Temporal Characteristics of High-Frequency Oscillations as a Biomarker of Human Epilepsy

by

Jared M. Scott

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biomedical Engineering)
at The University of Michigan
2021

Doctoral Committee:

Associate Professor William C. Stacey, Chair
Associate Professor Timothy Bruns
Professor Alfred Hero III
Professor James Weiland

Jared M. Scott

jaredmsc@umich.edu

ORCID iD: 0000-0001-9656-916X

Dedication

This dissertation is dedicated to my wife's Aunt Laurie, whose daughter suffers from a rare genetic form of epilepsy. I have seen firsthand the toll that epilepsy can exact, not only on those who suffer from regular seizures, but also the effect of these seizures on the family and friends who selflessly care for their loved ones day in and day out. Thinking and reflecting about Aunt Laurie in this way often motivated and inspired me when I encountered a tough problem through the course of my thesis. Though it is not enough, I hope it is some comfort to know that a passionate group of medical researchers across the globe are committed to ending the disease of epilepsy, and I am fortunate enough to have played a part in such a movement, however small my role has been.

Acknowledgments

There are many people in my life that have made the work of this thesis possible. First, I would like to thank Dr. William Stacey, my PhD advisor, for his insight and patience, and for bringing me into an engaging and creative lab. The support of my family and friends has also been crucial; they have been with me every step of the way, believing in me and encouraging me, even when the path forward was long. But most importantly, none of this work would have been possible without the love, support, and eternal patience of my amazing wife Kelley, who moved here with me just so that I could pursue a dream. To her I am forever grateful.

There was also significant institutional support that enabled this work, including that of the University of Michigan, Rackham Graduate School, the Biointerfaces Institute as well as the department of Biomedical Engineering. This support also includes several grants from the National Institutes of Health, the Robbins Family Research Fund, and Lucas Family Research Fund.

Table of Contents

List of Tables

List of Figures

Abstract

Epilepsy is a debilitating neurological disorder characterized by recurrent spontaneous seizures. While seizures themselves adversely affect physiological function for short time periods relative to normal brain states, their cumulative impact can significantly decrease patient quality of life in myriad ways. For many, anti-epileptic drugs are effective first-line therapies. One third of all patients do not respond to chemical intervention, however, and require invasive resective surgery to remove epileptic tissue. While this is still the most effective last-line treatment, many patients with 'refractory' epilepsy still experience seizures afterward, while some are not even surgical candidates. Thus, a significant portion of patients lack further recourse to manage their seizures – which additionally impacts their quality of life.

High-frequency oscillations (HFOs) are a recently discovered electrical biomarker with significant clinical potential in refractory human epilepsy. As a spatial biomarker, HFOs occur more frequently in epileptic tissue, and surgical removal of areas with high HFO rates can result in improved outcomes. There is also limited preliminary evidence that HFOs change prior to seizures, though it is currently unknown if HFOs function as *temporal* biomarkers of epilepsy and imminent seizure onset. No such temporal biomarker has ever been identified, though if it were to exist, it could be exploited in online seizure prediction algorithms. If these algorithms were clinically implemented in implantable neuromodulatory devices, improvements to quality of life for refractory

epilepsy patients might be possible. Thus, the overall aim of this work is to investigate HFOs as potential temporal biomarkers of seizures and epilepsy, and further to determine whether their time-varying properties can be exploited in seizure prediction.

In the first study we explore population-level evidence for the existence of this temporal effect in a large clinical cohort with refractory epilepsy. Using sophisticated automated HFO detection and big-data processing techniques, a continuous measure of HFO rates was developed to explore gradual changes in HFO rates prior to seizures, which were analyzed in aggregate to assess their stereotypical response. These methods resulted in the identification of a subset of patients in whom HFOs from epileptic tissue gradually increased before seizures.

In the second study, we use machine learning techniques to investigate temporal changes in HFO rates within individuals, and to assess their potential usefulness in patient-specific seizure prediction. Here, we identified a subset of patients whose predictive models sufficiently differentiated the preictal (before seizure) state better than random chance.

In the third study, we extend our prediction framework to include the signal properties of HFOs. We explore their ability to improve the identification of preictal periods, and additionally translate their predictive models into a proof-of-concept seizure warning system. For some patients, positive results from this demonstration show that seizure prediction using HFOs could be possible.

These studies overall provide convincing evidence that HFOs can change in measurable ways prior to seizure start. While this effect was not significant in some individuals, for many it enabled seizures to be predicted above random chance. Due to data limitations in overall recording duration and number of seizures captured, these findings require further validation with much larger high-density intracranial EEG datasets. Still, they provide a preliminary framework for the eventual use of HFOs in patient-specific seizure prediction with the potential to improve the lives of those with refractory epilepsy.

Chapter I: Introduction

During the transition toward sleep each night, the physiological rhythms of the human body gradually slow until the body is mostly motionless. The brain does not follow this trend, however. After wakefulness departs, neural activity in the brain is renewed and even intensifies: this is evident first in the waves of high amplitude synchronous activity of slow wave sleep, and then in the dramatic flourish of dissociated activity known as rapid eye movement sleep (REM sleep) [1]. To a parent observing a child dreaming under REM sleep, this veritable symphony in the brain might not be physically evident, except perhaps for the light flicker of the eyes beneath their lids. The arrival of such a moment for a relieved parent, however, is often the result of a practiced nightly routine that might involve pajamas and story books. Indeed for anyone, this nightly lapse in consciousness – and the act of sleep in general – requires some amount of forethought and preparation to ensure the body's safety while the mind drifts between dream states, blissfully unaware.

At first glance, it might seem that the serene sleep of a child would have nothing in common with the pronounced convulsions of someone experiencing an epileptic seizure. If the two are compared in a more abstract sense, however, there are intriguing parallels between the phenomena of sleep and seizures. Both sleep and seizures are characterized by significant transitions of bodily state that exert powerful control over cardiovascular and musculoskeletal functioning [2], [3]. In the brain, these state

changes also result in distinct patterns of neural activity that display periods of increased synchrony [4]. Crucially, sleep and seizures are both capable of significantly altering individual awareness and level of consciousness [4]. Despite the apparent similarities between sleep and seizures, however, they have diametrically opposed effects on the body: sleep generally rebuilds and restores vital physiological processes, but seizures instead disrupt them and even damage them over many occurrences. Spontaneous seizures that occur repeatedly – which are a defining characteristic of the disease of epilepsy – are associated with increased risk of injury and even death [5].

In addition to negative physiological effects, the unseen psychosocial costs of recurrent seizures are significant [6]. Seizures strike suddenly and without warning, which can leave unconscious or recovering individuals vulnerable to their surroundings, and at the mercy of passersby to deliver aid. This is in stark contrast to sleep and the nightly bedtime rituals of parent and child described above. Thus, seizures represent an *unplanned* and pathological gap in consciousness that cannot be sufficiently prepared for. For those with recurrent seizures and the disease of epilepsy, this looming and seemingly random threat is a significant detractor to overall wellness and quality of life [7]. Epilepsy does this in a variety of ways, from the disease's persistent social stigma [8], to its negative impact on individual financial health and education possibility [9], and by something as simple as not being able to drive a car [7]. Psychosocial effects of epilepsy also extend significantly to family members, caretakers, and relatives [10].

Still, the chief complaint of those with uncontrolled epilepsy is the random nature of seizures [11]. If it were hypothetically possible to know the exact timing of an individual's

seizures in any given day, life's activities could be planned or rerouted with more ease, and more importantly, patient safety could be addressed by preparing for the seizures themselves. While seizures likely do not occur in this rigidly deterministic fashion [12], it is likely that even limited foresight of future seizure timing could significantly improve patient quality of life. This idea is at the center of this work's ultimate end goal: to eventually address patient quality of life by investigating and developing a novel biomarker of seizure timing in the context of seizure prediction and probabilistic forecasting.

Background and motivation

Seizures: A fundamental property of neuronal networks is their ability to exhibit seizure-like activity [13]; this is found in simple organisms like fruit flies [14] and zebrafish [15], and it is also conserved in more complex and developed organisms like mammals [13]. Shown in the intracranial EEG waveforms in Figure 1 below, a seizure itself is the paroxysmal firing of large neuronal populations in synchrony [3]. Seizures manifest physically in myriad ways [16], and their severity ranges from brief periods of impaired awareness to complete loss of consciousness with accompanying rhythmic convulsions of limbs and body. This variety is in part because a seizure is not in and of

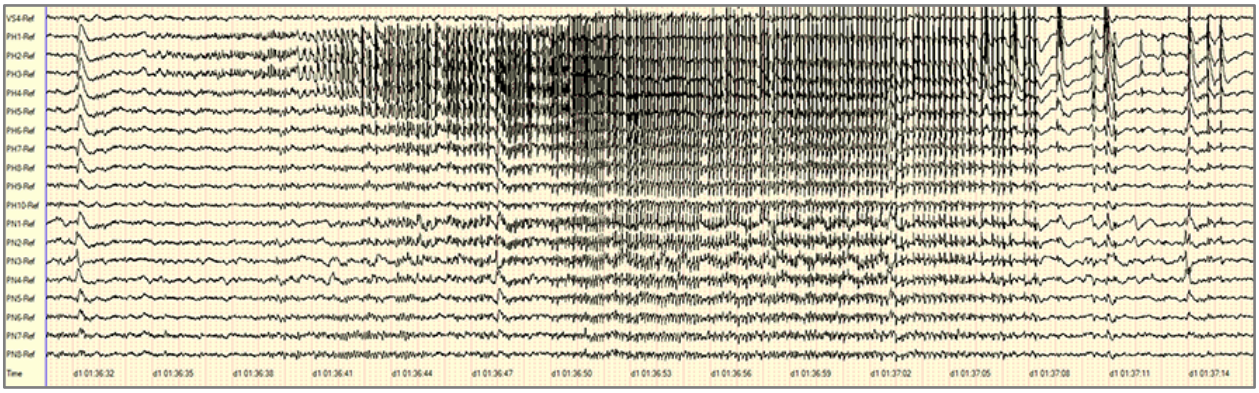

FIGURE 1: Example seizure waveform. This seizure originated in data recorded from patient UMHS-0026. Overall the seizure lasts approximately 30 seconds. It begins with a large spike and DC shift in many channels; subsequently it evolves and increases in amplitude until reaching a maximum at seizure offset, where several bursts of spikes follow.

itself a disease, but rather an epiphenomenon of underlying neuronal pathology [3]. Approximately 1 in 10 people will have at least one unprovoked seizure in their lifetime [17], [18], and the probability of having additional seizures increases every time one is experienced [19].

Epilepsy: The disease of epilepsy is defined as the recurrence of spontaneous seizures, and it is one of the most prevalent neurological disorders in the world, affecting 60 million total [3], [17], [18]. It is also a highly heterogeneous disease with many different types of syndromes, each of which can result from a variety of etiologies [20], [21]. After diagnosis – which is conducted with a thorough review of patient history and commonly includes scalp electroencephalogram monitoring (EEG) to identify the presence of interictal epileptic discharges (IEDs) – affected patients are typically put on a regimen of anti-epileptic drugs, which generally function by calming aberrant electrophysiological processes in the brain [22].

Treatments for refractory epilepsy: Approximately one third of individuals with epilepsy do not respond to anti-epileptic drugs (AEDs), however [23]. These patients

with ‘refractory epilepsy’ typically move through a variety of AED combinations before their treating clinicians arrive at the conclusion that other more invasive treatments are necessary. There are two last-line therapies available in this regard: 1) resective surgery, which removes epileptic tissue thought to be generating seizures, or 2) implantable neuromodulatory devices, which deliver electrical impulses either in response to certain electrical patterns (referred to as closed-loop or responsive stimulation [24]), or in a continuous fashion using a predefined stimulus (open-loop stimulation).

Unmet clinical need in refractory epilepsy: Resective surgery remains the gold-standard treatment for refractory epilepsy, as the effectiveness of devices still falls below [25]–[28]. After one year, however, only 60% of resected patients are seizure-free [29]. Further, many patients with refractory epilepsy are not surgical candidates either because their seizures originate from multiple foci, or because the seizure-generating tissue overlaps with important and vital brain areas related to speech, movement, or sight. Thus, there is significant unmet clinical need in the sizable portion of refractory epilepsy patients still experiencing seizures.

High-frequency oscillations: High-frequency oscillations (HFOs) are a more recently discovered electrical biomarker of epilepsy that have begun to address this clinical need. HFOs are brief (< 50 ms) and somewhat rare neuronal events that occur in frequencies from 80 - 500 Hz [30]; there is also evidence that they occur in higher frequencies as well [31]. Though they can be associated with normal physiological processes involving memory and vision [32], [33], brain tissue with comparatively high

HFO rates is associated with epileptogenicity [21], [34], [35]. Importantly, surgical removal of high HFO tissue can result in better clinical outcomes and increased seizure freedom [36]–[40].

Using microwire electrodes, pathological HFOs were originally discovered approximately 20 years ago in humans and in animal models of mesial temporal lobe epilepsy (MTLE) [41]. It was quickly confirmed that HFOs are recordable with standard intracranial grid, and depth electrodes [42]; recently, they have also been identified with standard scalp EEG as well [43]. In the literature, HFOs are commonly subdivided into 'ripples' (80 - 200Hz) and 'fast ripples' (250 - 500 Hz). Initially it was thought that fast ripples had more of a pathological association [44], but recent evidence questions this assumption [45]–[47] To date there are no reliable methods to separate physiological and pathological HFOs [48]. The exact biological mechanism for HFO generation is unknown, but both epileptic ripples and fast ripples are currently thought to be the result of synchronous firing by principal pyramidal cells [49].

Recording HFOs: There have been significant advances in the technology used to record, and process HFOs since they were first discovered [42], [50]. Initially, clinical EEG systems could not adequately address the difficulties involved with recording HFOs. In order to accurately discern an HFO waveform, a sampling rate of at least 2 kHz is required [42], [50], but they are more easily identified as sampling rate increases [51]; until somewhat recently, most clinical systems were limited to sampling frequencies of 256 Hz. HFOs are also low amplitude events compared to their surrounding backgrounds; this signal-to-noise ratio was also a challenge that more

modern amplifiers had to overcome. But given the increasing demand for HFO recordings, clinical EEG manufacturers have responded with the development of advanced amplifiers that have low noise floors and high-sampling frequencies. [34]

Automated identification of HFOs: HFOs were originally identified manually in short 10-minute clips of data by visual inspection of raw and filtered waveforms [52], [53]. As recordings grew in duration and number, however, it was clear that this labor-intensive method was prohibitive of more advanced analyses with more patients. Using visual inspection as validation, a number of automated HFO detectors were developed (see [54] for a review) that could parse a huge amount of intracranial data in a small fraction of the amount of time it would take a human reviewer. These detectors have allowed the proliferation of numerous HFO studies in the literature today; many of these have contributed to the generally accepted understanding that HFOs are spatial biomarkers of epileptic tissue.

Thesis motivation – HFOs. While the spatial aspect of HFOs somewhat predominates the literature, other facets and research questions have gone unexplored that could still show significant clinical potential. Notably, there are few studies that have explored how HFOs change over time, and fewer still that have investigated such changes prior to seizure onset. The only two studies in this regard [55], [56] were limited by small patient cohorts and few seizures per patient. Still, both identified significant changes in preictal (meaning prior to seizure) HFOs in some patients. However, they came to opposing conclusions as to whether these changes were stereotyped, and additionally whether such preictal changes were truly different from changes in HFOs

observed during interictal periods (meaning between seizures). Thus, it is largely unknown if HFOs are temporal biomarkers of seizure onset and whether they can be used in seizure prediction or forecasting. This open question is the primary motivation for the novel investigation of HFOs presented in this thesis.

