection of many low amplitude events by ADA which were not detected by the RMS detector. Every additional detected event raises the rate by one event per minute, similar to the effect of false positive detections when using conventional automated algorithms. This can drastically change the relative rates in SOZ and nSOZ channels (especially for ADA, which more frequently detects low amplitude events in nSOZ), and therefore alter the prediction of SOZ location. This is supported by the fact that the CV of the rate was higher than for amplitude or duration, indicating a higher degree of variability across segments of iEEG.

In contrast to the rate, the aHFA and cHFO amplitudes were significantly higher in the SOZ compared to the nSOZ, with relatively consistent values across patients. Several other studies have suggested that HFO amplitude differs between SOZ and nSOZ channels. Amplitude was shown to be significantly higher in SOZ compared to non-SOZ electrodes, regardless of contact location, during interictal, preictal, and ictal periods (Guragain et al., 2018), but others showed that this difference was not significant during nonictal periods (Malinowska et al., 2015).

Furthermore, pathological HFOs were found to have higher amplitude than physiological-induced HFOs(Matsumoto et al., 2013). However, because the reported differences were often small, amplitude has not received much attention as a possible biomarker of the SOZ. Here, we found that the amplitude of cHFO and aHFA and the rate of cHFO provided robust separation between SOZ and nSOZ in all 11 patients. However, the amplitude exhibited less variability over time and more consistency across patients. This suggests that amplitude may be another promising candidate for a biomarker of the SOZ. Further validation with a larger cohort of patients, more comprehensive inclusion of iEEG electrodes, and comparison to surgical outcome are needed to explore this hypothesis and will be the subject of future investigations.

CHAPTER V

Amplitude of High Frequency Oscillations as a Biomarker of the Seizure Onset Zone

5.1 Introduction

High frequency oscillations (HFOs), defined as distinguishable spontaneous electrographic patterns consisting of at least four cycles in the frequency range of 80–500 Hz, are a promising biomarker for the seizure onset zone (SOZ). (Burnos et al., 2014; Cho et al., 2014; Dumpelmann et al., 2015; Fedele et al., 2016; Julia Jacobs et al., 2008; Liu et al., 2016; Nicolás von Ellenrieder et al., 2012; M. Wu et al., 2018) HFOs are commonly categorized as ripples, ranging from 80 to 250 Hz, (G Buzsáki et al., 1992; Garcia Barba, 2014) and fast ripples, ranging from 250 to 500 Hz. (Pizzo et al., 2016; Urrestarazu et al., 2007) Both ripples and fast ripples occur more frequently in the SOZ than outside the SOZ in human intracranial EEG (iEEG) recordings. (Burnos et al., 2014; Cho et al., 2014; Dumpelmann et al., 2015; Fedele et al., 2016; Julia Jacobs et al., 2008; Liu et al., 2016; Nicolás von Ellenrieder et al., 2012; M. Wu et al., 2018) Moreover, retrospective studies have reported that removal of brain regions with high occurrence of interictal HFOs is correlated to a positive outcome after epilepsy surgery. (Cho et al., 2014; Fedele et al., 2016; Fujiwara et al., 2012; Guragain et al., 2018; Haegelen et al., 2013; Julia Jacobs et al., 2010; J. Y. Wu et al., 2010) HFOs may therefore be a valuable tool for localization of epileptogenic tissue to aid in surgical planning. In addition, because HFO analysis is typically done on interictal iEEG data, this marker represents an opportunity to reduce the duration of the pre-surgical monitoring period. (J Jacobs et al., 2012; Jobst & Cascino, 2015)

The HFO rate (number of HFOs per minute per channel) has been almost exclusively used as the metric to test HFOs as a biomarker of the SOZ.(Alkawadri et al., 2014; Cho et al., 2014; Dumpelmann et al., 2015; Fedele et al., 2016; Iimura et al., 2017; Julia Jacobs et al., 2008; Liu et al., 2016; Malinowska et al., 2015; Malkki, 2014; Pail et al., 2017; Nicolás von Ellenrieder et al., 2012; Wang et al., 2013; M. Wu et al., 2018) However, there are several limitations to this method. The HFO rate varies depending on spatial localization which is difficult to find the standard rate for all brain regions.(Crépon et al., 2010; Dümpelmann et al., 2012; Dumpelmann et al., 2015; Ferrari-Marinho et al., 2015; Frauscher et al., 2018; Khosravani et al., 2009) The rate is also time-dependent, as it increases during sleep(Bagshaw et al., 2009; Dumpelmann et al., 2015) and has different mean values in interictal, ictal, and preictal states.(Khosravani et al., 2009; Malinowska et al., 2015; Pearce et al., 2013) This can lead to variation in the estimated location of the SOZ over time.(Gliske et al., 2018) Moreover, while it is understood that a high rate of HFOs may indicate epileptogenic tissue, no standard threshold for this value has been established. The reported rates of HFOs in the SOZ vary over several order of magnitude; for example, one subject may exhibit an HFO rate of 1-2 per minute,(Cho et al., 2014; Murphy et al., 2017) while another exceeds one hundred HFOs per minute per channel.(Julia Jacobs et al., 2008) This measurement is also confounded by the occurrence of physiological HFOs, which are associated with healthy cognitive processing.(Lachaux et al., 2012) The characteristics of physiological and pathological HFOs overlap,(Matsumoto et al., 2013) making it a challenge to detect solely pathological events, and the rates of physiological HFOs vary across the different regions of the brain.(Frauscher et al., 2018) Lastly, the estimate of HFO rate will be a function of the energy threshold used for detection (whether visual or automated), and every false positive detection will raise this value by one. For example, in the comparison of four well known

automatic detectors,(Zelmann et al., 2012) including RMS, LineLength, Hilbert, and Montreal Neurological Institute and Hospital (MNI) detector, the false detection rate varied between 16.7% (RMS detector) and 75.7% (MNI detector).

In addition to rate, differences in other HFO features, such as amplitude, peak frequency, and duration have been reported.(Alkawadri et al., 2014; Matsumoto et al., 2013; Pail et al., 2017; van Klink et al., 2016; Wang et al., 2013) In particular, the HFO amplitude was significantly higher in SOZ than nSOZ channels for both ripples(Alkawadri et al., 2014; Pail et al., 2017) and fast ripples,(Pail et al., 2017) although one study showed no significant for fast ripples.(Alkawadri et al., 2014) In addition, the pathological HFOs showed higher mean spectral amplitude than induced physiological HFOs.(Matsumoto et al., 2013; Nagasawa et al., 2012) This relationship appears to be relatively consistent over time, as the mean amplitude of HFOs in the SOZ was shown to be significantly higher than in non-SOZ (nSOZ) electrodes in interictal, preictal, and ictal periods.(Malinowska et al., 2015) While these studies suggest that HFO amplitude may have value as a biomarker of the SOZ, the reported differences between SOZ and nSOZ electrodes were typically small due to the different in comparison condition, and the detection scheme. For example, the choice of threshold will affect the estimate of event amplitude. As the threshold increases, so will the estimated mean amplitude for both SOZ and nSOZ channels.

Here, we tested this hypothesis by measuring the amplitude of HFOs in human iEEG using two automatic detection algorithms. First, we measured the amplitude and rate of anomalous high frequency events identified using an anomaly detection algorithm (ADA) in Chapter IV. ADA integrates several machine learning techniques to identify events that stand out from the background; most importantly, it does this without requiring prior assumptions about

event shape or amplitude. Second, we repeated the analysis using a conventional HFO detector based on root-mean-square (RMS) amplitude. For both detectors, we compared the rate and amplitude of detected events in SOZ and nSOZ electrodes and measured the classification accuracy based on these two features.

5.2 Methods

5.2.1 Patients and recordings

The iEEG recordings used in this study were collected from subjects with medically refractory epilepsy undergoing surgical evaluation at University of California, Irvine, Medical Center between April 2015 and December 2017. The selected subjects fulfilled the following inclusion criteria: (1) intracranial monitoring resulted in the localization of the SOZ by epileptologists to one or more electrodes; (2) electrode locations were confirmed by co-registering pre- and post-implantation structural T1-weighted magnetic resonance imaging (MRI) scans; (3) availability of at least one full night of iEEG data (minimum six hours) containing no seizures; and (4) a minimum sampling frequency of 2 kHz. From 36 subjects who underwent intracranial monitoring to localize the SOZ for possible surgical resection, 11 adult subjects fit these criteria (five female, 38.2 ± 16.9 years old). The data for 10 subjects had a 5 kHz sampling rate, and one subject had a 2 kHz sampling rate. The seizure onset times and SOZ were determined independently by experienced epileptologists. Fifty-five electrodes within the SOZ and 122 nSOZ electrodes were selected for analysis. The nSOZ channels were those that were clearly located within gray matter; channels located in white matter or containing continuous electrographic artifact were excluded. Multiple three-minute segments of iEEG data were selected from each patient using the following procedure: (1) segments were clipped from

overnight iEEG recordings between 11 PM and 6 AM; (2) segments were at least one hour away from a seizure onset time; and (3) segments were at least 15 minutes apart, resulting in the selection of two to three segments per hour. The number of segments for each subject ranged from 7 to 17. In total, 118 segments of iEEG (10.7 ± 2.8 segments per subject or 32.1 ± 8.4 minutes per subject) were analyzed in this study.

5.2.2 Anomaly detection algorithm (ADA)

The ADA identifies high frequency events that stand out from the background signal without requiring assumptions about the appearance of those events. The algorithm has been described in detail Chapter IV, so here we provide a brief summary of the procedure. The algorithm used in this study was identical to ADA in Chapter IV except we added one extra condition to the classification step, which enabled us to vary the number of events detected by the algorithm in order to evaluate the robustness of the metrics in Section 5.3.4. All data analysis was done in MATLAB 2018b (MathWorks, MA) using custom-written code. The ADA consists of three steps:

(1) Preprocessing: We re-referenced the iEEG to a bipolar montage and then flattened the power spectrum using the Simple Diff method.(Gardner et al., 2007; Roehri et al., 2016) We filtered the resulting signals using an 80 Hz high-pass finite impulse response digital filter and a zero-phase digital filtering procedure (function `filtfilt` in MATLAB).