Seizure prediction: There is tremendous clinical potential in the pursuit of accurate seizure prediction. As such, the prediction of epileptic seizures has been the focus of much research over the past two decades. The first prediction studies used large numbers of EEG-derived features in small data sets to prove that prediction was possible [57]. These studies lacked statistical validation that their prediction algorithms performed better than a random chance predictor, however. After this misstep, a set of statistical requirements and associated methods were adopted by the field to add rigor and consistency to future prediction studies [58]–[60]. These methods include seizure time surrogates, where labels on seizure and non-seizure periods are randomly permuted [61]; comparisons using a random Poisson predictor [62], and the AUC metric, which has since been used in many studies [63]–[65] and is used throughout this thesis.

Neurovista and prospective prediction: In 2013, the defining achievement in seizure prediction came from the development of an implantable intracranial monitoring device, called the 'Neurovista' device [66]. It resulted in chronic recordings of high-quality ambulatory iEEG for 13 patients; in some patients these recordings spanned months and close to years. In addition to recording intracranial data, the device also functioned as a seizure forecasting system, and would alert the wearer to their current

seizure risk. This aspect of the device used signal features derived from incoming intracranial EEG data as the input to a true prospective prediction algorithm – which, to this day, stands as the only such instance of its kind. The iEEG data acquired during this study has since been used in numerous prediction studies [64], [65], [67]–[71]; some of those algorithms evaluated seizure risk retrospectively by evaluating all data at once, and others used pseudo-prospective approaches to evaluate seizure risk moving forward continuously in time, mimicking the original Neurovista method.

Thesis motivation – seizure prediction. Along with our evolving understanding of epilepsy and seizures, the knowledge gained from early and more recent seizure prediction studies is shaping the field's future. A recent review [12] of seizure prediction written by seminal experts in the field detailed a number of strategies and research goals to be pursued in future work; these included the following: 1) further development of pre-seizure electrical biomarkers, 2) the pursuit of patient- and possibly even seizure-specific prediction algorithms, and 3) the reformulation of seizure prediction into a probabilistic rather than deterministic framework. By exploring HFOs in the context of seizure prediction and probabilistic forecasting, the work of this thesis is directly motivated by all three items above: given the previously described variability of individual outcomes for HFO studies involving temporal changes in the preictal period [55], [56], it is possible that HFOs are a patient-specific biomarker of seizure onset, capable of use in seizure prediction. Additionally, recent evidence has shown that HFOs occurring at the beginning of a seizure can differentiate between two different seizure types [72], a finding which is relevant to seizure-specific prediction using HFOs.

Thesis aims and organization

With the ultimate goal of improving quality of life for patients with refractory epilepsy, this thesis addresses two general aims: 1) to evaluate HFOs as a temporal biomarker of seizure onset and epilepsy, and 2) to assess the relevance of these findings to seizure prediction. These aims are addressed across the findings of three independent studies.

Using a big-data framework that was developed in concert with state-of-the-art HFO identification and processing methods, the first study in Chapter II investigates temporal properties of preictal HFO rates at the population level to identify the existence of this effect, and its potential prevalence within a large clinical cohort. We develop a novel continuous measure of HFO rate (cHFO), and use this to analyze preictal and interictal time periods together to identify stereotypical patterns – first in individuals, then compared at the population level – that differentiate pre-seizure data.

The second study in Chapter III addresses HFOs as a potential patient-specific biomarker of imminent seizure onset. We characterize fluctuations in the distribution of HFO rates through time across several sliding windows of different duration. These data are used to train logistic regression models – a framework chosen for its probabilistic output – to identify differences in preictal versus interictal HFOs. The predictive performance of these models with unseen held-out data is cross-validated and then assessed with the AUC metric, which affords an overall idea of how well the predictive classifiers can differentiate preictal from interictal periods.

The third study in Chapter IV is a significant expansion of the techniques and analyses presented in Chapter III. It extends the data used for prediction and forecasting to include information about HFO signal features. The predictive performance of these models is then compared with the performance of models in Chapter III to assess their increased utility in seizure prediction. Finally, this study demonstrates the practical potential of HFOs in seizure prediction and forecasting by providing an implementation of a seizure-warning system. This study is meant to show that HFOs could help forecast oncoming seizures, and motivates the use of high resolution EEG in future devices.

Overall this thesis presents several significant contributions to HFO research and seizure prediction in epilepsy. The techniques and methodologies for HFO identification and processing used in all three studies represent significant advances in the use of HFOs in a big data framework – particularly for HFO data around seizures, whose accuracy had not sufficiently been addressed. The use of this framework with the HFO data of a large clinical cohort provides further validation of these results. The use of HFOs in seizure prediction is a novel idea that, prior to this thesis, has not been investigated in the literature. It is hoped that these findings can serve as preliminary waypoints to further research on the temporal aspects of HFOs, especially in the context of seizure prediction.

Chapter II: Preictal Variability of High-Frequency Oscillation Rates in Refractory Epilepsy

J. M. Scott, S. Ren, S. V. Gliske, and W. C. Stacey, "Preictal variability of high-frequency oscillation rates in refractory epilepsy," *Epilepsia*, 2020.

Abstract

Objective: High-frequency oscillations (HFOs) have shown promising utility in the spatial localization of the seizure onset zone for patients with focal refractory epilepsy. Comparatively few studies have addressed potential temporal variations in HFOs, or their role in the preictal period. Here, we introduce a novel evaluation of the instantaneous HFO rate through interictal and peri-ictal epochs to assess their usefulness in identifying imminent seizure onset. **Methods:** Utilizing an automated HFO detector, we analyzed intracranial electroencephalographic data from 30 patients with refractory epilepsy undergoing long-term presurgical evaluation. We evaluated HFO rates both as a 30-minute average and as a continuous function of time and used nonparametric statistical methods to compare individual and population-level differences in rate during peri-ictal and interictal periods. **Results:** Mean HFO rate was significantly higher for all epochs in seizure onset zone channels versus other channels. Across the 30 patients of our cohort, we found no statistically significant differences in mean HFO rate during preictal and interictal epochs. For continuous HFO rates in seizure onset zone channels, however, we found significant population-wide increases

in preictal trends relative to interictal periods. Using a data-driven analysis, we identified a subset of 11 patients in whom either preictal HFO rates or their continuous trends were significantly increased relative to those of interictal baseline and the rest of the population. **Significance:** These results corroborate existing findings that HFO rates within epileptic tissue are higher during interictal periods. We show this finding is also present in preictal, ictal, and postictal data, and identify a novel biomarker of preictal state: an upward trend in HFO rate leading into seizures in some patients. Overall, our findings provide preliminary evidence that HFOs can function as a temporal biomarker of seizure onset.

Introduction

High-frequency oscillations (HFOs) have shown promise in clinical epilepsy research as a biomarker of epileptic tissue. Defined as short bursts of neural activity > 80 Hz, HFOs occur more frequently in epileptic tissue [30], [35] Numerous studies have shown that HFOs accurately delineate the seizure onset zone and potentially improve surgical outcomes [36]–[40]. Although most HFO studies concentrate on localization of abnormal channels, there is interest in characterizing other aspects of HFOs and epilepsy [43]. As high-frequency activity has been shown to increase prior to seizure onset both clinically and in experimental models, [73]–[75] some have also hypothesized a link between HFOs, the mechanisms of ictogenesis, and preictal brain states [49], [72], [74]–[80].

The existence of a preictal state is still unproven, but growing evidence suggests it is measurable in many patients [58], [81], [82]. One notable study found differences in preictal electroencephalogram (EEG) occurring even hours before seizure onset [82]. However, very few studies address HFOs in the preictal period. Early work with small cohorts showed that preictal HFOs have subtle changes in the preictal period, such as spectral and rate changes [55] or alterations in HFO features [56]. Newer hardware and software now make HFO research much more robust, allowing high-quality, larger datasets [56], [77], [83]–[86]; the role of HFOs as a preictal biomarker can now be answered with much higher rigor. To our knowledge, there is no study of peri-ictal HFO rates using modern equipment and algorithms to acquire a robust sample size. This has halted further progress toward our understanding of the temporal evolution of HFOs and their relationship to mechanisms of seizure generation and termination. Furthermore, it has prevented the adoption of HFOs as a temporal biomarker.

We designed this study to directly address these deficits. Here, we analyze >11 million automatically detected HFOs from the entire intracranial EEG record of 30 patients. We adapt the analysis to generate the first robust comparison of peri- and interictal HFO rates. We find a subset of patients in whom HFO rates change up to 30 minutes prior to seizures, which we suggest can be used as a temporal biomarker of impending seizure onset in future seizure prediction applications.

Methods

Patient population: Data were acquired from all consecutive patients at the University of Michigan who had intracranial EEG (iEEG) monitoring for refractory epilepsy with at least 4096 Hz sampling rate from 2016 to 2018. For inclusion in the study, patients had to have a total recorded time of at least 24 hours, during which at least 1 seizure occurred. Additionally, we required sufficient metadata regarding channel mappings, seizure times, and other clinical data. This produced a total of 30 patients for the study. Electrodes implanted for monitoring included a mix of subdural grids, conventional depth electrodes, and stereo-EEG electrodes. Channels were labeled as "seizure onset zone," and seizure onset/offset times were determined, according to the official clinical report of the treating epileptologist. Channels were labeled as lying within "resected volume" by consultation with the neurosurgeon and comparison of pre- and postoperative imaging (when available). Prior to data acquisition, full institutional review board approval was obtained, as well as written consent from patients to share their deidentified data. All EEG data were acquired with a Quantum amplifier (Natus Medical) with a sampling rate of 4096 Hz. Further summary of the patient population can be found in Table 1.

Data processing and analysis: All data were analyzed using custom C++ and MATLAB (MathWorks) packages and scripts. As seen in Figure 2, our data analysis workflow consisted of three main components: automated HFO detection, indexing and windowing operations, and statistical analysis of mean and continuous HFO rates. These individual steps are described below.

TABLE 1: Clinical data for first study. Abbreviations: M/F: male, female, L/R: left / right, T: temporal, P: parietal, F: frontal, Occ: occipital, DNET: dysembryoplastic neuroepithelial tumor, NF1: neurofibromatosis type 1 tumor, NR: not resected, CD: cortical dysplasia, MTS: medial temporal sclerosis, PVNH: periventricular nodular heterotopia, PMG: polymicrogyria, VNS: vagal nerve stimulator.

Subject	Age	Sex	ILAE outcome	Seizure focus (hemisphere, region)	Pathology / *implant type*	Number of intracranial channels				Total recorded time	Total number HFOs	Mean HFO rate (#/min./channel)		HFO mean frequency (median, Hz)		Number of recorded seizures		Responder subset membership		
						total	ECoG	depth	SOZ	(hours)		SOZ	OUT	SOZ	OUT	total	used	mean rate	slope SOZ	slope OUT
UMHS-0019	59	F	II	R T	Gliosis	106	106	0	2	168.8	400,123	1.99	0.17	156.2	161.6	5	1	X	X	
UMHS-0020	45	F	II	R T	MTS	25	0	25	9	171.2	55,311	0.36	0.13	172.7	221.8	7	7			
UMHS-0021	30	M	II	R T	Gliosis, PVNH, PMG	46	0	46	13	179.5	459,037	1.91	0.47	169.9	166.8	9	6			
UMHS-0022	40	M	I	L T	CD, MTS	38	0	38	3	160.8	72,486	1.38	0.06	190.0	182.7	8	5		X	
UMHS-0023	29	M	NR	L T, P	PVNH / *Neuropace*	69	41	28	29	164.3	354,931	0.83	0.34	157.0	166.4	20	9			
UMHS-0024	31	M	NR	L, R T	*Neuropace*	75	55	20	16	177.2	1,124,176	2.62	1.24	152.1	154.6	28	11			
UMHS-0025	17	F	II	L T	Gliosis	20	0	20	5	207.7	269,638	1.77	0.88	161.6	172.8	10	3			X
UMHS-0026	22	F	NR	R T	PVNH	52	0	52	3	246.2	390,187	1.52	0.51	165.3	166.3	40	7	X	X	X
UMHS-0027	26	M	NR	L Diffuse	*VNS*	91	81	10	3	205.2	1,212,921	2.98	2.19	148.3	154.0	97	8			X
UMHS-0028	14	F	I	R T	Tumor: Glioma	53	47	6	5	79.7	198,968	2.39	0.37	154.3	159.2	7	4	X		
UMHS-0029	48	M	NR	L T, Occ.	*Neuropace*	91	91	0	22	226.3	819,880	0.61	0.72	159.3	168.1	14	7			
UMHS-0030	5	M	III	L T	MTS, Gliosis	100	100	0	2	146	378,824	1.01	0.56	152.3	169.0	33	12			
UMHS-0031	13	M	I	L T	Gliosis, Tumor: NF1	99	99	0	6	180	371,855	0.75	0.24	150.4	159.4	9	6			
UMHS-0032	41	F	I	R F	CD	32	0	32	3	184.3	382,400	2.45	0.64	159.4	170.5	8	6	X	X	X
UMHS-0033	5	F	II	R Ins.	CD, Gliosis	74	0	74	4	120.7	150,963	0.97	0.30	169.8	219.7	28	19			
UMHS-0034	33	F	I	R F	Gliosis	32	0	32	11	136.3	455,089	2.41	1.18	172.2	167.3	17	16			
UMHS-0035	50	F	I	L T	Gliosis	57	57	0	2	162.7	122,451	0.67	0.19	147.9	172.4	7	6			
UMHS-0036	43	M	NR	L, R T	CD / *Neuropace*	54	0	54	2	172.5	335,274	1.36	0.60	151.8	163.6	18	12			
UMHS-0037	14	M	I	L F	Tumor: DNET	50	0	50	-	219.7	229,207	-	0.30	-	157.3	34	22			
UMHS-0038	28	M	II	L T	MTS, Gliosis	61	61	0	-	178.7	746,718	-	1.16	-	156.5	7	2			
UMHS-0039	47	M	NR	R P	CD / *Neuropace*	90	0	90	10	155.2	233,050	0.99	0.22	160.6	184.0	19	7			
UMHS-0040	14	F	I	L P	CD, Gliosis	63	55	8	8	196.7	386,462	0.37	0.64	158.7	170.1	7	7			X
UMHS-0041	32	F	I	R F	CD	71	0	71	9	176.5	73,589	0.30	0.04	166.7	191.0	36	3			
UMHS-0043	28	M	II	R T	Gliosis	86	0	86	9	182.2	279,124	0.75	0.33	170.9	226.8	46	5	X		
UMHS-0044	45	F	NR	L T, P	*Neuropace*	76	0	76	6	170.2	385,032	1.24	0.45	155.4	179.6	13	4			
UMHS-0045	17	F	NR	L, R T	*Neuropace*	94	0	94	15	331.5	645,420	0.76	0.24	167.3	185.8	6	6			
UMHS-0046	23	F	I	L F	CD	30	0	30	9	139.3	16,061	0.12	0.03	166.1	210.8	17	8			
UMHS-0047	48	F	II	R T	Gliosis	70	0	70	3	301.7	417,307	0.65	0.22	155.0	196.8	1	1	X		
UMHS-0048	22	F	NR	L, R T	*Neuropace*	86	0	86	8	141.8	271,327	2.29	0.25	164.6	178.0	23	3	X		
UMHS-0049	53	F	NR	L, R T	*Neuropace*	94	0	94	15	176.8	179,259	0.63	0.11	179.6	166.9	17	4			
					TOTALS / averages	1985	793	1192	232	5459.5	11,417,070	1.29	0.49	162.0	178.1	591	217	7	4	5

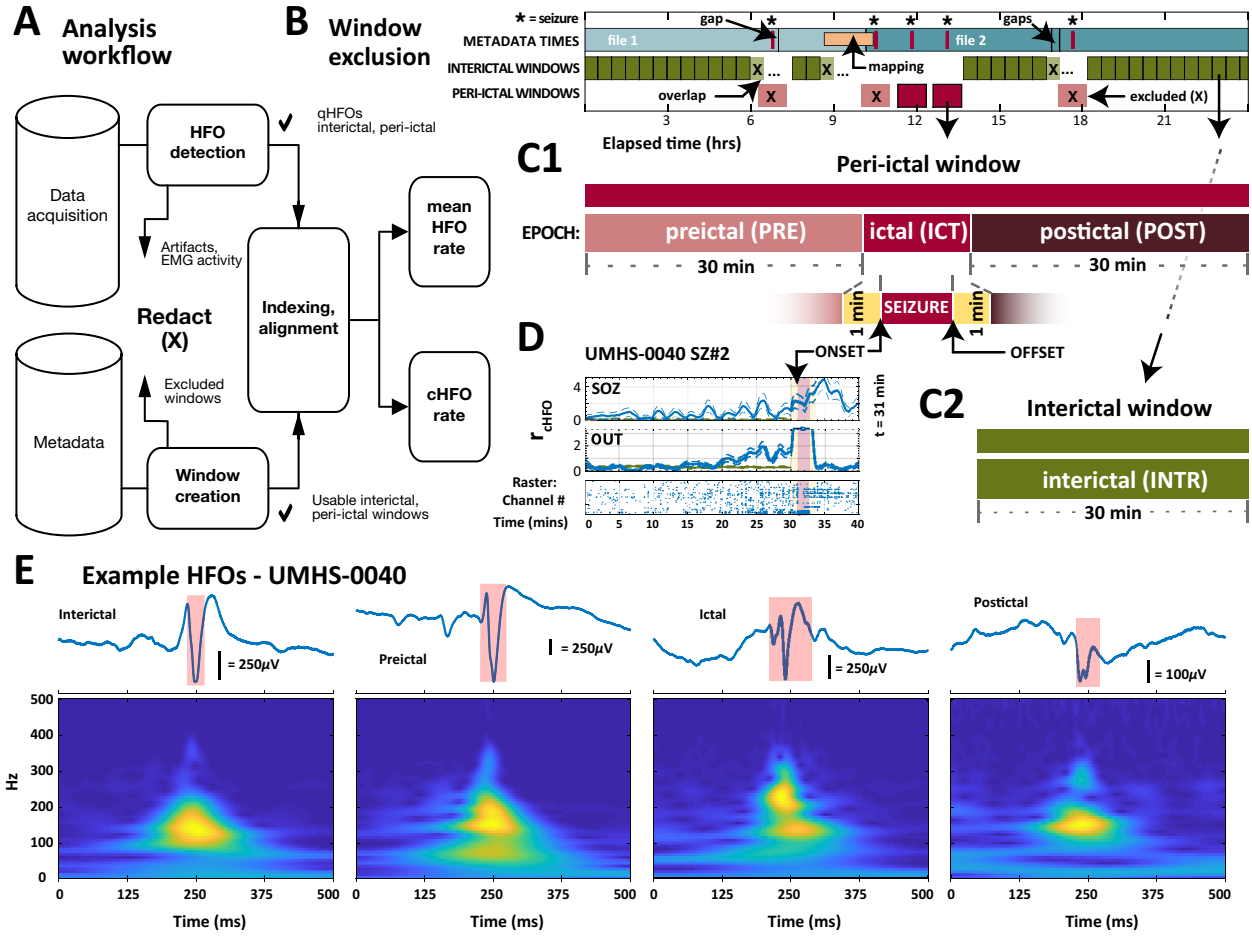

FIGURE 2: Schematic diagram showing overall data analysis workflow. A, Quality high-frequency oscillation (HFO) detections (quality HFOs [qHFOs]) and their respective interictal and peri-ictal windows of analysis are aligned in time to compute mean and continuous HFO rate. EMG, electromyographic. B, Analysis windows are created from patient metadata and excluded from further analysis if overlap occurs with a number of conditions that would bias results. C1, Remaining peri-ictal windows are further divided into preictal, ictal (which includes a 1-minute buffer on either side of the clinically marked seizure time), and postictal epochs. C2, Remaining interictal windows are defined as 30-minute epochs. D, Continuous HFO rate (cHFO) computed from a single seizure in an individual patient is shown for seizure onset zone channels (top row, SOZ) and nonepileptic channels (middle row, OUT). cHFO rates were computed from discrete HFO detections, shown as a raster plot of preictal detections (bottom row) and organized by channel index. This patient (UMHS-0040) was a member of the "slope responder" subset of patients and showed preictal increases in cHFO rate as onset approached. Here cHFO rate is defined as HFOs per minute per channel. Dotted lines indicate ±1 standard deviation; blue denotes preictal cHFO rate, and green denotes interictal cHFO rate for comparison. The peri-ictal window was truncated for display purposes at 40 minutes. E, Example HFO detections for the same patient in interictal, preictal, ictal, and postictal periods are visualized in time-frequency plots, each computed with the Morse wavelet.

Automatic HFO detection and electromyographic artifact removal: For automated HFO detection, we used a previously validated HFO detector [86]. Briefly summarized, we use the highly sensitive "Staba" detector [52] on band-passed (80-500

Hz) data, then redact detections likely to be due to artifacts, leaving more specific "quality HFOs" (qHFOs). We also applied an additional, published artifact rejection method designed to redact activity associated with scalp muscle artifact, which can produce many false-positive detections in the lateral temporal lobes [85]. All HFOs discussed in this work were subjected to this full process.

Adjusting HFO detector for peri-ictal periods: All resulting HFOs for a given patient were labeled as either interictal baseline or peri-ictal, which we defined to include the full period from 30 minutes prior to 30 minutes after a seizure. Interictal HFOs were indexed into a successive series of interictal windows whose individual duration was 30 minutes. Peri-ictal detections were further subdivided into three continuous epochs: preictal, ictal, and postictal. We defined the preictal and postictal epochs as beginning 30 minutes before and ending 30 minutes after the ictal epoch, respectively. The ictal epoch was defined by the clinical mark of beginning and end, as well as an additional 1-minute buffer before and after the seizure. This buffer was added to mitigate potential inconsistencies in clinically marked seizure times, which can vary between clinicians [87], [88]. A schematic showing the exact timing of these epochs is given in Figure 2C.

Most automated HFO detectors are designed for interictal data, where the EEG baseline is assumed to be relatively stable over time; the HFO detection algorithm compares with the baseline EEG every 10 minutes, which is assumed to be interictal [52]. However, including peri-ictal data presents a new challenge, because a seizure changes the "baseline" significantly and disrupts the threshold for HFO detection. To

address these considerations, we used two simple modifications to our HFO detection process during peri-ictal periods.

The first modification was designed to align the 10-minute windows correctly to ensure ictal data were not present in the preictal epochs. This did not change the method of HFO detection, merely the start and stop times for the preictal epochs. During peri-ictal periods, the baseline was referenced to the start of the seizure, that is, the HFO detector was started 31 minutes prior to each seizure onset, which includes the aforementioned 1-minute buffer. From this point, the detector ran in successive 10-minute segments until reaching the end of the postictal epoch as we have defined it above. Aligning the qHFO detector in this manner ensured that ictal EEG activity did not contaminate the preictal baseline threshold used to identify HFOs. Note that if baseline also increased preictally, this would lead to fewer HFOs being detected during the preictal period. Thus, the results herein are a conservative estimate of preictal HFOs.

Second, we fixed the “baseline” threshold used for ictal and postictal HFO detection to the value of the 10-minute preictal segment just prior to the ictal period. This ensured that ictal and postictal rates were scaled to preictal baseline, rather than ictal activity. This was necessary because ictal data typically have a much higher baseline root mean square value than the preictal portion that precedes it, and our understanding of “increased HFO rates,” as well as the automated detector, is based upon comparison with interictal baseline. This method ensured the ictal and postictal rates would be referenced to the preictal baseline, prior to any ictal activity.

Window exclusion and alignment: Because the peri-ictal and interictal data have different reference points, it is possible that the windows overlap with each other or with periods of unreliable data. To ensure data quality and no overlap, we excluded windows that could be unreliable (Figure 2B). Specifically, we redacted windows that had overlap with any of the following conditions: (1) any other window, (2) file start or stop times, (3) gaps in recorded data of 1 minute or more, and (4) known extraoperative mapping procedures or other similar periods of poor data quality. Windows meeting any of these conditions were labeled unusable and excluded from further analysis. After this procedure, there were 217 seizures available for processing in the 30 patients. Remaining windows were then sorted according to type (i.e., interictal baseline or peri-ictal) and aligned in time, which allowed comparison of HFO times across all windows. Grouping these windows then allowed computation of average HFO rates as described below.

Computing HFO rate: Our analysis utilized two different representations of HFO rate: mean HFO rate and continuous HFO (cHFO) rate. These values were computed across two groups of intracranial channels: seizure onset zone channels (hereafter abbreviated SOZ), and all channels that were outside of both the SOZ and the volume of resected tissue (RV), which we denote OUT. Note that there is usually a great deal of overlap between SOZ and RV, but RV often has many channels that were not in the SOZ, and may not contain all of the SOZ, depending upon clinical circumstances. Mean HFO rate was computed as the average over all usable windows and was defined as

the total number of HFOs divided by the product of the number of channels and total duration of the windows used.

cHFO rate: The Nelson-Aalen hazard rate: A robust analysis of temporal characteristics of HFOs requires information on their rate as a function of time, rather than simply an average over long epochs. We estimated HFO rates as a continuous function of time (cHFO rates), with the nonparametric Nelson-Aalen hazard rate model, and smoothed its output with kernel methods [89]–[91]. In a general sense, the Nelson-Aalen model gives the risk of an event's occurrence as a function of time, which is equivalent to its instantaneous rate [90]. This method has been used to quantify oscillatory activity during sleep [92] as well as the risk of seizures over time [93].

Kernel smoothing methods can translate discrete events into continuous estimates of rate, but they require the selection of a bandwidth parameter, which generally controls how jagged or smooth the estimate appears. We fixed this parameter at 1 minute for all patients, which prevented ictal HFOs from influencing preictal cHFO rates as the kernel window moved forward in time.

We computed cHFO rates with the Nelson-Aalen model in the same general manner as mean HFO rates, with one exception. Instead of using all interictal windows in the Nelson-Aalen computation, we restricted their number to be equal to the number of usable peri-ictal windows, choosing them at random from all usable interictal windows. While this allowed us to characterize interictal cHFO rates with the same temporal scale as peri-ictal cHFO rates, it also meant that interictal cHFO rates were

only calculated from a portion of the available data. To mitigate this, we repeated the calculation 10 times with different random selections and report the average of all 10 as the final estimate.

Final analysis and statistical tests: After determining mean and continuous HFO rates for all patients, we compared interictal and peri-ictal rates across all patients. We assessed patient-wise differences in mean HFO rate across channel groups (SOZ, OUT) and epochs (interictal, preictal) with the Wilcoxon signed rank test, using the appropriate Bonferroni correction. We also used the Kolmogorov-Smirnov test to compare differences in the population distributions of mean HFO rate across channel groups.

The cHFO rate is a continuous variable that estimates the instantaneous rate at every point in time. We first analyzed these results visually and noticed two clear groups of patients: (1) most patients had essentially stable cHFO rates preictally, which were similar to the interictal values; and (2) some patients had preictal cHFO rates that were larger than the interictal values and appeared to increase leading to the seizure. To quantify this difference, we fit a line to preictal and interictal cHFO trajectories using least squares linear regression. We compared slopes of these lines within and across patients with the Wilcoxon signed rank test and further compared their overall distributions for different channel groups with the Kolmogorov-Smirnov test.

For both analyses, we used an unbiased, data-driven approach to identify natural clusters of outliers in the distributions by applying a kernel density estimator to the

population distribution, then identifying local minima that distinguished any anomalous cluster, similar to our previous methods [86]. These minima were used as thresholds to identify putative responders.

Results

Our automated HFO detector was run on the iEEG data of 30 patients (15 male, 15 female) from the University of Michigan health system. Patients in the study represented a diverse clinical cohort with a variety of ages, seizure foci, and epileptic etiologies. In total, > 11.4 million HFOs from nearly 2,000 iEEG channels were detected and analyzed across > 225 days of iEEG data. Further patient summary can be found in Table 1.