(2) Constructing the distance matrix: A small 1.5 ms non-overlapping window was used to downsample the preprocessed iEEG signals; the value for each small window was the average voltage of all data points within the window. This reduced the computation time of the

algorithm. For event detection, we defined a large sliding window of 50 ms with 25 ms overlap, based on the duration of a typical HFO event.

The distance, or dissimilarity, between all pairs of large windows in one channel of the iEEG recording was calculated using dynamic time warping (DTW) via the MATLAB command `dtw`.(Chiba, 1978; Paliwal et al., 1982) DTW measures the similarity between two segments of time series data, and unlike Euclidian distance, it is robust to phase differences between the two signals. These pairwise distances were stored in a two-dimensional symmetric matrix, which we will refer to as the distance matrix (DM).

(3) Clustering and classification: A linkage function was used to create a hierarchical cluster tree via unweighted average distance linkage method. The maximum number of clusters was set to seven, but the results were robust to alteration of this value (see Figure 5.7).

To identify the anomalous events, we started by selecting the main background cluster (BGC), which is defined as the single cluster with the largest number of members. Each member is one 50 ms window of data. In the original algorithm, all clusters except the BGC were defined as anomalous (Chapter IV). Here, the anomalous clusters were selected by comparing the inter-cluster distances (ICD) to a separation threshold calculated from the pairwise distances of all members within the BGC:

$$\text{Separation threshold} = \bar{d} + N \sqrt{\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n (d_{ij} - \bar{d})^2}$$

where d_{ij} is the pairwise distance between windows i and j in the BGC; \bar{d} is the mean of all pairwise distances d_{ij} ; n is the number of members of BGC; and N is the number of standard deviations (SD). The $ICD(k)$ is the mean distance between the k th cluster, C_k , and the BGC:

$$ICD(k) = \sum_{i \in C_k} \sum_{j \in BGC} d_{ij} / (m \times n),$$

where d_{ij} is the distance between the i th member of C_k and the j th member of the BGC, and m and n are the number of members in C_k and the BGC, respectively. Any cluster with an ICD exceeding the separation threshold was marked as an anomalous cluster. Finally, all anomalous clusters were grouped together and any large windows within one channel that overlapped in time were merged into a single event. The default value of N was set to one, which the detected events were the exact same population as the original ADA, and we varied this value to test the robustness of the results.

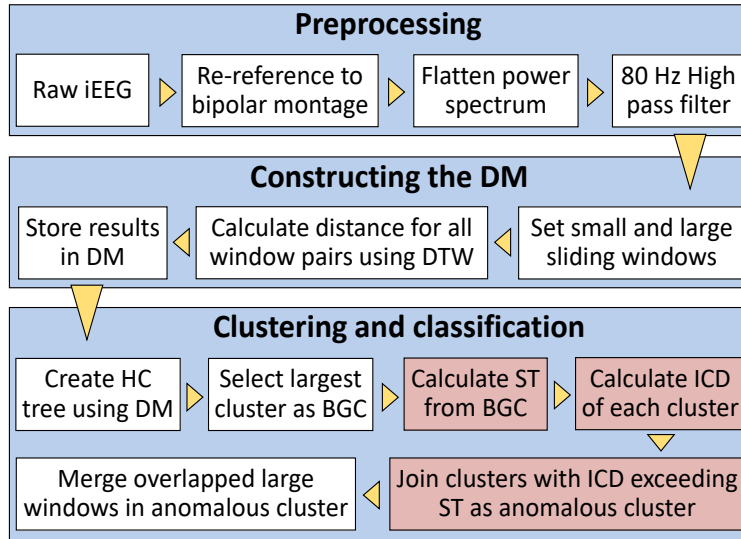


Figure 5.1 Flow chart for modified ADA.

The algorithm consists of preprocessing, constructing the DM, and clustering and classification (blue boxes). The small white boxes represent the major processing steps for ADA, and the small red boxes show the added steps that were not in the original algorithm. The arrowheads show the flow of the algorithm. Abbreviations are DM: distance matrix, DTW: dynamic time warping, HC: hierarchical cluster, BGC: background cluster, ST: separation threshold, ICD: inter-cluster distance.

5.2.3 RMS detector

The RMS detector was originally used to identify HFOs in microelectrode recordings in the human temporal lobe.(Cimbálník, Hewitt, Worrell, & Stead, 2018; Staba et al., 2007, 2004) Since then, it has been widely adopted as an automated HFO detection algorithm and has been applied to a variety of intracranial electrode types across all regions of the brain.(Cho et al., 2014; Gardner et al., 2007; Gliske et al., 2016; Murphy et al., 2017; M. Wu et al., 2018) Here, the RMS detector was used to obtain the benchmark accuracy for classification of SOZ and nSOZ channels based on HFO rate, as a basis of comparison for HFO amplitude. Full details of the automated detection algorithm can be found in Staba et al. 2002,(Staba et al., 2002) here, we implemented the algorithm with all parameters as defined in the original publication.

5.2.4 Characteristics of detected events

In Chapter IV, we previously showed that the ADA and RMS detector identify different subsets of events, with some degree of overlap. More specifically, although ADA does not place specific thresholds on the amplitude or the number of oscillations, approximately 40% of all events in SOZ and 11% of events in nSOZ were detected by both ADA and the RMS detector. The other ADA events were low to intermediate amplitude oscillations and other unique patterns, detected more frequently in nSOZ than SOZ electrodes. Hereafter, we will refer to events detected by ADA as anomalous high frequency activity (aHFA), and events identified by the RMS detector will be called conventional HFOs (cHFO).

The characteristics of aHFA and cHFO analyzed in this study consisted of (1) the event rate, defined as the average number of detected events per minute per channel, (2) the amplitude,

defined as the average value of the upper amplitude envelope, obtained from the Hilbert transform, and (3) the coefficient of variation (CV), a statistical measure of dispersion of data around the mean, defined as the SD divided by the mean. The rate and amplitude of both aHFA and cHFO were evaluated for their accuracy as a biomarker for localization of the SOZ. The CV was used to quantify the stability of the rate and amplitude measurements across recording segments, channels, and subjects.

5.2.5 Evaluation of biomarker performance

The ability to differentiate SOZ and nSOZ was evaluated via the receiver operating characteristic (ROC) curve. Positives (P) were defined as the electrodes within the SOZ, as identified by epileptologists based on seizures captured via intracranial monitoring, and negatives (N) consisted of all brain regions outside the SOZ. The ROC curve was created by plotting the sensitivity, defined as $\text{TrueP}/(\text{TrueP}+\text{FalseN})$, against false positive rate (FPR), defined as $\text{FalseP}/(\text{FalseP}+\text{TrueN})$, for a range of amplitude and rate thresholds. This was done using two different methods: (1) A segment-based ROC curve was created by pooling all segments together and using the threshold to classify individual segments, and (2) a channel-based ROC curve was created by applying the threshold to the average value for each channel, where the average was calculated across all iEEG segments in that channel. Then the performance of the metric was assessed using the area under the ROC curve (AUC) and the sensitivity and FPR at the optimal cut-point. The AUC ranges from zero to one, with one indicating perfect performance. The optimal cut-point, the point with the best balance between sensitivity and FPR, was defined as the point that minimized the Euclidean distance between the ROC curve and the (0% FPR, 100% sensitivity) point at the upper left corner of the ROC plot.

Note that, in the case where an HFO cannot be detected (zero HFO rate) either in SOZ or nSOZ, the amplitude in that segment will be represented as zero amplitude.

We employed the Wilcoxon signed-rank test and Wilcoxon rank sum test to determine statistical significance for paired and unpaired samples, respectively, and the significance for all analyses was set at $p < 0.05$.

5.3 Results

5.3.1 Incidence of detected high frequency events

Note that the default setting, the separation threshold was set to $N=1$, modified ADA detected the exact same events as the original ADA. In total, 21294 cHFOs were detected by the RMS detector and 22311 aHFA were detected by ADA across all iEEG electrodes. These events were contained within 1,958 three-minute segments of data (598 from SOZ and 1,360 from nSOZ). The population of aHFA consisted of 6335 events in SOZ and 15976 events in nSOZ channels. For the RMS detector, 14,008 cHFOs were in SOZ channels and 7,286 were in nSOZ channels.

5.3.2 The amplitude of high frequency events is higher in SOZ and is stable over time

Consistent with prior studies, the rate of cHFO within SOZ channels (7.8 ± 5.0 per minute) was significantly higher than in nSOZ channels (1.8 ± 2.9 per minute) (Figure 5.2A). The aHFA rate was not significantly different between SOZ and nSOZ channels. Surprisingly, the amplitude of detected events was significantly higher in SOZ than nSOZ channels for both ADA and RMS detection (Figure 5.2B). Specifically, the amplitude of aHFA was $40.1 \pm 32.5 \mu\text{V}$ in the SOZ and

7.2±9.1 μV in the nSOZ, and the cHFO amplitude was 37.7±33.7 μV in the SOZ and 6.2±8.6 μV in the nSOZ.

The values of these metrics fluctuated over time, as previously reported by others. (Gliske et al., 2018) However, in evaluating these fluctuations, we found that the amplitude measurement was more consistent across segments of data (i.e. more stable over time), evidenced by significantly lower CV values for the amplitude of aHFA and cHFO compared to cHFO rate (Figure 5.2C).




Figure 5.2 The cHFO rate, cHFO amplitude, and aHFA rate are higher in the SOZ, but the measurements of amplitude are more stable over time.

(A) Rate and (B) amplitude of detected events from the ADA and RMS detectors. Each datapoint represents the mean value from one segment of iEEG. (C) CVs for rate and amplitude, calculated across all iEEG segments for each channel. All features are shown separately for SOZ (red) and nSOZ channels (black). * p -value < 0.05, Wilcoxon rank sum test.

5.3.3 Amplitude enables more accurate classification of SOZ and nSOZ than rate for individual subjects

The individual subject results also reflect these differences in rate and amplitude; this is crucial, as HFO analysis will typically be done on single subjects if used clinically for SOZ

localization. The heatmaps in Figure 5.3 depict the differences between SOZ and nSOZ channels for all iEEG segments for each individual subject. Visually, both the amplitude and rate of detected events were higher in the SOZ than the nSOZ channels, as expected. However, the amplitude measurements typically exhibited a greater difference between SOZ and nSOZ channels than were seen for the standard metric of cHFO rate. In particular, the amplitude values in nSOZ channels were lower relative to the SOZ values and were more consistent over time.

To assess the performance of all three metrics as potential SOZ biomarkers, we calculated the segment-based ROC curves for all individual subjects. The ROC curves for amplitude were generally closer to the upper left-hand corner, suggesting that amplitude provided better separation between SOZ and nSOZ channels than cHFO rate (Figure 5.4A-C). To statistically compare the performance of the three metrics, we used a Wilcoxon signed-rank test to assess the differences in AUC, as well as sensitivity and FPR at the optimal cut-point of individual subjects (Figure 5.4D). The AUCs of the aHFA amplitude and cHFO amplitude (mean values of 0.959 and 0.96, respectively) were significantly higher than the AUC for the cHFO rate (0.912). Furthermore, the FPR values at the optimal cut-points of the aHFA amplitude (mean of 5.6%) and cHFO amplitude (6.2%) were significantly lower than for cHFO rate (13.3%). For the sensitivity, only the cHFO amplitude (mean of 93.6%) was significantly higher than the cHFO rate (86.0%) at a p -value of 0.05. The sensitivity for aHFA amplitude (93.6%) was not significantly different than the sensitivity for cHFO rate, likely due to the presence of two outlier points in the distribution (Figure 5.4D). In comparing the ADA and RMS detectors, the aHFA amplitude had a higher mean AUC, higher mean sensitivity, and lower mean FPR than cHFO amplitude, but these differences were not statistically significant.

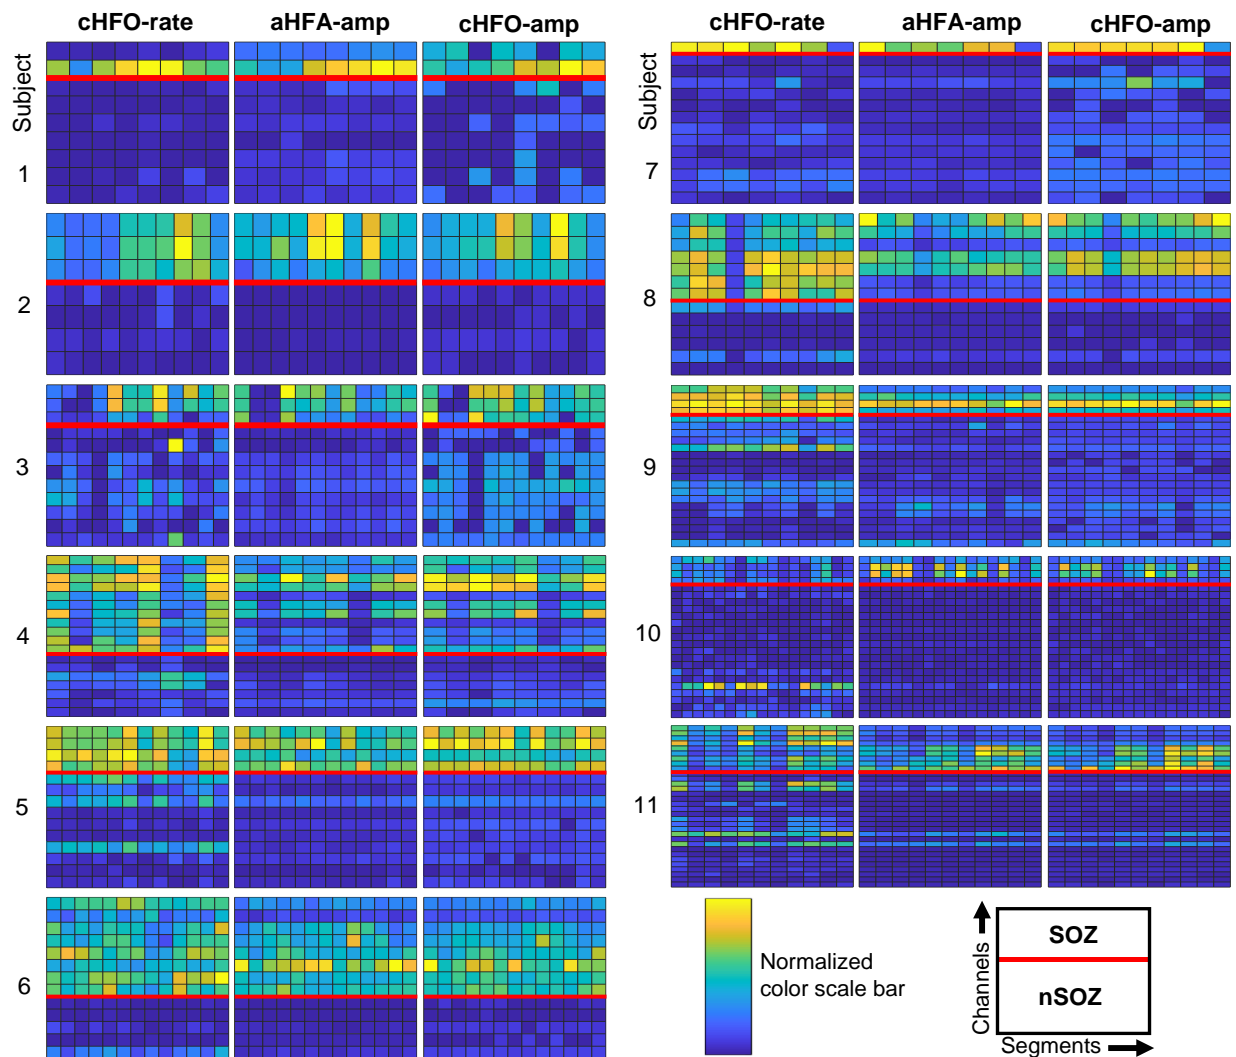


Figure 5.3 Amplitude exhibits greater differences between SOZ and nSOZ channels than cHFO rate.

Heatmaps illustrate the cHFO rate (left sub-panel), aHFA amplitude (center sub-panel) and cHFO amplitude (right sub-panel) for all channels and all iEEG segments. Each set of three sub-panels shows data from one subject. In each sub-panel, each column represents data from one 3-minute iEEG segment for all channels. Each row shows the data from a single channel, with the channels divided into SOZ (above red line) and nSOZ (below red line). The color represents the mean rate or amplitude of the detected events during one 3-minute segment of iEEG, normalized by dividing by the maximum value within each individual subject.

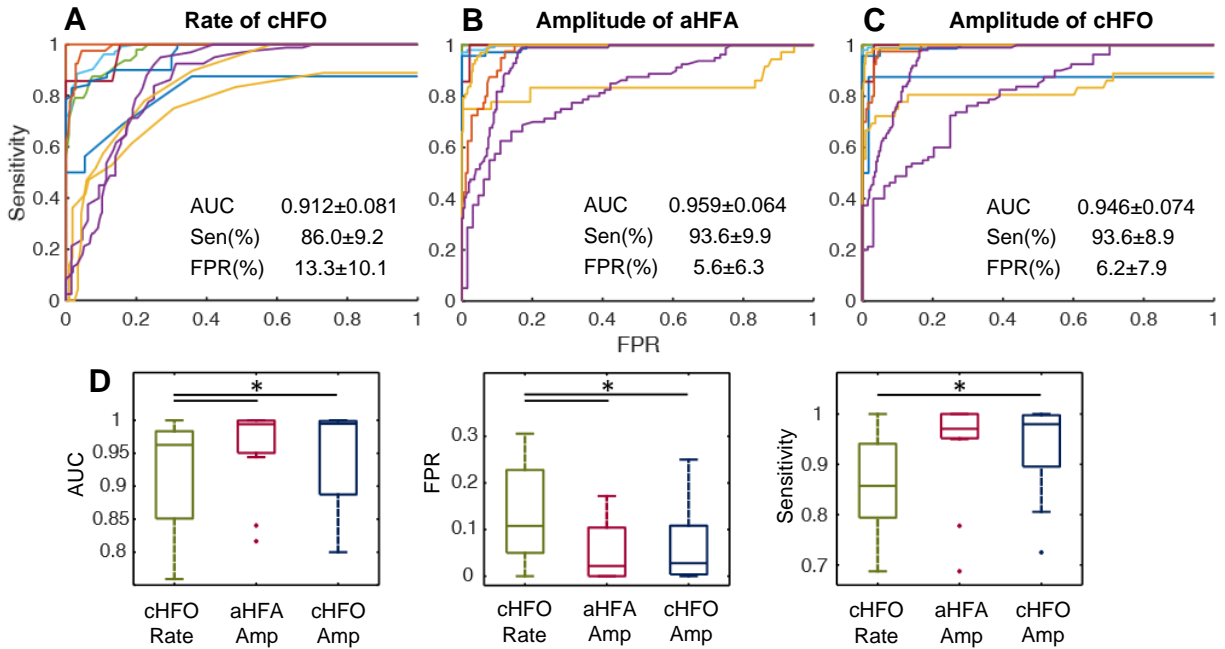


Figure 5.4 Classification of SOZ and nSOZ channels in individual subjects, amplitude provides better performance than cHFO rate.