Comparison of mean HFO rates: We first compared mean HFO rate across all the temporal epochs, an analysis that previously has been restricted almost exclusively to interictal periods. As shown in numerous prior studies, we found that SOZ channels had significantly higher mean rates than OUT channels for interictal and preictal epochs (Figure 3A, $P < .001$). Similar results occurred in ictal and postictal epochs (not shown, $P < .001$). We also compared mean HFO rates in different epochs across our population (not shown); ictal periods had much higher HFO rates than all other epochs (SOZ, OUT: $P < .001$), whereas postictal rates were quite variable among different patients but on average tended to be slightly higher than either interictal or preictal epochs, although

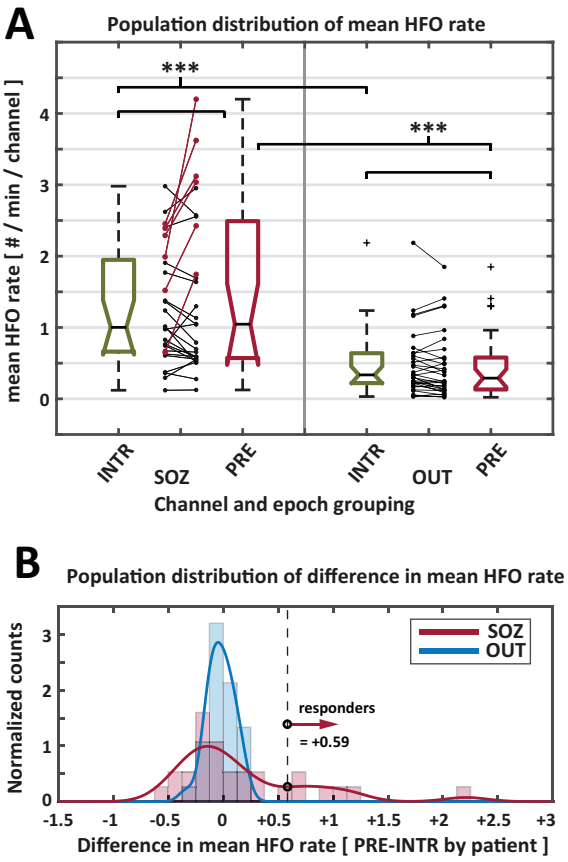

*FIGURE 3: Population comparisons of mean HFO rate. A, Population box plots of mean high-frequency oscillation (HFO) rate comparing interictal (INTR) and preictal (PRE) epochs, organized by channel group (seizure onset zone channels [SOZ], nonepileptic channels [OUT]). No statistical difference in mean HFO rate during interictal and preictal periods was found; mean rate in SOZ channels was significantly higher than OUT channels for all epochs (ictal and postictal, not shown: P < .001). Statistical comparisons performed (Wilcoxon signed rank test) are denoted by brackets at the top of each panel; asterisks show statistical significance, ***P < .001. Differences in raw data during interictal and preictal epochs are visualized per patient between box plot groups: "mean rate responders"—patients with increased difference in preictal rate in SOZ channels—are shown with red lines, whereas other patients are shown with black lines. B, Smoothed and binned population distributions of the difference in preictal versus interictal mean HFO rate are shown by channel group. OUT channels (blue) are unimodal, but SOZ channels are bimodal and show the presence of a "mean rate responder" patient subset (red), each having a difference in rate of 0.58 HFOs/min/channel .*

this did not reach significance in all groups (data not shown). The primary analysis was to compare inter- and preictal HFO rates. When averaged across all patients, there was no statistically significant difference in mean HFO rate between interictal and preictal epochs for either SOZ or OUT channel groups. In certain patients, however, we noticed that preictal rates were significantly higher than their interictal values, especially in the

SOZ. This led to the possibility that specific patients might have large differences between inter- and preictal HFO rates that are not seen when averaged across all patients. We plotted the distribution among all patients of the difference between preictal and interictal rates for both channel groups. As shown in the histograms of Figure 3B, the distribution for OUT channels is centered at zero and is unimodal. In contrast, the distribution for SOZ channels appears significantly skewed to the right, with several patients comprising the right tail of the distribution. This suggested that a distinct subset of "responder" patients in our cohort had significant in- creases in preictal HFO rates in the SOZ. Although these patients were too few to allow statistical tests to find strong independence of the SOZ and OUT distributions (Kolmogorov-Smirnov test, $P = .072$), they are clearly outliers in the SOZ distribution. The threshold to identify these outliers (first local minimum in the distribution of SOZ channels) was 0.58 HFOs/min/channel, yielding seven total "mean rate responders"—individuals for whom the difference in mean HFO rate for preictal and interictal epochs was much higher than the rest of the population. Patients who are within this subset are marked in Table 1 and labeled red in Figure 3A.

Comparison of continuous HFO rates: We used the Nelson-Aalen hazard rate model to estimate HFO rate as a continuous function of time (cHFO rates). The result of this analysis for a single patient is shown in Figure 4, which superimposes the interictal and preictal cHFO rates for visual comparison. Calculating the cHFO rate creates a time-dependent function, which we evaluated mathematically (see next section). We first made visual observations of these functions, comparing the cHFO trajectories

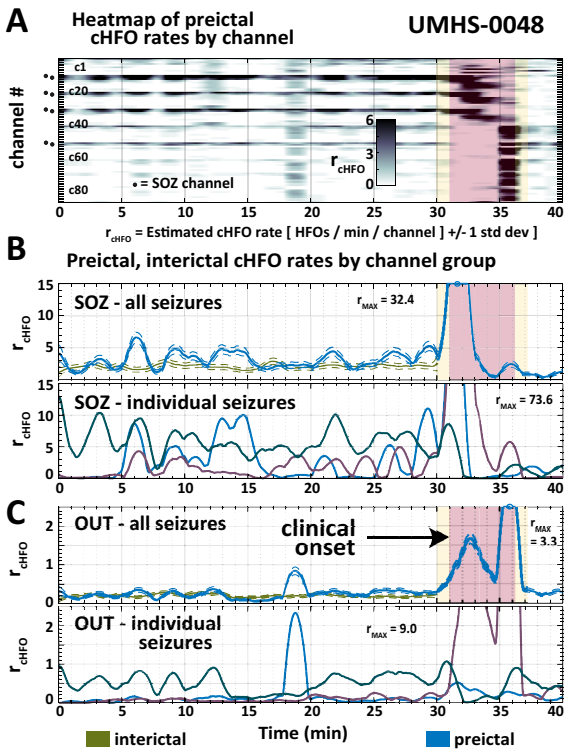

FIGURE 4: Example cHFO rate analysis for single patient. Example of continuous high-frequency oscillation (cHFO) rate analysis (Nelson-Aalen hazard rate estimate) for a single patient across multiple seizures, comparing preictal (blue) and interictal (green) epochs. This patient's preictal cHFO rates were on average higher than interictal rates. The scaled heatmap of cHFO rates (A) shows the contribution of individual channels to estimates computed from seizure onset zone channels (SOZ; B) and nonepileptic channels (OUT; C). Plots beneath B and C both show cHFO trajectories by individual seizure (without interictal reference). cHFO rate is defined as HFOs per minute per channel and is shown in the top rows of B and C with ±1 standard deviation (dotted lines). Yellow rectangles show the 1-minute ictal buffer, and red rectangles indicate the clinical duration of a given patient's longest seizure. The peri-ictal window was truncated for display purposes at 40 minutes.

between interictal and preictal periods. As seen in Figure 4, this patient's preictal cHFO rate is generally higher than the interictal rate.

In our visual observations, we saw significant temporal variability in preictal cHFO trajectories within our patient cohort across channel groups and epochs. We identified patients with preictal cHFO trajectories that were similar to interictal ones (examples in Figure 5A). There were patients with increased preictal cHFO activity over interictal baseline (examples in Figure 5B,C); of these, some had distinct bursts of preictal

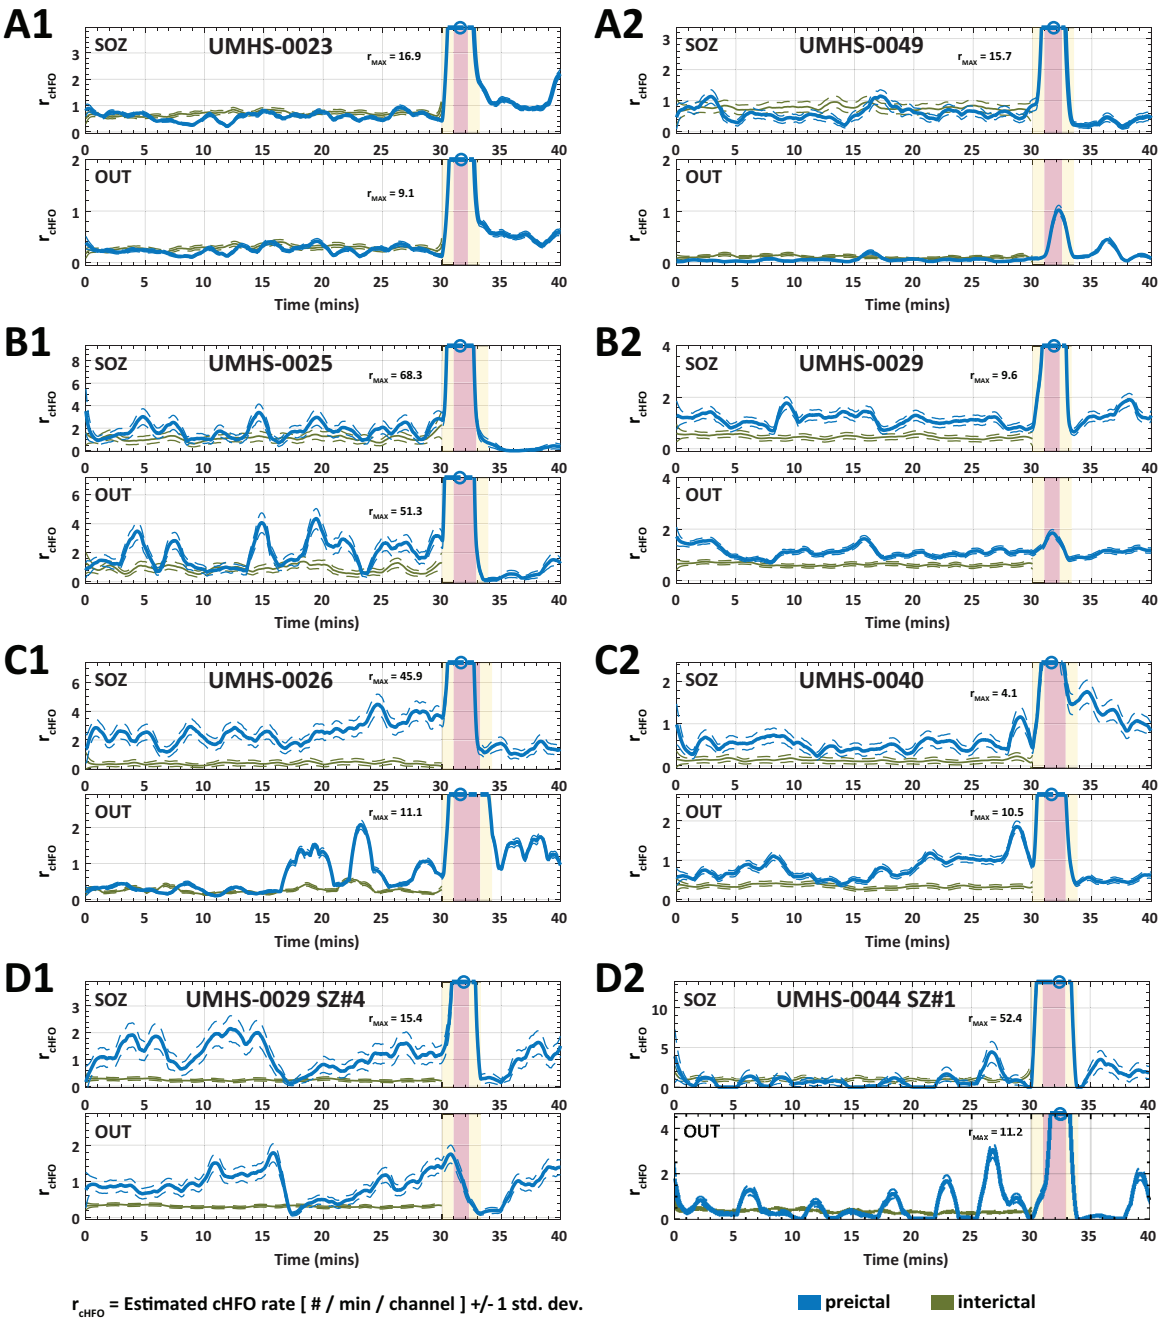

FIGURE 5: Variability of observed preictal continuous high-frequency oscillation (cHFO) rates. A, Many patients had few significant differences between interictal and preictal cHFO rates (example patients given in A1 and A2). B, Other patients displayed increased preictal cHFO trends relative to those of interictal periods; of these, periodic bursts of HFOs were evident in some (B1), whereas others showed more sustained increases in preictal HFO rates over interictal (B2). C, Two patients with gradually increasing preictal HFO rates were also identified. D, Examples of individual seizures in different patients, whose preictal cHFO rates also gradually increased toward onset, similarly to the average preictal trends of C. Here, cHFO rate is defined as HFOs per minute per channel. Visual formatting of all subfigures herein is the same as shown in Figure 3B,C. OUT, nonepileptic channels; SOZ, seizure onset zone channels; SZ, seizure.

cHFOs, and others had more sustained increases (Figure 5B). We also identified patients with preictal cHFO trajectories that appeared to increase gradually, leading to seizure onset (Figure 5C). These preictal trends were averaged across many seizures, but were also observed prior to individual seizures (Figure 5D). Even limited to visual inspection, these various changes were visible in at least 12 of the 30 patients. These example visual observations of preictal cHFO trends in various patients motivated further in-depth quantitative analysis, which we describe in detail below. Also, note that Figure 5 shows two patients (UMHS-0029 and -0040) in whom the HFO rate is higher in OUT compared with SOZ. As seen in Table 1, these were the only two patients who had this effect, which occurred when averaging over the entire region rather than selecting specific high-rate channels within the SOZ. Patient UMHS-0029 was not a responder, and UMHS-0040 had an atypical response described below.

Statistical significance of temporal trends: The visual observations in the previous section suggested that perhaps the change in the rate as seizures approach, rather than simply the magnitude, was associated with impending seizures. To quantify the temporal trends shown in Figure 5C, we compared the cHFO rates as mathematical functions. We used linear regression to fit a line to the 30-minute trajectory of cHFOs in the average preictal and interictal windows in each patient. These values are shown as population box plots in Figure 6A, where we define slope (Δ*cHFO rate*) as the change in HFO rate over 30 minutes, with rate given as HFOs per minute per channel. A number of patients had high preictal slope in SOZ channels, whereas across the population,

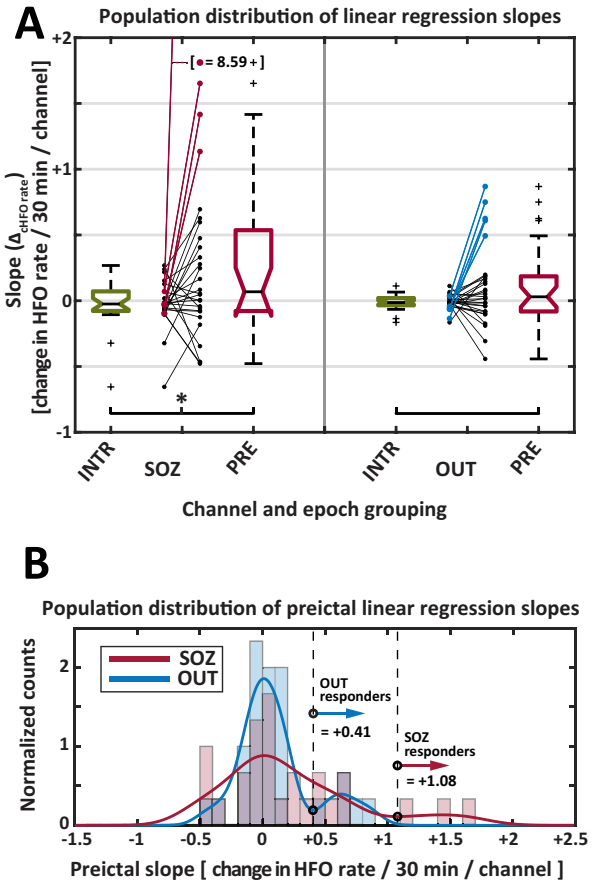

*FIGURE 6: Population comparisons of cHFO rate regression slopes. A, Population box plots of regression slopes fitted to continuous high-frequency oscillation (HFO) rates of interictal (INTR) and preictal (PRE) epochs, organized by channel group (seizure onset zone channels [SOZ], nonepileptic channels [OUT]). As a population, increased preictal slopes were observed only in SOZ channels (*P < .05, Wilcoxon signed rank test). Differences in raw data during interictal and preictal epochs are visualized per patient between box plot groups; ''slope responders''—patients with increasing preictal continuous HFO (cHFO) rates in SOZ and OUT channels—are shown with red and blue lines respectively, whereas other patients are shown with black lines. B, Smoothed and binned population distributions of preictal cHFO regression slopes are shown by channel group; both SOZ and OUT distributions are bimodal. OUT slope responders (blue) have a slope threshold of +0.41 over 30 minutes, and SOZ slope responders (red) have a slope threshold of +1.08 over 30 minutes. Here, we define cHFO regression slope ($\Delta_{cHFOrate}$) as the change in HFO rate over 30 minutes, where HFO rate is defined previously as HFOs/minute/channel.*

interictal slopes were close to zero. We compared the distributions with a signed rank test, which takes pairwise differences between the preictal and interictal periods for each patient. The SOZ had a significant increase in slope (median Δ*cHFO rate*, *PRE*–*INTR*= 0.13, $P < .05$), whereas in OUT there was no appreciable difference

(median Δ*cHFO rate*, *PRE–INTR*= 0.01, *P* = .15). As seen in Figure 6A, the differences were primarily due to certain patients with higher rate who were different from the rest of the group. To identify these potential outliers, we used a strategy similar to that shown in Figure 3B; we made a histogram of preictal slopes, fit them with a kernel density estimator, and looked for natural thresholds. In this case, the preictal distributions were statistically different from interictal ones for both channel groups (Kolmogorov-Smirnov: SOZ, OUT: $P < .05$, $P < .01$). The threshold for outliers, that is, "responders," was OUT Δ*cHFO rate*= +0.41, SOZ Δ*cHFO rate* = +1.08. This gave a total of four patients in the "SOZ slope responder" subset, and five in the "OUT slope responder" subset (individuals marked in Table 1, and colored lines in Figure 6A). The responders were chosen solely on the basis of their preictal slopes being outliers, but note that the difference with interictal Δ*cHFOrate,PRE–INTR* in each case was also very high. We thus conclude that the preictal change in cHFO rates is a novel potential biomarker of seizure onset.

Relationship of responders with clinical metadata: We evaluated whether any of the three responder groups (mean rate, n = 7; SOZ rate, n = 4; OUT rate, n = 5) were correlated with clinical factors from Table 1. Of these responders, four had International League Against Epilepsy class I outcomes, four had class II, and three did not have resections (Table 1). We could not find any consistent demographical or etiological factor that was associated with a particular "responder" subset of patients; the rate of class I outcomes was similar to that of the whole group, and there were not enough patients to have sufficient power to identify specific differences in other factors such as

location and pathology. We analyzed whether these results in 30 patients would be likely to apply to the larger epilepsy population. We evaluated this with a binomial confidence interval, with 30 samples and 11 successes ("responders"); the 95% confidence interval is 20%-56% (6-16 patients). Considering that as low as 38% of patients with refractory epilepsy achieve lasting seizure freedom after surgery, [94], [95] we feel this responder rate is likely to have significant clinical impact as a biomarker. It is highly likely to be present in a large number of patients in larger studies.

Discussion

We performed a systematic analysis of time-varying HFO rates in a large cohort of patients with refractory epilepsy, robustly comparing interictal and peri-ictal rates for the first time. Our analysis of mean HFO rate found no difference between preictal and interictal rates at a population level. Despite this, we used a data-driven approach to identify a putative subset of patients who are "mean rate responders," in whom there was a large difference between preictal and interictal rates. We also found that mean HFO rate was highest in SOZ channels, which corroborates existing findings that interictal HFOs localize epileptic tissue, [35], [43], [96] although we have confirmed it for preictal, ictal, and postictal epochs as well. Mean ictal HFO rates were significantly higher than rates for other epochs, a finding also supported in the literature [56], [97], [98].

Prior HFO work has been based upon average rates over long windows (i.e., 10 or 30 minutes). Here, we investigated peri-ictal HFO trends as a continuous function of

time (cHFO rate), which estimates the “hazard rate” of HFOs occurring at any given moment in time. Despite little evidence of population-wide stereotypy, this revealed many varied and unique temporal patterns of peri-ictal cHFO trajectories among individuals. In our statistical analysis of cHFO rates, we compared the relative magnitude of preictal and interictal cHFO trends by their linear slope and again used their underlying distributions to identify two subsets of patients (“SOZ slope responders” and “OUT slope responders”) with increased preictal cHFO activity relative to other patients.

These results are supported by previous findings, although there have been relatively few papers dealing with the effects of preictal HFOs. Early work found that HFOs had significant preictal changes in small cohorts of patients [55], [56]. Other studies investigated high-frequency activity, but not necessarily discrete HFOs, and found similar results. One found that increases in 60-100 Hz power preceded seizure onset by as much 20 minutes in patients with refractory neocortical epilepsy [74]. Another showed that a predictive classifier of preictal state performed well in a subset of seven of 53 patients, each of whom showed distinct changes in preictal high-frequency activity that were coupled with slower brain rhythms [99]. The authors noted that their algorithm might have been successful in more patients if their cohort were more homogenous. Our work has quite similar results with HFOs; in our clinically diverse population, there were distinct subsets of patients in whom HFO rate reliably increased prior to seizures, albeit in different but complementary ways.

We did not identify any factors to predict which patients would be “responders”; however, it is important to point out that this is not a major concern, because the potential use case for HFOs as a temporal biomarker would require intracranial monitoring, which can be used to identify and train an algorithm post hoc. Thus, we do not anticipate that clinical metadata alone could be used to stratify which patients could be candidates. However, we did a deep analysis of the OUT slope responder group, as this indicated patients in whom HFO data suggested possible epileptic pathology outside of the SOZ. UMHS-0026 and -0032 were responders in all three groups, suggesting HFOs were strong biomarkers across all recorded channels. The other three, however, were only OUT slope responders. Two of them (UMHS-0025 and -0040) had secondary foci identified by the treating clinicians that were not included in the final SOZ channels. The other (UMHS- 0027) had seizures with diffuse onsets. From this cohort, we hypothesize that high preictal change in HFO rate may be associated with the seizure-generating tissue, and may suggest an independent method of using HFOs to identify the epileptogenic zone. In other words, the OUT slope responders may indicate a previously unrecognized method to use HFOs to identify the epileptogenic zone.

This analysis has some clear limitations. HFO occurrence is not a linear phenomenon, so applying a linear regression to the rate cannot capture the complex brain dynamics that produce it, and we make no claim that it was the “best fit” to the data. This function was chosen as the simplest method to characterize a generic increase in HFO rate during the preictal period across patients. Our goal was to

investigate gradual changes in preictal HFO rate across many seizures; accounting for nonlinear factors that would better model these variable cHFO trends was beyond the scope of this study. This analysis was designed to determine whether HFO rates were related to seizure onset, but was not designed to "predict seizures," as it averaged preictal behavior across many seizures. Furthermore, this work analyzed only the HFO rate; there are numerous additional features of the HFOs such as amplitude, spectral content, and duration [100] that will enrich this analysis in future work. There is also evidence of preictal EEG changes that may be applicable to HFOs, [56], [81], [101] and seizures themselves undergo changes in dynamical states, which may also affect HFOs [13], [102], [103]. These varied features provide a rich environment for future analyses, using robust methods to compare interictal and preictal data, to assess HFOs as a potential seizure prediction biomarker [59], [60].

Conclusion: Our investigation found that peri-ictal HFO rates and trends vary significantly across patients and even within individuals. We found a subset of patients in whom HFOs could be a valuable tool to identify the preictal state. This potential biomarker could be utilized in future studies on seizure prediction, focusing on in-depth characterization of interictal variability of HFO rates and greater numbers of seizures. Additionally, such work could better define the role of pathologic high-frequency activity in the mechanisms of seizure generation and its implications for the disease of epilepsy as a whole.

Chapter III: Viability of Preictal High-Frequency Oscillation Rates as a Biomarker for Seizure Prediction

Accepted for publication: Frontiers in Human Neuroscience, 12/07/2020

Abstract

Motivation: There is an ongoing search for definitive and reliable biomarkers to forecast or predict imminent seizure onset, but to date most research has been limited to EEG with sampling rates < 1000 Hz. High-frequency oscillations (HFOs) have gained acceptance as an indicator of epileptic tissue, but few have investigated the temporal properties of HFOs or their potential role as a predictor in seizure prediction. Here we evaluate time-varying trends in preictal HFO rates as a potential biomarker of seizure prediction. **Methods:** HFOs were identified for all interictal and preictal periods with a validated automated detector in 27 patients who underwent intracranial EEG monitoring. We used LASSO logistic regression with several features of the HFO rate to distinguish preictal from interictal periods in each individual. We then tested these models with held-out data and evaluated their performance with the area-under-the-curve (AUC) of their receiver-operating curve (ROC). Finally, we assessed the significance of these results using non-parametric statistical tests. **Results:** There was variability in the ability of HFOs to discern preictal from interictal states across our cohort. We identified a subset of 10 patients in whom the presence of the preictal state could be successfully predicted better than chance. For some of these individuals, average AUC in the held-

out data reached higher than 0.80, which suggests that HFO rates can significantly differentiate preictal and interictal periods for certain patients. **Significance:** These findings show that temporal trends in HFO rate can predict the preictal state better than random chance in some individuals. Such promising results indicate that future prediction efforts would benefit from the inclusion of high-frequency information in their predictive models and technological architecture.

Introduction

One of the most debilitating aspects of epilepsy is the uncertainty patients feel, not knowing when the next seizure will occur. Though seizures themselves account for an extremely small percentage of an individual's time, [66] the constant threat of a seizure can make the planning of normal day-to-day activities an impossibility for some [7]. This has led many investigators to search for methods to predict when seizure might occur. [12], [58], [62], [104], [105]

While 'seizure prediction' has been an attractive research subject for decades, early efforts had many unforeseen challenges. While there was evidence that EEG changed in the minutes or hours before seizures, [58] it was difficult to prove that these measures could work prospectively. A major breakthrough occurred when rigorous statistics were developed—the key was to show that a given algorithm could outperform random chance [59], [60]. Several studies then followed using this method, and were able to show that intracranial EEG signals could predict the preictal state better than chance. [64]–[66] Critical in that work was the unprecedented collection of months of continuous

EEG in a clinical trial in Australia, which allowed for rigorous long-term statistics [64], [66]. That dataset has become a crucial tool in later work, including international competitions, [64] as prediction algorithms have made many further improvements [65], [67], [70], [106]. However, the data also have two important limitations: the data were acquired at low sampling rate (200 Hz) that does not allow analysis of high resolution EEG signals; and more importantly, since the trial ended no similar chronic recordings have been collected.

Thus, while there have been many very promising results in the field of seizure prediction, most work has been focused on a single dataset of long term, low resolution intracranial EEG. The results have proven that seizure prediction is possible in many patients, but clearly are far from optimal. One potential avenue for further improvement is the possibility that higher resolution EEG could hold greater information. In particular, over the past 20 years it has become increasingly apparent that High Frequency Oscillations (HFOs) are a powerful biomarker of epilepsy [30], [35], [43]. HFOs consist of short (< 100 ms) oscillations in the 80-500 Hz frequency band, and require sampling rates of at least 2000 Hz for accurate identification [51]. HFOs are more likely to occur in the epileptogenic zone [35] and may help guide surgical decisions [38]–[40], [107]. One relatively unexplored aspect of HFOs is that their characteristics can also change in the 30 minutes prior to seizure initiation in certain individuals [55], [56]. These preliminary studies were constrained by small patient cohorts and datasets that were not as specific as currently-available methods [83], [108]. Nevertheless, the evidence from those studies motivate using HFOs to identify the preictal state.

Utilizing population-level inference and a large clinical dataset, our group recently found several features of HFO rates that were highly correlated with the preictal state [109]. In that work, we averaged the HFO response over all available data per patient and compared the responses during interictal and preictal epochs; several patients had significant results. However, in order to utilize HFOs to identify the preictal state prospectively, a different analysis is necessary. The HFO response in a given segment of data must be compared individually to that of other segments, rather than in aggregate as in that prior work.

Robust implementation of seizure detection algorithms requires several months of continuous recording, as was accomplished by the Neurovista trial in Australia [66]. Such data with sufficient sampling rate to detect HFOs is currently impossible to attain. Until such devices are available, the only alternative is to utilize inpatient intracranial EEG monitoring, which lasts less than 2 weeks. Although such data are vastly inferior, they are also the only current option. Until implantable devices with >1000 Hz sampling rate are available, the role of HFOs in the specific context of seizure prediction must first be evaluated using only the limited intracranial monitoring data available, which is our goal herein.

With this study, we evaluate the preliminary usefulness of HFOs in patient-specific seizure prediction. We employ state-of-the-art automated HFO detection methods on the entire recorded intracranial EEG data of a clinically-diverse cohort of 27 patients. With more than 10 million detected HFOs in this dataset, we use various features of HFO rates as predictors in patient-specific preictal classification models. With robust

machine learning methods and statistical techniques to validate our results, we find that 10/27 patients have excellent classifier performance. These results are limited due to the short recording periods, but were very promising. While the technology does not yet exist that would allow a full prospective analysis using high resolution data, these results motivate future studies that incorporate such technology in the next generation of seizure prediction devices.

Methods

Patient population: To form our patient cohort, we looked at all patients with refractory epilepsy who had undergone intracranial EEG (iEEG) monitoring at the University of Michigan from 2016 – 2018. In order to ensure that sufficient data was available for training and testing our models, we required patients with the following: 1) a defined seizure onset zone, 2) at least three recorded seizures that were each preceded by non-zero HFO rates, and 3) the availability of at least 24 hours of data; applying these criteria to the 32 available patients resulted in 27 patients. The study was approved by the local IRB, and all patients in the study consented to have their EEG data de-identified for later analysis. Of note, all data were acquired under standard clinical procedures, and the current work was done retrospectively: no data from this research had any effect on the clinical care. Further summary of the patient population is found in Table 2.

TABLE 2: Clinical data for second study. Abbreviations: M/F: male, female, L/R: left / right, T: temporal, P: parietal, F: frontal, Occ: occipital, DNET: dysembryoplastic neuroepithelial tumor, NF1: neurofibromatosis type 1 tumor, NR: not resected, CD: cortical dysplasia, MTS: medial temporal sclerosis, PVNH: periventricular nodular heterotopia, PMG: polymicrogyria, VNS: vagal nerve stimulator.