ROC curves are shown for (A) cHFO rate, (B) aHFA amplitude, and (C) cHFO amplitude for classification of SOZ and nSOZ channels. The ROC curves were calculated using the segment-based scheme, and each line in the subfigures represents an individual subject. (D) Boxplots of AUC, FPR, and sensitivity at the optimal cut-point for all 11 subjects computed from the ROC curves in (A-C). Results are shown for all three metrics: cHFO rate (green), aHFA amplitude (red), and cHFO amplitude (blue). The amplitude metrics have significantly higher AUC and lower FPR than the cHFO rate, and the cHFO amplitude has significantly higher sensitivity than cHFO rate. * p-value < 0.05, Wilcoxon signed-rank test.

5.3.4 Amplitude enables more accurate classification of SOZ and nSOZ than rate at the group level

The overall difference in performance between the amplitude and rate metrics can be observed by generating the segment-based (Figure 5.5A) and channel-based (Figure 5.5B) ROC curves for all patients pooled together. Both methods gave comparable results. The segment-based ROC curves indicated similar performance between the amplitude of aHFA and cHFO (AUC: 0.946 and 0.939, respectively), and both outperformed the rate of cHFO (AUC: 0.879) (Figure 5.5A). Similarly, for the channel-based method, the AUC values were 0.956 for aHFA amplitude, 0.951 for cHFO amplitude, and 0.912 for cHFO rate (Figure 5.5B). The segment-based method demonstrates the performance across measurements from different points in time, and the channel-based method matches the potential clinical use of the metric, where individual channels would be classified as SOZ or nSOZ during surgical planning. Both approaches indicate the robustness of the results when using a common cut-point across all subjects, as opposed to the individualized cut-points used in Figure 5.4.

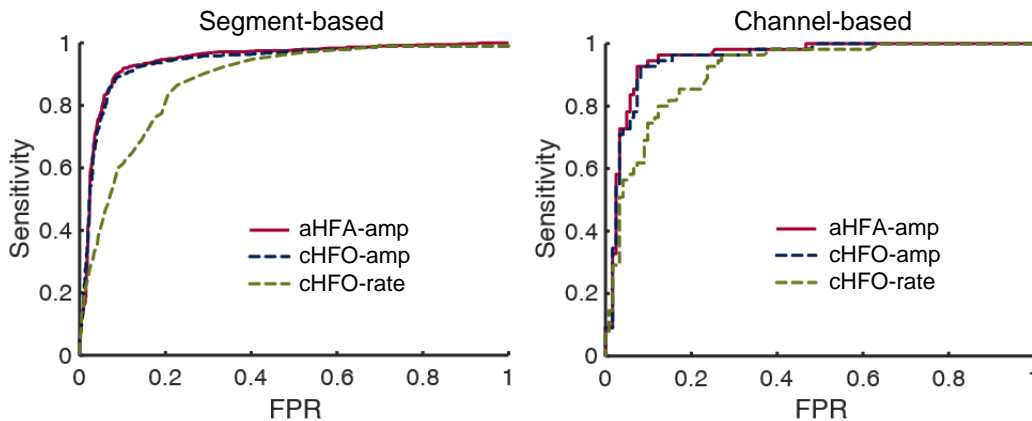


Figure 5.5 Amplitude provides superior classification of SOZ and nSOZ channels with data pooled across all subjects.

ROC curves are shown for cHFO rate (green), aHFA amplitude (red), and cHFO amplitude (blue). (A) ROC curves calculated using the segment-based method. (B) ROC curves calculated

using the channel-based method. The amplitude metrics result in higher AUC, higher optimal sensitivity, and lower optimal FPR using both evaluation methods.

5.3.5 The amplitude metric is more robust than rate to changes in detection sensitivity

In addition to the advantage in the measurement consistency (Figure 5.2C), we found that the amplitude was robust to changes in the detection parameters, while the rate was not. We varied the detection threshold for each algorithm (the value of N in the separation threshold for ADA and the number of standard deviations used for the RMS threshold in the RMS detector), and we measured the resulting AUC values after classification of SOZ and nSOZ channels (Figure 5.6). Here, we pooled all segments together and used the cut-point to classify individual segments to illustrate the overall performance of the metrics. The AUC values for the aHFA rate and cHFO rate were highly dependent on the threshold. For example, varying the value of N in ADA caused the AUC for aHFA rate to increase from 0.394 ($N=1$) to 0.848 ($N=7$). Similarly, when increasing the energy threshold in the RMS detector, the AUC increased from 0.563 to 0.901 ($N=9$) for cHFO rate. This shows that optimization of the threshold is critical to obtaining high accuracy when using the rate of events as a metric for classification. In contrast, the AUC values computed using the aHFA amplitude (range of 0.817 ($N=11$) to 0.946($N=1$)) and cHFO amplitude (range of 0.927($N=1$) to 0.783 ($N=17$)) were relatively robust to the alteration of the threshold.




Figure 5.6: The amplitude metric is robust to changes in the detector sensitivity.

For ADA, the number of standard deviations above the mean in the separation threshold was varied from $N=1$ to $N=11$. For the RMS detector, the threshold applied to the RMS amplitude was varied from one to 15 standard deviations above the mean. Results are shown for the amplitude and rate of aHFA and cHFO, and the segment-based calculation was used for the AUC. Within the range of $N < 10$, the performance is consistently high for the amplitude metric, but the results for rate are highly variable and depend on the choice of threshold.

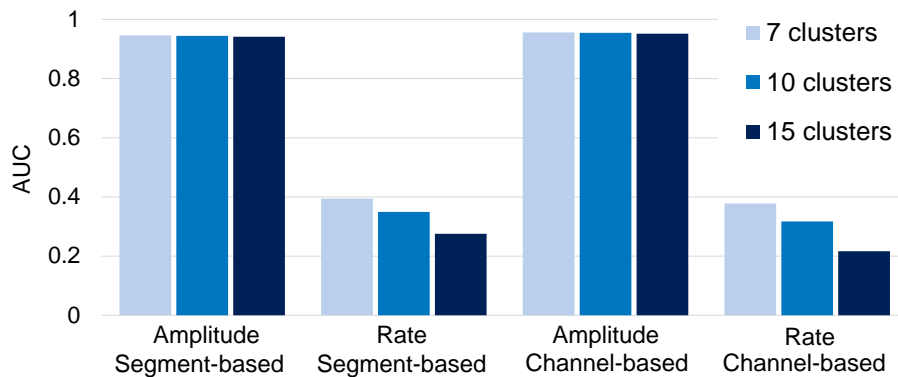


Figure 5.7 The robustness of the metrics to the changes of the maximum number of clusters in clustering and classification process.

For ADA and RMS detector, the separation threshold and RMS threshold were set as default setting (ADA: $N=1$, and RMS: $N=5$). The maximum number of clusters was varied from seven

(light blue) to fifteen (dark blue). Results show the AUC calculated from the ROC curves of the amplitude and rate of aHFA and cHFO using the segment-based and channel-based. The performance is consistently high for the amplitude metric of both aHFA and cHFO. On the contrary, the results for rate are variable and depend on the choice of maximum number of clusters; however, the affect is far less than the choice of threshold. The result also confirmed the finding that the change in detected population, which can be seen from the variation of AUCs of the rate, was not affected the amplitude as a metric of SOZ separation.

5.4 Discussion

Here, we present evidence that the HFO amplitude is superior to the HFO rate as a candidate biomarker of the SOZ due to its increased localization accuracy, measurement stability, and robustness to changes in detection parameters. Consistent with prior studies, we found that the rate of cHFO can be used to classify SOZ and nSOZ electrodes (average AUC of 0.912, with 86% sensitivity and 13% FPR at the optimal cut-point). However, in a head-to-head comparison, the use of aHFA amplitude for classification led to a 5% increase in AUC, 9% increase in sensitivity, and 58% decrease in FPR relative to the cHFO rate. The performance using cHFO amplitude was similar to that of aHFA amplitude. Moreover, we found that the amplitude measurements were more consistent over time, as indicated by lower CV values, and the classification performance using amplitude was robust to changes in the detection threshold.

Previous literature supports the idea that HFO amplitude is a potential marker of the SOZ. The amplitude of conventional HFOs was shown to be significantly higher in SOZ than nSOZ channels for both ripples(Alkawadri et al., 2014; Pail et al., 2017) and fast ripples,(Pail et al., 2017) although one report suggested that the difference was not significant for fast

ripples.(Alkawadri et al., 2014) The studies showed that pathological HFOs have higher mean spectral amplitude, longer mean duration, and lower mean frequency than induced physiological HFOs.(Matsumoto et al., 2013; Nagasawa et al., 2012) However, the reported amplitude distributions for SOZ and nSOZ channels are often overlapping.(Alkawadri et al., 2014; Matsumoto et al., 2013) The classification of SOZ and nSOZ channels using amplitude was usually inferior to the rate of HFOs.(Alkawadri et al., 2014; Matsumoto et al., 2013; Nagasawa et al., 2012) In addition, the amplitude of HFOs (especially ripples) in multiregional or nonlocalizable epilepsy have lower amplitude than those associated with focal neocortical epilepsy.(Wang et al., 2013) In Von Ellenrieder 2016, the difference of the amplitude between HFOs marked within SOZ was significantly higher than nSOZ, which matches with our result.(Nicolá Von Ellenrieder, Frauscher, Dubeau, & Gotman, 2016) However, there were studies using actual amplitude as a main metric to distinguish the SOZ and nSOZ.