Subject	Age	Sex	ILAE outcome	Seizure focus (hemisphere, region)	Pathology / *implant type*	Number of intracranial channels				Total recorded time (hours)	Total number HFOs	Mean HFO rate (#/min./channel)		Number of seizures				Responder window subset (window duration, min.)		
						total	ECoG	depth	SOZ			SOZ	OUT	total	used	training	testing	30	15	10
UMHS-0018	41	M	Ib	L F	CD	32	0	32	4	59.8	108,510	4.18	0.54	3	3	2	1			
UMHS-0019	59	F	II	R T	Gliosis	106	106	0	2	168.8	170,946	2.30	0.19	5	3	2	1			
UMHS-0020	45	F	II	R T	MTS	25	0	25	9	171.2	54,254	0.38	0.12	7	7	5	2			
UMHS-0021	30	M	II	R T	Gliosis, PVNH, PMG	46	0	46	13	179.5	394,398	1.98	0.50	9	7	5	2			
UMHS-0023	29	M	NR	L T, P	PVNH / *Neuropace*	69	41	28	29	164.3	390,134	0.86	0.37	20	9	6	3			
UMHS-0024	31	M	NR	L, R T	*Neuropace*	75	55	20	16	177.2	1,649,380	3.40	1.71	28	11	7	4			
UMHS-0025	17	F	II	L T	Gliosis	20	0	20	5	207.7	270,125	1.75	0.86	10	5	3	2			
UMHS-0026	22	F	NR	R T	PVNH	52	0	52	3	246.2	382,201	1.28	0.45	40	10	7	3	X	X	X
UMHS-0027	26	M	NR	L Diffuse	*VNS*	91	81	10	3	205.2	1,601,359	1.90	1.41	97	11	7	4			
UMHS-0028	14	F	I	R T	Tumor: Glioma	53	47	6	5	79.7	140,782	2.95	0.42	7	6	4	2	X	X	X
UMHS-0029	48	M	NR	L T, Occ.	*Neuropace*	91	91	0	22	226.3	847,560	0.60	0.71	14	7	5	2			
UMHS-0030	5	M	III	L T	MTS, Gliosis	100	100	0	2	146	330,614	0.98	0.56	33	21	14	7	X		X
UMHS-0031	13	M	I	L T	Gliosis, Tumor: NF1	99	99	0	6	180	263,676	1.17	0.39	9	4	3	1			
UMHS-0032	41	F	I	R F	CD	32	0	32	3	184.3	295,865	3.79	0.96	8	6	4	2			
UMHS-0033	5	F	II	R Ins.	CD, Gliosis	74	0	74	4	120.7	233,883	1.40	0.38	28	8	5	3		X	X
UMHS-0034	33	F	I	R F	Gliosis	32	0	32	11	136.3	448,718	2.58	1.26	17	16	11	5	X		
UMHS-0035	50	F	I	L T	Gliosis	57	57	0	2	162.7	108,147	0.73	0.21	7	4	3	1		X	
UMHS-0036	43	M	NR	L, R T	CD / *Neuropace*	54	0	54	2	172.5	347,928	1.34	0.60	18	12	8	4			
UMHS-0039	47	M	NR	R P	CD / *Neuropace*	90	0	90	10	155.2	266,422	1.02	0.23	19	9	6	3			
UMHS-0040	14	F	I	L P	CD, Gliosis	63	55	8	8	196.7	323,180	0.38	0.66	7	7	5	2		X	
UMHS-0041	32	F	I	R F	CD	71	0	71	9	176.5	43,350	0.27	0.04	36	3	2	1			
UMHS-0043	28	M	II	R T	Gliosis	86	0	86	9	182.2	386,967	1.34	0.42	46	16	11	5		X	X
UMHS-0044	45	F	NR	L T, P	*Neuropace*	76	0	76	6	170.2	414,195	1.29	0.47	13	5	3	2			
UMHS-0045	17	F	NR	L, R T	*Neuropace*	94	0	94	15	331.5	631,551	0.79	0.25	6	6	4	2		X	
UMHS-0046	23	F	I	L F	CD	30	0	30	9	139.3	16,575	0.15	0.04	17	5	3	2			
UMHS-0048	22	F	NR	L, R T	*Neuropace*	86	0	86	8	141.8	404,972	2.76	0.33	23	8	5	3	X	X	X
UMHS-0049	53	F	NR	L, R T	*Neuropace*	94	0	94	15	176.8	287,303	0.98	0.16	17	5	3	2			
					TOTALS / averages	1798	732	1066	230	4658.6	10,812,995	1.58	0.53	544	214	143	71	5	8	6

Data acquisition: All intracranial recordings were sampled at 4,096 Hz with a Quantum amplifier (Natus Medical Inc., Pleasanton, CA); the electrodes implanted for monitoring consisted of subdural grid, depth and stereo-EEG electrodes, as deemed appropriate for each patient during standard clinical care. All recordings were referenced to a lab-standard instrument reference placed midway between Fz and Cz when first recorded, and then were re-referenced for HFO detection using common average referencing [86], which was applied to all electrodes of the same type, e.g. all depths or all grids or strips together. The treating epileptologist determined which channels comprised the seizure onset zone (SOZ channels), as well as the onset and offset times of all seizures; we obtained these metadata through the official clinical report for a given patient. Channels within the resected volume of tissue (RV channels) were identified and labeled through consultation with the neurosurgeon and by pre- and post-op imaging comparisons if available. Any channel that was not labeled as an SOZ or RV channel was labeled as an OUT channel. Note that a seizure prediction algorithm should have knowledge of the SOZ and OUT channels available, as it must be trained on previous seizures and would be implemented after these studies are completed. It is also important to note that the SOZ is what was determined by the reading clinician and does not depend upon being the true epileptogenic zone. We incorporated the analysis of OUT channels as a conservative way to account for diagnostic uncertainty and see if other channels also had useful information. Channels labeled as RV that did not overlap with the SOZ were not used in our analysis, in order to maintain a more conservative analysis.

Data analysis: All data analysis was conducted with custom MATLAB (Mathworks, Natick, MA) and C++ functions and scripts. As described in detail below and shown in the block process diagram of Figure 7, this analysis consisted of several steps: first, automated HFO detection was performed on all patient data. Then, several features across consecutive time windows of varying duration were computed from HFO rates. These features were used to train a logistic regression model to distinguish preictal versus interictal states. The algorithm was cross validated with held-out data and compared versus random chance. Model performance was quantified using ROC curves.

Automated HFO detection: All HFOs were identified with a validated automated detector [86] with additional modifications described further below. In summary, this detector is based upon the original 'Staba' RMS-based detector [52] which then increases the specificity by redacting detections that overlap in time with several EEG artifacts such as sharp transients, electrical interference and noise, and artifacts from signal filtering. To further increase HFO specificity, we excluded detected events with waveforms consistent with features of muscle (EMG) artifact, using another validated algorithm [85] as in our previous work [109]. Of note, these algorithms have previously been shown to be similar to human reviewers [86], [110].

We also modified the data processing pipeline of our automated detector to ensure that it functioned appropriately within the unique constraints of seizure prediction. Most automated detectors operate by processing incoming EEG data in successive epochs of fixed length, e.g. 10 minutes, and then assess the background activity of the entire

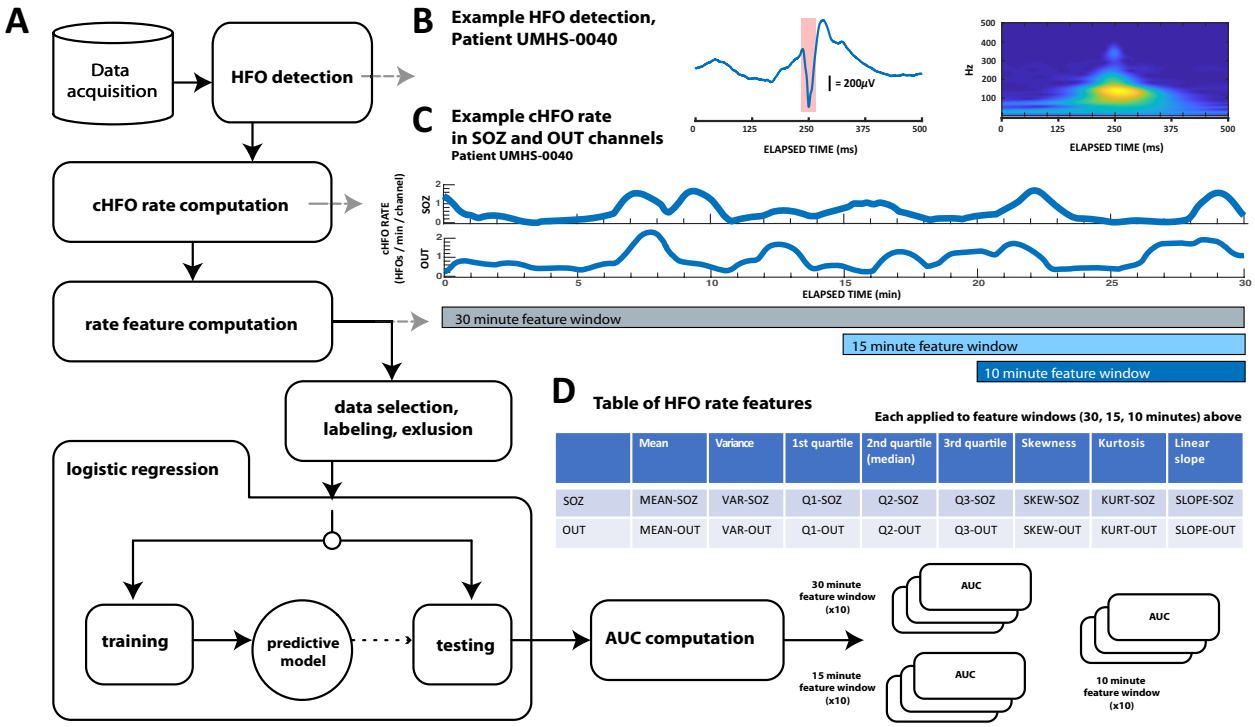

	Mean	Variance	1st quartile	2nd quartile (median)	3rd quartile	Skewness	Kurtosis	Linear slope
SOZ	MEAN-SOZ	VAR-SOZ	Q1-SOZ	Q2-SOZ	Q3-SOZ	SKEW-SOZ	KURT-SOZ	SLOPE-SOZ
OUT	MEAN-OUT	VAR-OUT	Q1-OUT	Q2-OUT	Q3-OUT	SKEW-OUT	KURT-OUT	SLOPE-OUT

FIGURE 7: Block process diagram for study analysis. Schematic diagram showing overall data analysis workflow. (A) General analysis workflow. After automated HFO detection, continuous HFO rates (cHFO rate) are computed in both the SOZ and OUT channel groups. Next, several statistical quantities (features of HFO rate) are computed from cHFO rates in three 'feature windows' of different durations: 30-, 15-, and 10- minute feature windows. After labeling this feature data as either preictal or interictal, observations that remain after an exclusion process are randomly divided into training and test data sets. Training data is used to train predictive LASSO logistic regression models, which are then tested with unseen testing data. The performance of each model with this testing data is assessed by computing the test AUC value, which, when averaged over 10x cross-validation runs for each of the three feature windows, are finally compared across patients; these results are visualized in Figure 9. (B) Example HFO detection, 'responder' patient UMHS-0040. The HFO waveform is displayed on the left, while its time-frequency decomposition (computed with the Morse wavelet) is visualized on the right. (C) Example of cHFO rates computed for patient UMHS-0040. Continuous HFO rates (cHFO rate - defined as HFOs / min / channel) are computed in both the SOZ and OUT channel groups separately. The rate features used in the proceeding Table D are computed from these cHFO trajectories in 30-, 15- and 10-minute segments. (D) Table of rate features. Eight features are applied to cHFO rates per channel group (SOZ and OUT channel groups), which yields a total of 16 rate features. Abbreviations shown in this table are used throughout the text.

epoch to determine a threshold for detecting HFOs within that epoch. That process cannot happen in real-time nor (pseudo)prospectively, because evaluating a potential HFO at a specific point in time requires knowledge of background activity that has yet to occur. Such a process would not be possible for prospective seizure prediction, in which there should be no knowledge of the future. To address this constraint, we modified the

detection algorithm to work prospectively. First, we approximated real-time detection by only detecting HFOs for 30 seconds at time. Second, we still used 10 minutes of EEG to calculate the background, but use the *previous* 10 minutes of EEG data, relative to the end of each of data segment. In effect, the algorithm is identical to the previous one except it only reports the HFOs that are detected during the final 30 seconds of a 10 minute segment, and the same process is repeated by sliding the 10 minute window forward 30 seconds. One outcome of this is that the first HFOs detected in any given data file start after the first 10 minutes of recording. With these adaptations, our automated HFO detection was better suited to the constraints of seizure prediction, and more closely resembled a real-time process. Further – and perhaps most importantly for preictal HFO detection – these changes also prevented seizure activity from influencing the detector. We compared these results to those of the original detector, and there was no appreciable difference in HFO rate (data not shown), which is expected since there were no changes inherent to the detector itself, but rather how it was fed data.

Computation of HFO rate: In order to investigate temporal variations in HFO rate with sufficient resolution, we approximated HFO rate (which we define as the number of HFOs per minute per channel) in both SOZ and OUT channel groups as a continuous function of time (cHFO rate). The cHFO rate was obtained by calculating the estimated HFO rate during one minute of data, then sliding the one-minute window forward one second and recalculating. This sliding window method approximates a continuous HFO rate with a 1 second time resolution. The sliding window was applied to all SOZ or OUT channels, which were grouped separately. For a given window segment and channel

group, the HFO rate was computed by summing the number of HFOs occurring across all channels of the same group; this value was then divided by the total number of channels in that respective group, which resulted in an estimate of the average cHFO per channel within each group (SOZ or OUT).

Features of HFO rate: The advantage to using cHFO rate as computed above - rather than averaging it over longer periods - is that the temporal resolution of cHFO rates can reveal fluctuations and patterns in HFOs down to the scale of a second - which could be important in characterizing preictal trends. We quantified the temporal variation of cHFO rates with several descriptive statistics, including mean, variance, linear slope, quartiles, skewness and kurtosis across a given epoch of time. We also compared linear trends in cHFO rates using the slope extracted from linear regression applied to cHFO rates for a given epoch of time. All these values were computed separately in SOZ and OUT channel groups across three different epochs of time: 30, 15, and 10 minutes, which we call 'feature windows'. The feature windows were designed to account for possible differences in seizure horizons between patients, as we hypothesized that the duration of any preictal state would not be constant across the entire cohort. All features were computed from the start of a given data file in consecutive 1-minute intervals. Each feature window was analyzed independently of the others throughout the entirety of the study.

Feature data labeling and exclusion: In machine learning, classification algorithms used in prediction need labeled observations of data in order to train their models. In this case, we label data as either interictal or preictal. Based on our prior

data showing HFO features changing up to 30 minutes prior to seizures, [56], [109] we defined the ‘preictal period’ as the 31 minutes prior to the start of the seizure. The extra minute occurs because we inserted a buffer of one minute just prior to seizure onset, which accounts for some interrater variability in seizure onset time [87].

For each of the feature windows (10-, 15-, or 30-minutes), the ‘preictal’ windows were defined as the last window immediately prior to the seizure, but not including any of the 1 minute just before seizure onset. Because the calculations slide forward in 1-minute steps, this means each ‘preictal’ feature window ends between 1-2 minutes prior to the clinician-determined seizure onset time. For each feature window length, we only included the one ‘preictal’ window immediately before the seizure. Because our prior data suggested up to 30 minutes could be considered as the physiological preictal period, to be conservative we ignored data during that period that was not in the ‘preictal’ feature window. Data from those times (the two previous 10-minute windows and one previous 15-minute window) were discarded from both the preictal and interictal analysis.

‘Interictal’ was defined as all data starting 11 minutes after a seizure until 31 minutes prior to the next seizure, which allows a one-minute buffer for uncertainties in the start/stop times of the seizure. We note that some research has shown that the preictal state may extend beyond 30 minutes [81], [82], so this definition is conservative and may not capture all differences. We calculated an ‘interictal’ feature window for every consecutive epoch (i.e. every 30 min for the 30-min feature window; every 10 min for the 10-min feature window).

There were other limited circumstances that we excluded from analysis. To ensure that seizures were evaluated independently of other seizures, such as when multiple seizures occur sequentially, we redacted preictal observations falling within peri-ictal extent (11 min postictal or 31 min preictal) of other seizures. Further, we also redacted any observation that overlapped with periods of incomplete or missing data, which could result from gaps within a file or from a file's end. Finally, considering our modifications to the HFO detector, any data observation overlapping with the first 10 minutes of a given data file was also redacted, as HFOs are not detected for the first 10 minutes.