We hypothesized that the reported differences in amplitude were small in most studies due to the choice of automated detection algorithm and the selection of detection parameters together with the difference in condition of the comparison. For example, if an energy or amplitude threshold is used as a criterion for detection, the choice of threshold will affect the estimate of event amplitude. As the threshold increases, so will the estimated mean amplitude for both SOZ and nSOZ channels. In Nagasawa 2012 and Alkawadri 2014, the detection method was relied on the visual marking by reviewers; therefore, the marked events were high in amplitude in both SOZ and nSOZ.(Alkawadri et al., 2014; Nagasawa et al., 2012) Therefore, we tested this hypothesis using two complementary detection techniques. We employed the RMS detector to represent standard threshold-based detection. This detector was used to obtain the performance benchmark for classification accuracy, as it has been one of the most widely used

automatic HFO detection algorithms in studies of high frequency activity.(Cho et al., 2014; Gardner et al., 2007; Gliske et al., 2016; Murphy et al., 2017; M. Wu et al., 2018) However, it is important to note that none of the prior studies reporting differences in amplitude used this detection algorithm. For comparison, we used the ADA, which is based on machine learning techniques and does not require the selection of a strict amplitude or energy threshold. In addition, the comparison conditions also heavily affect the results. In Matsumoto 2013 and Pail 2017, the amplitude used in the studies were the normalized amplitude and relative amplitude, respectively,(Matsumoto et al., 2013; Pail et al., 2017) and the result showed that the difference was not as great as the study from Von Ellenrieder 2016.(Nicolá Von Ellenrieder et al., 2016) Nagasawa 2012 and Matsumoto 2013 also reported the difference between pathological and induced physiological HFOs, which tend to be higher amplitude than non-induced events(Matsumoto et al., 2013; Nagasawa et al., 2012).

The ADA and RMS detectors identified an overlapping set of events in the high frequency iEEG, with subsets of events detected only by ADA and only by RMS (Chapter IV). Generally, the amount of overlap between aHFA and cHFO increased as the threshold for each detection algorithm increased. At a threshold of $N=1$ (one SD above the mean), 82% of the aHFA in the SOZ were also detected by the RMS detector. This percentage increased to 91% for $N=5$. This indicates that the ADA detects events in the SOZ that fit the conventional, empirical definition of an HFO. In nSOZ channels, the percent overlap was lower, increasing from 21% to 50% for $N=1$ to $N=5$. This is because the ADA detected many low and intermediate amplitude events which did not exceed the RMS amplitude threshold.

Surprisingly, both the ADA and RMS detectors provided estimates of HFO amplitude that could be used to classify SOZ and nSOZ channels. The differences in performance between

the two detectors were not statistically significant. This suggests that existing detection algorithms can be used to measure HFO amplitude instead of rate, if the threshold is chosen carefully. However, in the long term, new detection algorithms and strategies for parameter optimization should be developed specifically for amplitude. For example, ADA offers several potential advantages over existing algorithms. It does not require any training or parameter optimization, it is fully unsupervised, and it does not require prior assumptions about the shape or amplitude of the high frequency events. This may have been reflected in the results, as Figure 5.5 shows a trend of better performance for ADA amplitude versus RMS amplitude. Analysis of a larger cohort of subjects and investigation into the causes of outliers can help determine whether this difference is real.

In addition to higher classification accuracy for HFO amplitude compared to HFO rate, we found that the performance using amplitude was more robust to changes in detection parameters. This may be because every additional detected event will increase the measurement of rate, but one event is unlikely to considerably alter the median amplitude across all events. For both ADA and the RMS detector, we showed that changing the threshold for detection had very little impact on the AUC (Figure 5.6). Therefore, ADA can be used as originally designed, without the addition of a separation threshold. In the ADA clustering and classification process, we used a maximum of seven clusters; however, we tested a range of seven to fifteen maximum clusters and found no significant difference in the amplitude AUC when pooling all segments together (Figure 5.7).

There were several limitations to our study. We analyzed data from 11 subjects, with approximately 30 minutes of iEEG per patient; however, this is a relatively small number of subjects, considering the diversity of epilepsy syndromes. Moving forward, it will be necessary

to obtain a larger cohort of subjects and analyze a larger percentage of electrodes for each one. Here, we analyzed only electrodes that were clearly localized to gray matter and either were in the SOZ or exhibited no epileptic activity. We excluded electrodes on the boundary between gray and white matter and electrodes with frequent epileptiform discharges or in regions of immediate seizure spread. In practice, these electrodes may locate in brain regions being considered for surgical resection, so it will be crucial to understand if HFO amplitude is a valid biomarker in these borderline cases. Furthermore, validation of amplitude as a biomarker of the epileptogenic zone must incorporate data regarding the exact resected brain region and long-term surgical outcome. Our dataset included one subject that ultimately did not have resective surgery and three subjects who received implanted neurostimulators, which precluded a comparison to surgical outcome. Future work will address these limitations, with the long-term goal of identifying robust, accurate biomarkers of the epileptogenic zone to aid surgical planning. Ultimately, this will increase the percentage of patients with refractory epilepsy who are seizure free after surgery and may also increase the number of patients who are offered surgery as a treatment option.

REFERENCES

- Alkawadri, R., Gaspard, N., Goncharova, I. I., Spencer, D. D., Gerrard, J. L., Zaveri, H., ... Hirsch, L. J. (2014). The spatial and signal characteristics of physiologic high frequency oscillations. *Epilepsia*, 55(12), 1986–1995. <https://doi.org/10.1111/epi.12851>
- Amiri, M., Lina, J. M., Pizzo, F., & Gotman, J. (2016). High Frequency Oscillations and spikes: Separating real HFOs from false oscillations. *Clinical Neurophysiology*, 127(1), 187–196. <https://doi.org/10.1016/j.clinph.2015.04.290>
- Andrade-Valenca, L. P., Dubeau, F., Mari, F., Zelmann, R., & Gotman, J. (2011a). Interictal scalp fast oscillations as a marker of the seizure onset zone. *Neurology*, 77(6), 524–531. <https://doi.org/10.1212/WNL.0b013e318228bee2>
- Andrade-Valenca, L. P., Dubeau, F., Mari, F., Zelmann, R., & Gotman, J. (2011b). Interictal scalp fast oscillations as a marker of the seizure onset zone. *Neurology*, 77(6), 524–531. <https://doi.org/10.1212/WNL.0b013e318228bee2>
- Avoli, M. (2001). Do interictal discharges promote or control seizures? Experimental evidence from an in vitro model of epileptiform discharge. *Epilepsia*, 42 Suppl 3, 2–4.
- Bagshaw, A. P., Jacobs, J., Levan, P., Dubeau, F., & Gotman, J. (2009). Effect of sleep stage on interictal high-frequency oscillations recorded from depth macroelectrodes in patients with focal epilepsy. *Epilepsia*, 50(4), 617–628. <https://doi.org/10.1111/j.1528-1167.2008.01784.x>
- Bautista, R. E. D., Cobbs, M. A., Spencer, D. D., & Spencer, S. S. (1999). Prediction of Surgical Outcome by Interictal Epileptiform Abnormalities During Intracranial EEG Monitoring in Patients with Extrahippocampal Seizures. *Epilepsia*, 40(7), 880–890. <https://doi.org/10.1111/j.1528-1157.1999.tb00794.x>
- Benbadis, S. R. (2001). Epileptic seizures and syndromes. *Neurologic Clinics*, 19(2), 251–270. [https://doi.org/10.1016/S0733-8619\(05\)70018-9](https://doi.org/10.1016/S0733-8619(05)70018-9)
- Berg, A. T., Berkovic, S. F., Brodie, M. J., Buchhalter, J., Cross, J. H., Van Emde Boas, W., ... Scheffer, I. E. (2010). Revised terminology and concepts for organization of seizures and epilepsies: Report of the ILAE Commission on Classification and Terminology, 2005-2009. *Epilepsia*, 51(4), 676–685. <https://doi.org/10.1111/j.1528-1167.2010.02522.x>
- Bernardo, D., Nariai, H., Hussain, S. A., Sankar, R., Salamon, N., Krueger, D. A., ... Wu, J. Y. (2018). Visual and semi-automatic non-invasive detection of interictal fast ripples: A potential biomarker of epilepsy in children with tuberous sclerosis complex. *Clinical Neurophysiology*, 129(7). <https://doi.org/10.1016/j.clinph.2018.03.010>
- Biro, G., Kachenoura, A., Albera, L., Bénar, C., & Wendling, F. (2013). Automatic detection of fast ripples. *Journal of Neuroscience Methods*, 213(2), 236–249. <https://doi.org/10.1016/j.jneumeth.2012.12.013>
- Blanco, J. A., Stead, M., Krieger, A., Stacey, W., Maus, D., Marsh, E., ... Worrell, G. A. (2011). Data mining neocortical high-frequency oscillations in epilepsy and controls. *Brain*, 134(10), 2948–2959. <https://doi.org/10.1093/brain/awr212>
- Blanco, J. A., Stead, M., Krieger, A., Viventi, J., Marsh, W. R., Lee, K. H., ... Litt, B. (2010). Unsupervised classification of high-frequency oscillations in human neocortical epilepsy and control patients. *Journal of Neurophysiology*, 104(5), 2900–2912. <https://doi.org/10.1152/jn.01082.2009>
- Bragin, A., Engel, J., & Staba, R. J. (2010). High-frequency oscillations in epileptic brain. *Current Opinion in Neurology*, 23(2), 151–156. <https://doi.org/10.1097/WCO.0b013e3283373ac8>
- Bragin, A., Engel, J., Wilson, C. L., Fried, I., & Buzsáki, G. (1999). High-frequency oscillations in human brain. *Hippocampus*, 9(2), 137–142. [https://doi.org/10.1002/\(SICI\)1098-1063\(1999\)9:2<137::AID-HIPO5>3.0.CO;2-0](https://doi.org/10.1002/(SICI)1098-1063(1999)9:2<137::AID-HIPO5>3.0.CO;2-0)
- Bragin, A., Engel, J., Wilson, C. L., Fried, I., & Mathern, G. W. (1999). Hippocampal and entorhinal cortex high-frequency oscillations (100-500 Hz) in human epileptic brain and in kainic acid-treated

- rats with chronic seizures. *Epilepsia*, 40(2), 127–137. <https://doi.org/10.1111/j.1528-1157.1999.tb02065.x>
- Bragin, A., Mody, I., Wilson, C. L., & Engel, J. (2002). Local Generation of Fast Ripples in Epileptic Brain. *The Journal of Neuroscience*, 22(5), 2012 LP – 2021. <https://doi.org/10.1523/JNEUROSCI.22-05-02012.2002>
- Bragin, A., Wilson, C. L., & Engel, J. (2007). Voltage depth profiles of high-frequency oscillations after kainic acid-induced status epilepticus. *Epilepsia*, 48(SUPPL. 5), 35–40. <https://doi.org/10.1111/j.1528-1167.2007.01287.x>
- Brázdil, M., Pail, M., Halánek, J., Plešinger, F., Cimbálník, J., Roman, R., ... Jurák, P. (2017). Very high-frequency oscillations: Novel biomarkers of the epileptogenic zone. *Annals of Neurology*, 82(2), 299–310. <https://doi.org/10.1002/ana.25006>
- Burnos, S., Hilfiker, P., Sürücü, O., Scholkmann, F., Krayenbühl, N., Grunwald, T., & Sarnthein, J. (2014). Human Intracranial High Frequency Oscillations (HFOs) Detected by Automatic Time-Frequency Analysis. *PLoS ONE*, 9(4), e94381. <https://doi.org/10.1371/journal.pone.0094381>
- Buzsáki, G., Horváth, Z., Urioste, R., Hetke, J., & Wise, K. (1992). High-frequency network oscillation in the hippocampus. *Science (New York, N.Y.)*, 256(5059), 1025–1027. <https://doi.org/10.1126/science.1589772>
- Buzsáki, György, Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nat Rev Neurosci*, 13(6), 407–420. Retrieved from <http://dx.doi.org/10.1038/nrn3241>
- Buzsáki, György, & Silva, F. L. da. (2012). High frequency oscillations in the intact brain. *Progress in Neurobiology*, 98(3), 241–249. <https://doi.org/10.1016/j.pneurobio.2012.02.004>
- Chaibi, S., Lajnef, T., Sakka, Z., Samet, M., & Kachouri, A. (2013). A comparison of methods for detection of high frequency oscillations (HFOs) in human intracerebral EEG recordings. *American Journal of ...*, 3(2), 25–34. <https://doi.org/10.5923/j.ajsp.20130302.02>
- Charupanit, K., & Lopour, B. A. (2017). A Simple Statistical Method for the Automatic Detection of Ripples in Human Intracranial EEG. *Brain Topography*, 30(6), 724–738. <https://doi.org/10.1007/s10548-017-0579-6>
- Charupanit, Krit, Nunez, M. D., Bernardo, D., Bebin, M., Krueger, D. A., Northrup, H., ... Lopour, B. A. (2018). Automated Detection of High Frequency Oscillations in Human Scalp Electroencephalogram. *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS, 2018-July*, 3116–3119. <https://doi.org/10.1109/EMBC.2018.8513033>
- Chiba, S. (1978). Dynamic Programming Algorithm Optimization for Spoken Word.pdf, (1).
- Cho, J. R., Koo, D. L., Joo, E. Y., Seo, D. W., Hong, S. C., Jiruska, P., & Hong, S. B. (2014). Resection of individually identified high-rate high-frequency oscillations region is associated with favorable outcome in neocortical epilepsy. *Epilepsia*, 55(11), 1872–1883. <https://doi.org/10.1111/epi.12808>
- Chu, C. J., Chan, A., Song, D., Staley, K. J., Stufflebeam, S. M., & Kramer, M. A. (2017). A semi-automated method for rapid detection of ripple events on interictal voltage discharges in the scalp electroencephalogram. *Journal of Neuroscience Methods*, 277, 46–55. <https://doi.org/10.1016/j.jneumeth.2016.12.009>
- Cimbálník, J., Hewitt, A., Worrell, G., & Stead, M. (2018). The CS algorithm: A novel method for high frequency oscillation detection in EEG. *Journal of Neuroscience Methods*, 293, 6–16. <https://doi.org/10.1016/j.jneumeth.2017.08.023>
- Crépon, B., Navarro, V., Hasboun, D., Clemenceau, S., Martinerie, J., Baulac, M., ... Le Van Quyen, M. (2010). Mapping interictal oscillations greater than 200 Hz recorded with intracranial macroelectrodes in human epilepsy. *Brain*, 133(1), 33–45. Retrieved from <http://brain.oxfordjournals.org/content/133/1/33.abstract>
- Curtis, M. De, & Avanzini, G. (2001). <de Cutris 2001 - Interictal spikes in focal epileptogenesis.pdf>, 63, 541–567.
- Curtis, M. De, Jefferys, J. G. R., & Avoli, M. (2012). Interictal Epileptiform Discharges in Partial

- Epilepsy Different IED Patterns in Epileptic Patients : Spikes , Spike Bursts , 1–24.
- Dümpelmann, M., Jacobs, J., Kerber, K., & Schulze-Bonhage, A. (2012). Automatic 80-250Hz “ripple” high frequency oscillation detection in invasive subdural grid and strip recordings in epilepsy by a radial basis function neural network. *Clinical Neurophysiology*, *123*(9), 1721–1731. <https://doi.org/10.1016/j.clinph.2012.02.072>
- Dümpelmann, M., Jacobs, J., & Schulze-Bonhage, A. (2015). Temporal and spatial characteristics of high frequency oscillations as a new biomarker in epilepsy. *Epilepsia*, *56*(2), 197–206. <https://doi.org/10.1111/epi.12844>
- Engel, J. (1993). Update on surgical treatment of the epilepsies. *Neurology*, *43*(8), 1612 LP – 1612. <https://doi.org/10.1212/WNL.43.8.1612>
- Engel, J., Bragin, A., Staba, R., Mody, I., Engel Jr, J., Bragin, A., ... Mody, I. (2009). High-frequency oscillations: What is normal and what is not? *Epilepsia*, *50*(4), 598–604. <https://doi.org/10.1111/j.1528-1167.2008.01917.x>
- Engel, J., & da Silva, F. L. (2012). High-frequency oscillations - Where we are and where we need to go. *Progress in Neurobiology*, *98*(3), 316–318. <https://doi.org/10.1016/j.pneurobio.2012.02.001>
- Fedele, T., van ‘t Klooster, M., Burnos, S., Zweiphenning, W., van Klink, N., Leijten, F., ... Sarnthein, J. (2016). Automatic detection of high frequency oscillations during epilepsy surgery predicts seizure outcome. *Clinical Neurophysiology*, *127*(9). <https://doi.org/10.1016/j.clinph.2016.06.009>
- Ferrari-Marinho, T., Perucca, P., Mok, K., Olivier, A., Hall, J., Dubeau, F., & Gotman, J. (2015). Pathologic substrates of focal epilepsy influence the generation of high-frequency oscillations. *Epilepsia*, *56*(4), 592–598. <https://doi.org/10.1111/epi.12940>
- Fisher, R. S., Boas, W. van E., Blume, W., Elger, C., Genton, P., Lee, P., & Engel, J. (2005). Epileptic Seizures and Epilepsy: Definitions Proposed by the International League Against Epilepsy (ILAE) and the International Bureau for Epilepsy (IBE). *Epilepsia*, *46*(4), 470–472. <https://doi.org/10.1111/j.0013-9580.2005.66104.x>
- Fisher, R. S., Cross, J. H., French, J. A., Higurashi, N., Hirsch, E., Jansen, F. E., ... Zuberi, S. M. (2018). Operational classification of seizure types by the International League Against Epilepsy: position paper of the ILAE Commission for Classification and Terminology. *Zeitschrift Fur Epileptologie*, *31*(4), 272–281. <https://doi.org/10.1007/s10309-018-0216-8>
- Frauscher, B., Bartolomei, F., Kobayashi, K., Cimbalnik, J., van ‘t Klooster, M. A., Rampp, S., ... Gotman, J. (2017). High-frequency oscillations: The state of clinical research. *Epilepsia*, *58*(8), 1316–1329. <https://doi.org/10.1111/epi.13829>
- Frauscher, B., von Ellenrieder, N., Zelmann, R., Rogers, C., Nguyen, D. K., Kahane, P., ... Gotman, J. (2018). High-Frequency Oscillations in the Normal Human Brain. *Annals of Neurology*, *84*(3). <https://doi.org/10.1002/ana.25304>
- French, J. A. (2007). Refractory epilepsy: Clinical overview. *Epilepsia*, *48*(SUPPL. 1), 3–7. <https://doi.org/10.1111/j.1528-1167.2007.00992.x>
- Fujiwara, H., Greiner, H. M., Lee, K. H., Holland-Bouley, K. D., Seo, J. H., Arthur, T., ... Rose, D. F. (2012). Resection of ictal high-frequency oscillations leads to favorable surgical outcome in pediatric epilepsy. *Epilepsia*, *53*(9), 1607–1617. <https://doi.org/10.1111/j.1528-1167.2012.03629.x>
- Garcia Barba, C. (2014). The Brain Topology of Fast Ripples, and their Correlation with Epilepsy. *Journal of Neurological Disorders*, *02*(04). <https://doi.org/10.4172/2329-6895.1000162>
- Gardner, A. B., Worrell, G. A., Marsh, E., Dlugos, D., & Litt, B. (2007). Human and automated detection of high-frequency oscillations in clinical intracranial EEG recordings. *Clinical Neurophysiology*, *118*(5), 1134–1143. <https://doi.org/10.1016/j.clinph.2006.12.019>
- Gliske, S. V., Irwin, Z. T., Chestek, C., Hegeman, G. L., Brinkmann, B., Sagher, O., ... Hegeman, G. L. (2018). Variability in the location of high frequency oscillations during prolonged intracranial EEG recordings. *Nature Communications*, *9*(1). <https://doi.org/10.1038/s41467-018-04549-2>
- Gliske, S. V., Irwin, Z. T., Davis, K. A., Sahaya, K., Chestek, C., & Stacey, W. C. (2016). Universal automated high frequency oscillation detector for real-time, long term EEG. *Clinical Neurophysiology*, *127*(2), 1057–1066. <https://doi.org/10.1016/j.clinph.2015.07.016>