Logistic regression model: We used a logistic regression model to classify preictal versus interictal data. Logistic regression determines the probability that given data is from a specific labeled class, and has been used in seizure prediction studies previously [111]. It also has the advantage of allowing us to analyze the relative contributions of each feature, rather than being a 'black box' approach. We trained models for each of the three feature windows (10, 15, 30 min) using 2/3 of the data and then testing on the remaining 1/3. This process was cross-validated 10 times for each feature window by randomly-selecting different interictal and preictal data, and re-running the training and testing step, for a total of 30 models per patient. Random selection, rather than chronological, was used because of the limitations of this dataset: unlike in the Neurovista dataset that had months for the recordings to stabilize, [112] our data is limited to 2 weeks of inpatient monitoring. This unavoidably leads to some variability over time due to various factors such as medication taper, sleep disturbances, and the settling of electrodes [112]–[114]. Here, we used random selection to reduce

the influence of these factors on overall model performance, but this also may reduce the effectiveness of the model.

In order to facilitate the models helping to determine which coefficients were most useful in forecasting seizures, we used LASSO logistic regression [111], [115], [116] to create the predictive models used in our study. Specifically, in Matlab we used the lassoglm function, with the following general syntax: lassoglm(XTrain, yTrain, 'binomial', 'CV', k), where XTrain is the feature vector, yTrain is a binary vector with '0' for interictal and '1' for preictal, and k is chosen as the number of seizures within the training data. This function inherently cross-validates the trained model based upon the number of seizures k, which reduces overfitting. In general, LASSO introduces a penalty on the absolute value of the coefficients, and optimizes the model by iterating through different penalty parameters to find the lowest error, while removing coefficients that have minimal effects [115]. Thus, one outcome of the training step is to identify which features were the most important for identification of the preictal state.

Assessing predictive performance: Each cross-validation iteration tests whether the predictive model can correctly classify novel preictal versus interictal data. We computed the ROC curve for each iteration, then computed the arithmetic mean of all the areas under the curve (AUC) across all ten iterations. A random predictor would have an AUC of 0.5, while a successful predictor should have an AUC higher than 0.5. We chose a nominal threshold of 0.6 to show the minimal improvement above 0.5 that would be meaningful. However, that threshold is subjective so we then tested the significance of each AUC using bootstrapping by randomizing preictal and interictal

labels (n=1000). The statistical significance of these average AUC was determined by taking the harmonic mean of the bootstrap *p*-values, [117] a procedure used in meta-analysis to combine *p*-values from multiple tests. Successful tests were those in which the average AUC was ≥ 0.6 and $p < 0.05$. We note that in clinical practice an AUC of 0.6 might be difficult to implement successfully on its own; however, it is comparable with prior seizure prediction work in standard EEG [12], [58], [62], [104], [105].

Results

Our heterogeneous patient cohort was comprised of individuals with a variety of ages, clinical etiologies and pathologies, and seizure foci. Out of 32 original patients in our database, four patients (UMHS-0037, -0038, -0042, -0047) were excluded either because of insufficient recorded seizures or undefined seizure onset zones. One patient in particular (UMHS-0022) had seizures with no HFOs prior to onset; this patient was also excluded, which left a total of 27 patients remaining for further analysis. Across these 27 patients, we detected more than 10 million HFOs across over 190 total days of intracranial EEG recordings. Over 210 seizures and 3,800 hours of interictal data (average of 8 seizures and 141 hours per patient) were used to train and test our classification models.

Comparison of test AUC values: We first assessed the general responses across all cross validation models in all patients. Over the 27 patients, with 30 models each

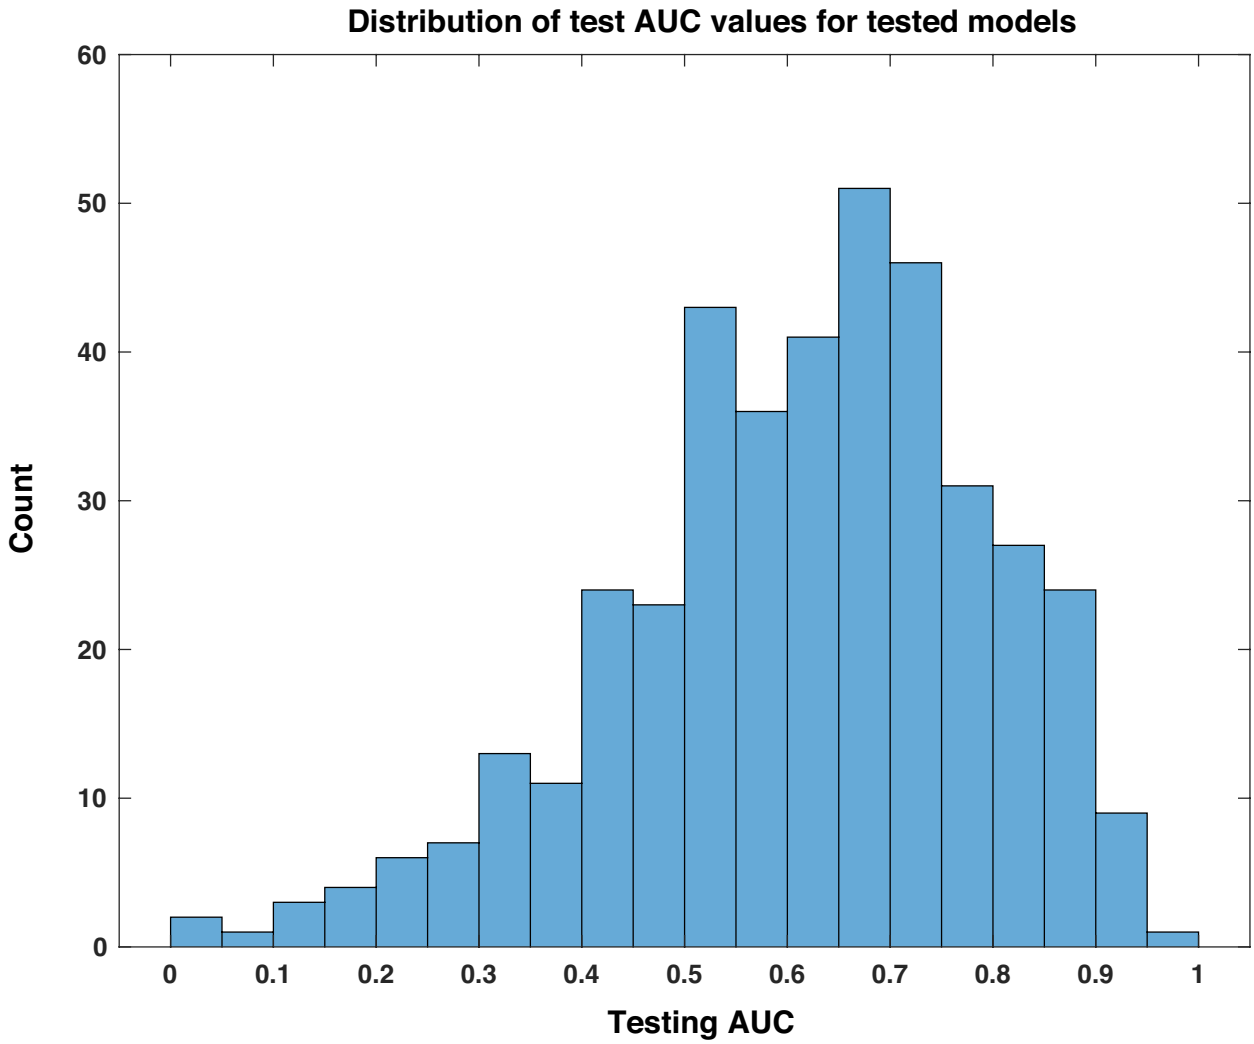

FIGURE 8: Distribution of test AUC values for tested models.This histogram of testing AUC values, computed for all tested models individually over all patients and feature windows, is skewed toward predictive performance that is better than random chance, i.e. values higher than 0.5.

(810 total), the model successfully converged to a solution in 403 instances (49.8%). The non-converging solutions are easily identified because all coefficients for HFO features are 0, and it is obvious that the model could not be used. In such cases, we conservatively assigned them a testing AUC value of 0.5 (and a bootstrap *p*-value equal to 1) – the same performance as a random predictor. The remaining patient models were composed of linear combinations of HFO rate features. As shown in the histogram of Figure 8, the distribution of test AUC values for these models overall showed

significant variability and spread from 0.5 (AUC test - maximum: 0.97, minimum: 0.024, median: 0.64). The skew of this distribution toward values greater than 0.5 suggests that a significant portion of models that used HFO features could perform better than random chance at identifying the preictal period.

We evaluated the consistency and reliability of this result within patients by determining if its average test AUC was at least 0.6 and if the average bootstrapped *p*-value was < 0.05. These values are shown with statistical significance noted in the bar plots of Figure 9. We found 10 out of the 27 patients had a significant response in at least one of the feature windows. We denote these 10 patients as 'responders,' and their average predictive response was robust and consistent. The presence of this subset of patients in our cohort suggests that there are measurable changes in preictal HFO rate preceding epileptic seizures that deviate from interictal trends. This finding shows that HFOs can act as a temporal biomarker of seizure onset in some patients.

Within the responder group, 4 were significant in only one feature window, while the rest had multiple. We compared the three windows (10, 15, 30 min) and found no evidence that the performance of one window was better than any other - either by how frequently it was significant in these patients, or by how high its overall performance was (Chi-square test: *p*-value = 0.61; Kruskal-Wallis test: $p = 0.737$). All responders and their significant windows are identified in Figure 9 and in Table 2. The p-values and associated asterisks indicating statistical significance in Figure 9 were based on individual bootstrap tests and not corrected for multiple comparisons.

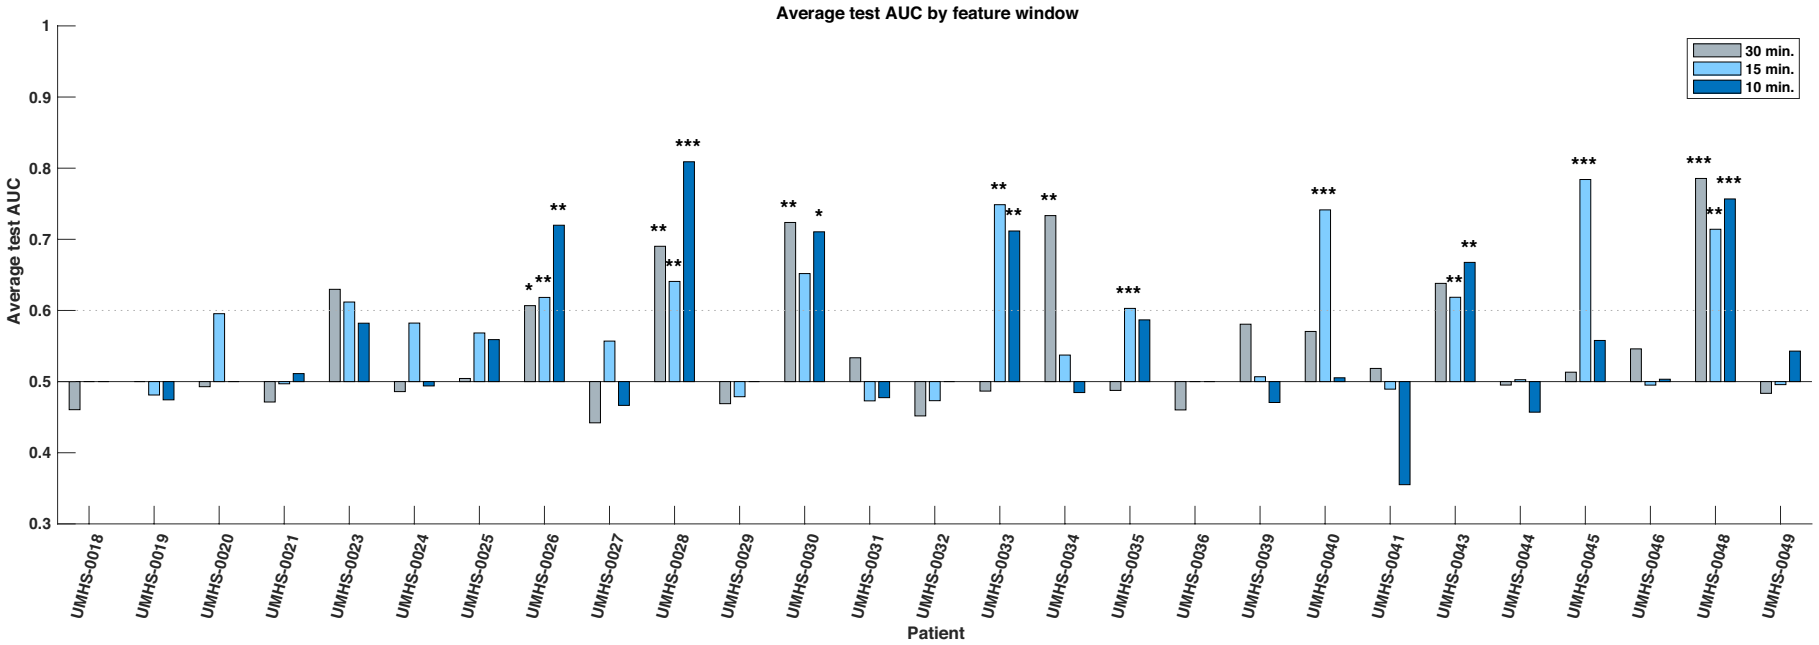

*FIGURE 9: Bar chart of average test AUC values by patient and feature window. 10 individual responder patients have significant predictive performance (average test AUC >= 0.6, significant average bootstrap test p-value < 0.05) in one or more feature windows. The statistical significance of the bootstrap test per feature window is indicated with asterisks: *, **, ***; p = < 0.05, < 0.01, < 0.001 respectively. Note that the significance is based upon how likely that patient's data could produce the given AUC by random chance, not whether the magnitude of the AUC itself is high.*

Significance of responder predictors: We investigated which features contributed to the significant predictive response observed in responder patients. Overall, both the combination and relative magnitude of HFO features in responder models varied significantly between patients, feature windows, and even between different cross-validation runs. Considering this variability, we could not evaluate feature importance directly by the raw coefficient values that resulted from LASSO logistic regression. Instead, we calculated how often a given feature was included among models - specifically, how often its corresponding coefficient was non-zero. In this manner, we considered the most commonly used features to be the most important to differentiating the preictal state from other interictal observations - whether its associated output coefficient was positive (which would indicate increased likelihood of an imminent seizure resulting from an increase in the feature's value) or negative (i.e. decreased seizure likelihood from a feature's increase). These frequencies of non-zero model coefficients per feature are shown by feature window in Figure 10, and are sorted in order from most to least common within responder models. Though we did not evaluate feature magnitude directly, we note that the medians of all responder SLOPE-SOZ features by patient and feature window were all positive, which reinforces our prior findings that gradually increasing HFO rates anticipate seizure onset [109].