- Gotman, J. (1991). Relationships Between Interictal Spiking and Seizures: Human and Experimental Evidence. *Canadian Journal of Neurological Sciences / Journal Canadien Des Sciences Neurologiques*, 18(S4), 573–576. <https://doi.org/10.1017/S031716710003273X>
- Grange, W., Haas, P., Wild, A., Lieb, M. A., Calame, M., Hegner, M., & Hecht, B. (2008). Detection of transient events in the presence of background noise. *Journal of Physical Chemistry B*, 112(23), 7140–7144. <https://doi.org/10.1021/jp7114862>
- Grenier, F., Timofeev, I., & Steriade, M. (2001). Focal Synchronization of Ripples (80–200 Hz) in Neocortex and Their Neuronal Correlates. *Journal of Neurophysiology*, 86(4), 1884–1898. Retrieved from <http://jn.physiology.org/content/86/4/1884.abstract>
- Guragain, H., Cimbalnik, J., Stead, M., Groppe, D. M., Berry, B. M., Kremen, V., ... Brinkmann, B. H. (2018). Spatial variation in high-frequency oscillation rates and amplitudes in intracranial EEG. *Neurology*, 90(8). <https://doi.org/10.1212/WNL.0000000000004998>
- Haegelen, C., Perucca, P., Châtillon, C.-E., Andrade-Valença, L., Zelmann, R., Jacobs, J., ... Gotman, J. (2013). High-frequency oscillations, extent of surgical resection, and surgical outcome in drug-resistant focal epilepsy. *Epilepsia*, 54(5), 848–857. <https://doi.org/10.1111/epi.12075>
- Iimura, Y., Jones, K., Hattori, K., Okazawa, Y., Noda, A., Hoashi, K., ... Otsubo, H. (2017). Epileptogenic high-frequency oscillations skip the motor area in children with multilobar drug-resistant epilepsy. *Clinical Neurophysiology*, 128(7), 1197–1205. <https://doi.org/10.1016/j.clinph.2017.03.031>
- International Federation of Societies for Electroencephalography, and C. N. (1983). *Recommendations for the practice of clinical neurophysiology*. Elsevier Publishing Company.
- Jacobs, J., Staba, R., Asano, E., Otsubo, H., Wu, J. Y., Zijlmans, M., ... Gotman, J. (2012). High-frequency oscillations (HFOs) in clinical epilepsy. *Progress in Neurobiology*, 98(3), 302–315. <https://doi.org/10.1016/j.pneurobio.2012.03.001>
- Jacobs, J., Staba, R., Asano, E., Otsubo, H., Wu, J. Y., Zijlmans, M., ... Gotman, J. (2012). High-frequency oscillations (HFOs) in clinical epilepsy. *Progress in Neurobiology*, 98(3), 302–315. <https://doi.org/10.1016/j.pneurobio.2012.03.001>
- Jacobs, Joshua, & Kahana, M. J. (2010). Direct brain recordings fuel advances in cognitive electrophysiology. *Trends in Cognitive Sciences*, 14(4), 162–171. <https://doi.org/10.1016/j.tics.2010.01.005>
- Jacobs, Julia, Golla, T., Mader, M., Schelter, B., D??mpelmann, M., Korinthenberg, R., & Schulze-Bonhage, A. (2014). Electrical stimulation for cortical mapping reduces the density of high frequency oscillations. *Epilepsy Research*, 108(10), 1758–1769. <https://doi.org/10.1016/j.eplepsyres.2014.09.022>
- Jacobs, Julia, LeVan, P., Chander, R., Hall, J., Dubeau, F., & Gotman, J. (2008). Interictal high-frequency oscillations (80-500 Hz) are an indicator of seizure onset areas independent of spikes in the human epileptic brain. *Epilepsia*, 49(11), 1893–1907. <https://doi.org/10.1111/j.1528-1167.2008.01656.x>
- Jacobs, Julia, Levan, P., Chtillon, C. D., Olivier, A., Dubeau, F., & Gotman, J. (2009). High frequency oscillations in intracranial EEGs mark epileptogenicity rather than lesion type. *Brain*, 132(4), 1022–1037. <https://doi.org/10.1093/brain/awn351>
- Jacobs, Julia, Zijlmans, M., Zelmann, R., Chatillon, C.-É., Hall, J., Olivier, A., ... Gotman, J. (2010). High-frequency electroencephalographic oscillations correlate with outcome of epilepsy surgery. *Annals of Neurology*, 67(2), 209–220. <https://doi.org/10.1002/ana.21847>
- Jefferys, J. G. R., Menendez de la Prida, L., Wendling, F., Bragin, A., Avoli, M., Timofeev, I., & Lopes da Silva, F. H. (2012). Mechanisms of physiological and epileptic HFO generation. *Progress in Neurobiology*, 98(3), 250–264. <https://doi.org/10.1016/j.pneurobio.2012.02.005>
- Jehi, L. (2018). The epileptogenic zone: Concept and definition. *Epilepsy Currents*, 18(1), 12–16. <https://doi.org/10.5698/1535-7597.18.1.12>
- Jobst, B. C., & Cascino, G. D. (2015). Resective epilepsy surgery for drug-resistant focal epilepsy: A review. *JAMA - Journal of the American Medical Association*, 313(3), 285–293. <https://doi.org/10.1001/jama.2014.17426>

- Keogh, E., Lin, J., & Fu, A. (2005). HOT SAX: Efficiently Finding the Most Unusual Time Series Subsequence. In *Proceedings of the Fifth IEEE International Conference on Data Mining* (pp. 226–233). Washington, DC, USA: IEEE Computer Society. <https://doi.org/10.1109/ICDM.2005.79>
- Khadjevand, F., Cimbalknik, J., & Worrell, G. A. (2017). Progress and Remaining Challenges in the Application of High Frequency Oscillations as Biomarkers of Epileptic Brain. *Current Opinion in Biomedical Engineering*, 4, 87–96. <https://doi.org/10.1016/j.cobme.2017.09.006>
- Khosravani, H., Mehrotra, N., Rigby, M., Hader, W. J., Pinnegar, C. R., Pillay, N., ... Federico, P. (2009). Spatial localization and time-dependant changes of electrographic high frequency oscillations in human temporal lobe epilepsy. *Epilepsia*, 50(4), 605–616. <https://doi.org/10.1111/j.1528-1167.2008.01761.x>
- Kobayashi, K., Watanabe, Y., Inoue, T., Oka, M., Yoshinaga, H., & Ohtsuka, Y. (2010). Scalp-recorded high-frequency oscillations in childhood sleep-induced electrical status epilepticus. *Epilepsia*, 51(10), 2190–2194. <https://doi.org/10.1111/j.1528-1167.2010.02565.x>
- Kucewicz, M. T., Cimbalknik, J., Matsumoto, J. Y., Brinkmann, B. H., Bower, M. R., Vasoli, V., ... Worrell, G. A. (2014). High frequency oscillations are associated with cognitive processing in human recognition memory. *Brain*. Retrieved from <http://brain.oxfordjournals.org/content/early/2014/06/11/brain.awu149.abstract>
- Lachaux, J.-P., Axmacher, N., Mormann, F., Halgren, E., & Crone, N. E. (2012). High-frequency neural activity and human cognition: past, present and possible future of intracranial EEG research. *Progress in Neurobiology*, 98(3), 279–301. <https://doi.org/10.1016/j.pneurobio.2012.06.008>
- Liu, S., Sha, Z., Sencer, A., Aydoseli, A., Bebek, N., Abosch, A., ... Ince, N. F. (2016). Exploring the time-frequency content of high frequency oscillations for automated identification of seizure onset zone in epilepsy. *Journal of Neural Engineering*, 13(2). <https://doi.org/10.1088/1741-2560/13/2/026026>
- Malinowska, U., Bergey, G. K., Harezlak, J., & Jouny, C. C. (2015). Identification of seizure onset zone and preictal state based on characteristics of high frequency oscillations. *Clinical Neurophysiology*, 126(8), 1505–1513. <https://doi.org/10.1016/j.clinph.2014.11.007>
- Malkki, H. (2014). Epilepsy [mdash] burning questions and emerging therapies. *Nature Reviews Neurology*, 10(5), 243.
- Matsumoto, A., Brinkmann, B. H., Matthew Stead, S., Matsumoto, J., Kucewicz, M. T., Marsh, W. R., ... Worrell, G. (2013). Pathological and physiological high-frequency oscillations in focal human epilepsy. *Journal of Neurophysiology*, 110(8), 1958–1964. <https://doi.org/10.1152/jn.00341.2013>
- Melani, F., Zemann, R., Dubeau, F., & Gotman, J. (2013). Occurrence of scalp-fast oscillations among patients with different spiking rate and their role as epileptogenicity marker. *Epilepsy Research*, 106(3), 345–356. <https://doi.org/10.1016/j.eplepsyres.2013.06.003>
- Misra, U. K., & Kalita, J. (2011). Management of provoked seizure. *Annals of Indian Academy of Neurology*, 14(1), 2–8. <https://doi.org/10.4103/0972-2327.78041>
- Moshé, S. L., Perucca, E., Ryvlin, P., & Tomson, T. (2015). Epilepsy: new advances. *The Lancet*, 385(9971), 884–898. [https://doi.org/https://doi.org/10.1016/S0140-6736\(14\)60456-6](https://doi.org/https://doi.org/10.1016/S0140-6736(14)60456-6)
- Murphy, P. M., Paternos, A. J. von, & Santaniello, S. (2017). A novel HFO-based method for unsupervised localization of the seizure onset zone in drug-resistant epilepsy. In *2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 1054–1057). <https://doi.org/10.1109/EMBC.2017.8037008>
- Nagasawa, T., Juhász, C., Rothermel, R., Hoehstetter, K., Sood, S., & Asano, E. (2012). Spontaneous and visually driven high-frequency oscillations in the occipital cortex: Intracranial recording in epileptic patients. *Human Brain Mapping*, 33(3), 569–583. <https://doi.org/10.1002/hbm.21233>
- Pail, M., Řehulka, P., Cimbálník, J., Doležalová, I., Chrastina, J., & Brázdil, M. (2017). Frequency-independent characteristics of high-frequency oscillations in epileptic and non-epileptic regions. *Clinical Neurophysiology*, 128(1), 106–114. <https://doi.org/10.1016/j.clinph.2016.10.011>
- Paliwal, K. K., Agarwal, A., & Sinha, S. S. (1982). A modification over Sakoe and Chiba's dynamic time warping algorithm for isolated word recognition. *Signal Processing*, 4(4), 329–333.

- [https://doi.org/10.1016/0165-1684\(82\)90009-3](https://doi.org/10.1016/0165-1684(82)90009-3)
- Pearce, A., Wulsin, D., Blanco, J. A., Krieger, A., Litt, B., & Stacey, W. C. (2013). Temporal changes of neocortical high-frequency oscillations in epilepsy. *Journal of Neurophysiology*, *110*(5), 1167–1179. <https://doi.org/10.1152/jn.01009.2012>
- Perruca, E. (2005). An Introduction to Antiepileptic Drugs. *Epilepsia*, *46*(4), 31–37.
- Pizzo, F., Frauscher, B., Ferrari-Marinho, T., Amiri, M., Dubeau, F., & Gotman, J. (2016). Detectability of Fast Ripples (>250 Hz) on the Scalp EEG: A Proof-of-Principle Study with Subdermal Electrodes. *Brain Topography*, *29*(3), 358–367. <https://doi.org/10.1007/s10548-016-0481-7>
- Roehri, N., Lina, J., Mosher, J. C., Bartolomei, F., Roehri, N., Lina, J., ... Bartolomei, F. (2016). Time-frequency strategies for increasing high frequency oscillation detectability in intracerebral, *2*(3), 1–12. <https://doi.org/10.1109/TBME.2016.2556425>
- Rosenow, F., & Lüders, H. (2001). Presurgical evaluation of epilepsy. *Brain*, *124*(9), 1683–1700. <https://doi.org/10.1093/brain/124.9.1683>
- Schevon, C. A., Trevelyan, A. J., Schroeder, C. E., Goodman, R. R., McKhann, G., & Emerson, R. G. (2009). Spatial characterization of interictal high frequency oscillations in epileptic neocortex. *Brain*. Retrieved from <http://brain.oxfordjournals.org/content/early/2009/09/10/brain.awp222.abstract>
- Schuele, S. U., & Lüders, H. O. (2008). Intractable epilepsy: management and therapeutic alternatives. *The Lancet Neurology*, *7*(6), 514–524. [https://doi.org/10.1016/S1474-4422\(08\)70108-X](https://doi.org/10.1016/S1474-4422(08)70108-X)
- Spencer, S., & Huh, L. (2008). Outcomes of epilepsy surgery in adults and children. *The Lancet Neurology*, *7*(6), 525–537. [https://doi.org/10.1016/S1474-4422\(08\)70109-1](https://doi.org/10.1016/S1474-4422(08)70109-1)
- Spring, A. M., Pittman, D. J., Aghakhani, Y., Jirsch, J., Pillay, N., Bello-Espinosa, L. E., ... Federico, P. (2017). Interrater reliability of visually evaluated high frequency oscillations. *Clinical Neurophysiology*, *128*(3), 433–441. <https://doi.org/10.1016/j.clinph.2016.12.017>
- Spring, A. M., Pittman, D. J., Aghakhani, Y., Jirsch, J., Pillay, N., Bello-Espinosa, L. E., ... Federico, P. (2018). Generalizability of High Frequency Oscillation Evaluations in the Ripple Band. *Frontiers in Neurology*, *9*(June), 510. <https://doi.org/10.3389/fneur.2018.00510>
- Staba, R. J., Frigetto, L., Behnke, E. J., Mathern, G. W., Fields, T., Bragin, A., ... Engel, J. (2007). Increased fast ripple to ripple ratios correlate with reduced hippocampal volumes and neuron loss in temporal lobe epilepsy patients. *Epilepsia*, *48*(11), 2130–2138. <https://doi.org/10.1111/j.1528-1167.2007.01225.x>
- Staba, R. J., Wilson, C. L., Bragin, A., Fried, I., & Engel, Jerome, J. (2002). Quantitative Analysis of High-Frequency Oscillations (80-500 Hz) Recorded in Human Epileptic Hippocampus and Entorhinal Cortex. *J Neurophysiol*, *88*(4), 1743–1752. Retrieved from <http://jn.physiology.org/content/88/4/1743>
- Staba, R. J., Wilson, C. L., Bragin, A., Jhung, D., Fried, I., & Engel, J. (2004). High-frequency oscillations recorded in human medial temporal lobe during sleep. *Annals of Neurology*, *56*(1), 108–115. <https://doi.org/10.1002/ana.20164>
- Talairach, J., & Bancaud, J. (1966). Lesion, "irritative" zone and epileptogenic focus. *Stereotactic and Functional Neurosurgery*, *27*(1–3), 91–94.
- Theodore, W. H., & Fisher, R. S. (2004). Review Brain stimulation for epilepsy. *The Lancet*, *3*(February), 111–118.
- Urrestarazu, E., Chander, R., Dubeau, F., & Gotman, J. (2007). Interictal high-frequency oscillations (10-500 Hz) in the intracerebral EEG of epileptic patients. *Brain*, *130*(9), 2354–2366. <https://doi.org/10.1093/brain/awm149>
- van Klink, N., Frauscher, B., Zijlmans, M., & Gotman, J. (2016). Relationships between interictal epileptic spikes and ripples in surface EEG. *Clinical Neurophysiology*, *127*(1), 143–149. <https://doi.org/10.1016/j.clinph.2015.04.059>
- Von Ellenrieder, Nicolás, Frauscher, B., Dubeau, F., & Gotman, J. (2016). Interaction with slow waves during sleep improves discrimination of physiologic and pathologic high-frequency oscillations (80-500 Hz). *Epilepsia*, *57*(6), 869–878. <https://doi.org/10.1111/epi.13380>

- von Ellenrieder, Nicolás, Andrade-Valença, L. P., Dubeau, F., & Gotman, J. (2012). Automatic detection of fast oscillations (40–200Hz) in scalp EEG recordings. *Clinical Neurophysiology*, *123*(4), 670–680. <https://doi.org/10.1016/j.clinph.2011.07.050>
- Wang, S., Wang, I. Z., Bulacio, J. C., Mosher, J. C., Gonzalez-Martinez, J., Alexopoulos, A. V., ... So, N. K. (2013). Ripple classification helps to localize the seizure-onset zone in neocortical epilepsy. *Epilepsia*, *54*(2), 370–376. <https://doi.org/10.1111/j.1528-1167.2012.03721.x>
- Worrell, G. a., Gardner, A. B., Stead, S. M., Hu, S., Goerss, S., Cascino, G. J., ... Litt, B. (2008). High-frequency oscillations in human temporal lobe: Simultaneous microwire and clinical macroelectrode recordings. *Brain*, *131*(4), 928–937. <https://doi.org/10.1093/brain/awn006>
- Worrell, G A, Jerbi, K., Kobayashi, K., Lina, J. M., Zelmann, R., & Le Van Quyen, M. (2012). Recording and analysis techniques for high-frequency oscillations. *Progress in Neurobiology*, *98*(3), 265–278. <https://doi.org/10.1016/j.pneurobio.2012.02.006>
- Worrell, Greg A., Parish, L., Cranstoun, S. D., Jonas, R., Baltuch, G., & Litt, B. (2004). High-frequency oscillations and seizure generation in neocortical epilepsy. *Brain*, *127*(7), 1496–1506. <https://doi.org/10.1093/brain/awh149>
- Wu, J. Y., Sankar, R., Lerner, J. T., Matsumoto, J. H., Vinters, H. V., & Mathern, G. W. (2010). Removing interictal fast ripples on electrocorticography linked with seizure freedom in children. *Neurology*, *75*(19), 1686 LP – 1694. <https://doi.org/10.1212/WNL.0b013e3181fc27d0>
- Wu, Joyce Y., Koh, S., Sankar, R., & Mathern, G. W. (2008). Paroxysmal fast activity: An interictal scalp EEG marker of epileptogenesis in children. *Epilepsy Research*, *82*(1), 101–108. <https://doi.org/10.1016/j.eplepsyres.2008.07.010>
- Wu, M., Wan, T., Ding, M., Wan, X., Du, Y., & She, J. (2018). A New Unsupervised Detector of High-Frequency Oscillations in Accurate Localization of Epileptic Seizure Onset Zones. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *26*(12), 2280–2289. <https://doi.org/10.1109/TNSRE.2018.2877820>
- Zelmann, R., Lina, J. M., Schulze-Bonhage, A., Gotman, J., & Jacobs, J. (2014). Scalp EEG is not a Blur: It Can See High Frequency Oscillations Although Their Generators are Small. *Brain Topography*, *27*(5), 683–704. <https://doi.org/10.1007/s10548-013-0321-y>
- Zelmann, R., Mari, F., Jacobs, J., Zijlmans, M., Chander, R., & Gotman, J. (2010). Automatic detector of High Frequency Oscillations for human recordings with macroelectrodes. *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC'10*, 2329–2333. <https://doi.org/10.1109/IEMBS.2010.5627464>
- Zelmann, R., Mari, F., Jacobs, J., Zijlmans, M., Dubeau, F., & Gotman, J. (2012). A comparison between detectors of high frequency oscillations. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *123*(1), 106–116. <https://doi.org/10.1016/j.clinph.2011.06.006>
- Zijlmans, M, Jacobs, J., Zelmann, R., Dubeau, F., & Gotman, J. (2009). High-frequency oscillations mirror disease activity in patients with epilepsy. *Neurology*, *72*(11), 979 LP – 986. <https://doi.org/10.1212/01.wnl.0000344402.20334.81>
- Zijlmans, Maeike, Jiruska, P., Zelmann, R., Leijten, F. S. S., Jefferys, J. G. R., & Gotman, J. (2012). High-frequency oscillations as a new biomarker in epilepsy. *Annals of Neurology*, *71*(2), 169–178. <https://doi.org/10.1002/ana.22548>