While there were some observed differences in which features were the most common between window durations, there were no statistically-significant differences in feature frequency across the three feature windows (Kruskal-Wallis: $p = 0.64$). In terms of the most important features, the linear slope of HFO rate in the SOZ (Slope-SOZ)

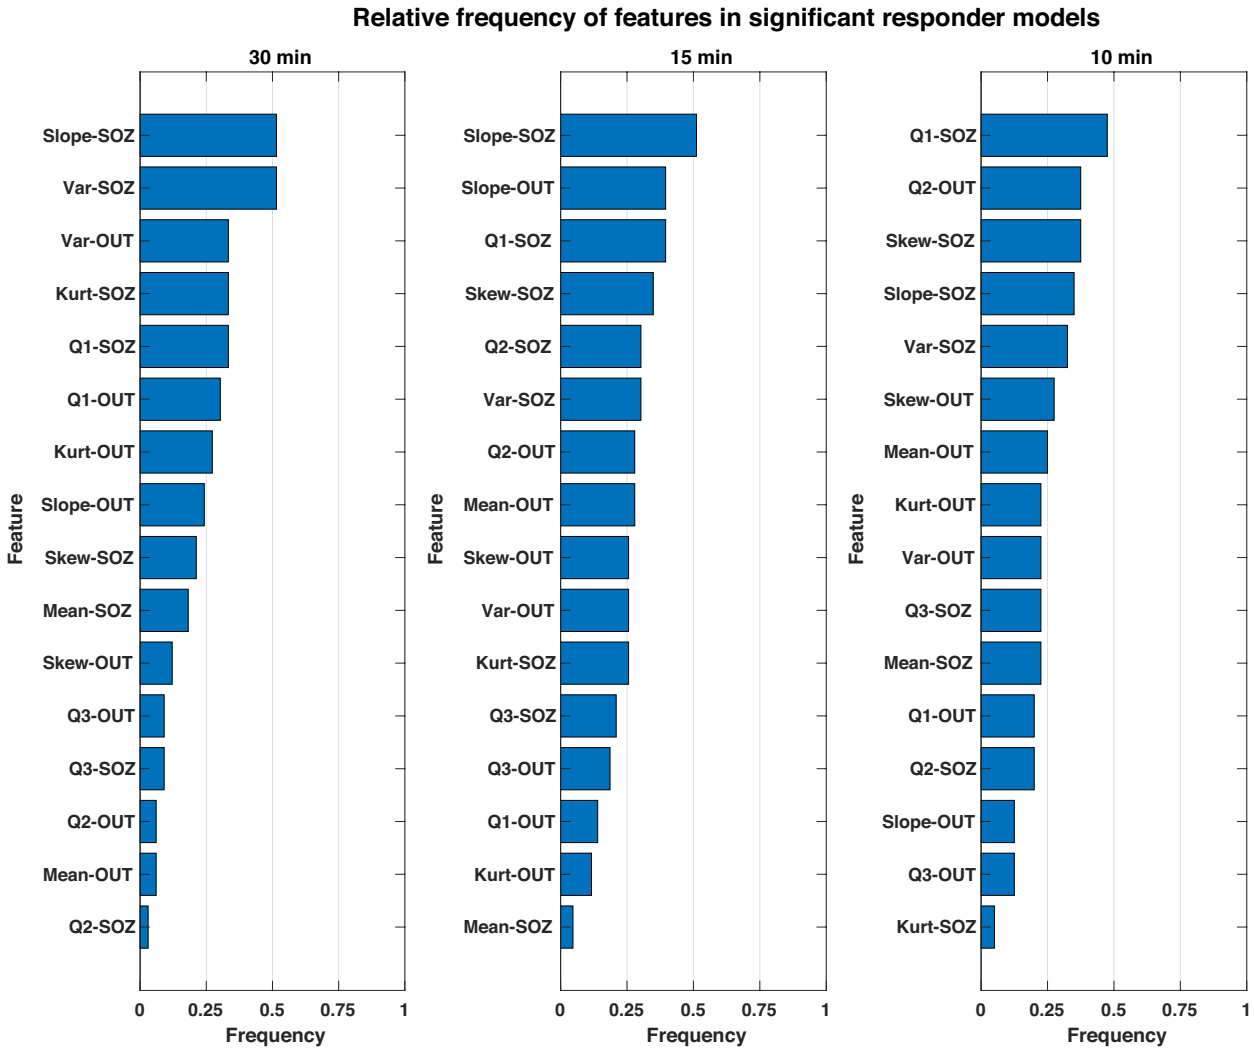

FIGURE 10: Relative importance of rate features in responders. Bar chart showing the relative frequency of each rate feature for only significant responder models. The features of HFO rate most important to discerning the preictal HFO response in responders are ranked in descending order (top to bottom) according to how often their respective model coefficients were non-zero for a given feature window. Overall, the most important feature was the SLOPE-SOZ feature, which was ranked first in both 30 and 15 minute feature windows. Also important were features in OUT channels, a novel finding that suggests HFOs outside epileptic tissue could still be involved in the process of seizure-generation.

was most important in both the 30 and 15 minute windows. Also common among important features were those computed from cHFO rates in OUT channels – channels that might be traditionally considered as less involved in pathological brain networks. Yet, there were no statistical differences in frequency between SOZ and OUT channel features (Rank-sum tests: $p = 0.34$, = 0.24, = 0.42 for 30, 15 and 10 minute windows

respectively), even though SOZ features were highest ranked across feature windows, with an average cumulative frequency almost 14% greater than that of OUT channel features. This suggests that HFO rates could be used to identify the preictal state regardless of their location.

Clinical factors of responders: Considering the clinical outcomes of responders, four were ILAE class I, two were class II, there was one class III, and the others were not resected. Comparing various clinical factors, there was no statistical evidence for differences in the composition of responder patients compared to the rest of the cohort. The ratio of temporal to extra-temporal seizure foci in responders was similar to other that of other patients (Fisher exact test: $p = 0.68$), and while there appeared to be a difference in the pathology of resected responders favoring gliosis, this was not significant in comparison to the rest of the cohort (Fisher exact test: $p = 0.14$). Despite lacking a clinical factor to differentiate this group from the rest of the population, based on our results, we estimate the relative proportion of responders in a given population is 19-55% of patients (95% binomial confidence interval with a test sample of 10/27), which demonstrates that patients with potential for significant HFO rate predictive performance could comprise a substantial portion of a large clinical cohort.

Discussion

In this first-of-its-kind study, we combined advanced automated HFO detection with the intracranial data of a large clinical cohort to investigate the potential use of high-frequency oscillations in seizure prediction. Across patients, we found significant

variation in the ability of time-varying properties of HFO rate to discern a preictal state. After applying a statistical benchmark to the average predictive performance of all models across our cohort, a subset of patient responders was identified that had consistent predictive performance better than random chance. The identification of these 10 individuals represents a novel finding and is our study's most important result. It provides firm support that high-frequency oscillations can function as a temporal biomarker of seizure onset, and additionally gives preliminary evidence that seizure prediction using HFOs is not only possible in a clinical context, it can hold significant potential for certain patients.

Another important outcome is the identification of *which* HFO rate features are the most useful. Ranked by their frequency in responder models across multiple windows of time, the most important predictive features of HFO rate included linear slope, variance, and the first quartile cHFO rate within the feature window. The most common feature was the linear slope, which measures gradual changes in HFO rate (either increasing or decreasing), suggesting that these changes are centrally important in determining if a seizure is imminent. One surprising finding was that even HFOs outside the SOZ were useful features. Note that it is not possible to compare relative magnitude of these feature coefficients directly because of the considerable model variability between patients, feature windows, and cross-validation runs. We analyzed the 10 responders and found that three of them had clinical situations in which the OUT channels were likely to be pathological. One patient had a known secondary seizure focus not included in the official SOZ (UMHS-0026), while another had high HFO activity in a non-resected

hippocampus that was likely dual pathology from a parietal lesion (UMHS-0040). However, the OUT features were not restricted just to those patients, and thus our finding of predictive value of HFO features outside the SOZ is an intriguing finding. This result suggests that HFOs even outside the SOZ provide important information on identifying impending seizures.

The test AUC values of responder patients we report are within the ranges presented in several seizure prediction studies, notably Brinkmann et al. 2016, Karoly et al. 2017, and Kuhlmann et al. 2018 [63]–[65]. There is one caveat to using the AUC metric in seizure prediction, as the inherent imbalance of interictal and preictal data can increase the reported specificity. In order to compare our work with other studies, however, this was an acceptable limitation for our analysis. While no prior work has evaluated HFOs for seizure prediction, there is evidence for a 'preictal state' [81]. HFOs have been shown to have different signal features [56] and changes in rate 30 minutes before seizures [109]. Further, some studies have shown distinct changes in high-frequency activity preceding seizure onset; some have also suggested that HFOs could be linked to seizure-generating mechanisms [74].

Despite our positive result, it must be noted that our overall methodology has a number of inherent constraints that limit our findings from being more widely applicable to seizure prediction in general. First, this analysis was based upon processing several minutes of data at a time (10, 15 or 30 minutes) rather than analyzing features of individual HFOs. There are a wide range of HFO features that could be incorporated into future prediction algorithms. Next, we note that 'true' seizure prediction would

involve choosing a specific algorithm and testing accuracy prospectively, which was not done here. Second, this method requires HFOs to be present and enough seizures to develop a predictive model; five of our cohort of 32 did not meet this standard. Finally, as stated before these data are limited to only 2 weeks immediately postoperatively during varied medication changes, which is known to be insufficient to have consistent EEG signals and sometimes even atypical seizures. Several of our patients had inconsistent results, but with so few seizures it is impossible to predict whether this would stabilize to an effective solution with more data. A much longer dataset under standard living conditions would be necessary to develop robust algorithms, but such data are not physically possible at present. Future work with a larger dataset could also incorporate additional features of the HFOs themselves (e.g. signal features such as frequency data), as well as previous prediction algorithms using standard EEG. This type of synergistic analysis on larger datasets could have much greater chance at a clinically-realizable seizure prediction algorithm.

Conclusion: Our results show that HFOs can function as a temporal biomarker of seizure onset. We show that changes in the HFO rate are capable of identifying the preictal state up to 30 minutes before a seizure in some patients. As a preliminary study, our findings are a foundation for future work pursuing individualized seizure-specific prediction efforts, which we envision could eventually function as a tool inside advanced implanted neuromodulation devices that utilize patient-specific and seizure-specific prediction methodologies. Advancement of this HFO seizure prediction framework, however, will require the availability of many chronic high-sampling rate

intracranial recordings. While this technology does not yet exist, recent technological improvements have brought it closer to realization - which is sufficient impetus to further investigate HFOs both as a temporal biomarker of epilepsy, and as a potentially powerful predictor of epileptic seizures.

Chapter IV: Signal Characteristics of Preictal High-Frequency Oscillations in Refractory Epilepsy

Abstract

Objective: High-frequency oscillations (HFOs) have become an important spatial biomarker of epileptic tissue and epilepsy in general. HFO signal features vary by patient, and recent evidence shows that HFO rates change in time prior to seizure onset; both are findings with relevance to patient-specific seizure prediction. Still, HFO signal features have never been evaluated in the context of seizure prediction. Here we analyze time-varying properties of both HFO rates and signal features to fully investigate the consistency of the preictal HFO response in a large clinical cohort, and to understand additional contributions of signal features. We also highlight the practical application of HFOs to seizure prediction with an implementation of a seizure advisory system. **Methods:** We analyzed the HFOs of 27 patients with refractory epilepsy who were being evaluated for resective surgery. We characterized changes in HFO rates and signal features over time, and assessed their preictal and interictal differences using cross-validated logistic regression models, whose predictive performance we compared the AUC metric. The implementation of our seizure advisory system used these models to generate seizure probabilities through continuous time, and applied two iteratively-determined probability thresholds to generate three discrete seizure warning levels (low, medium and high) in time. The performance of the advisory system was

assessed by comparing the percentage of time spent in each warning level, the percentage of seizures correctly identified in each warning level, as well as the average warning time until a seizure occurred. **Results:** There were 13 'responder' patients out of 27 with significant preictal HFO characteristics. Evidence that signal-based HFO characteristics could improve prediction performance was overall inconclusive; but for some patients the magnitude and consistency of prediction performance with HFO signal features was significantly increased. The performance of the seizure advisory system was within the range of several other notable prediction studies using this method. **Significance:** These findings further reinforce and expand evidence that HFOs are temporal biomarkers of seizure onset. They also demonstrate that HFO signal features can add meaningfully to prediction for some patients. While the seizure advisory system was presented chiefly as a proof-of-concept, its encouraging result in many patients is a powerful demonstration of the potential utility of HFOs within patient-specific seizure prediction.

Introduction

Approximately one third of patients with epilepsy do not respond to medication [23]. For these individuals, invasive resective surgery is often the only recourse to achieve lasting seizure freedom. Yet many still experience seizures after surgery, while others are not even candidates for resection [29]. Given their unmanaged seizures, these individuals represent a sizeable patient population whose quality of life could benefit significantly from an implanted therapeutic device capable of accurate seizure

prediction. Though these devices do not currently exist, their eventual clinical realization has been a driving motivation for many seizure prediction researchers, which has resulted in numerous milestone achievements over the past twenty years [104].

The field of seizure prediction is at a crossroads, however. With recent advancements in our understanding of seizure mechanisms, it is becoming more evident that seizure prediction in actual clinical practice will require patient- and seizure-specific adaptations [12], [62]. This results from the identification of individual-specific factors that can influence seizure mechanisms; endogenous examples include the individual variation of diurnal and multiday rhythms [101], the presence of multiple dynamical seizure types in a single individual [13], [103], and the phenomenon of seizure clustering [62]. So while the direction of future work in seizure prediction is clear, the technology to realize its clinical translation has lagged behind [62], as there is still only one chronic dataset currently available for prediction research, and it was created nearly ten years ago [66]. Perhaps more importantly, however, a reliable electrical biomarker of seizure onset has yet to be identified, let alone a more useful patient- and seizure-specific indicator [12], [62].

Seizure prediction, like clinical interpretation of EEGs, is dependent upon the available data and how they are interpreted. Historically, these data have been limited to EEG signals below 100 Hz. Most EEG devices sample in this range, and clinicians have learned to read EEG accordingly. In fact, one study demonstrated that higher resolution EEG did not change clinical interpretation at all: 100 Hz was clinically indistinguishable from 1000 Hz [118]. Nevertheless, these electrical biomarkers of

epilepsy have been extremely useful in diagnosing, managing and treating the disease [21]. However, the introduction of computational algorithms alongside high resolution hardware affords a new possibility for identifying signals that clinicians cannot see. The discovery of high-frequency oscillations (HFOs) - which occur more frequently in epileptic tissue - has aided surgical decision-making, [37] and has improved resective outcomes for some [39]. HFOs are now an established biomarker of epilepsy and epileptic tissue, but their clinical use was limited until only recently – when more sophisticated amplifiers capable of recording high-density low-noise EEG data were gradually installed at major centers around the world [42], [50]. Here, the clinical relevance of HFOs as a spatial biomarker was somewhat dependent on the technology that supported it. But given the mounting evidence that HFOs are highly correlated with epileptic tissue [43] along with the development of efficient automated detection algorithms and advancing hardware technology, in this study we argue the potential benefit of incorporating HFOs into the next generation of seizure prediction devices.

As justification for this argument, there is recent evidence that HFOs also function as temporal biomarkers of epilepsy and imminent seizure onset [109]. There is also evidence (presented in the study of Chapter III) that HFO rates could be useful in seizure prediction, especially for certain individuals. In addition to HFO rates, various signal features of HFO waveforms also have documented temporal aspects, as it has been shown that HFO signal features 1) vary in time and location over long periods, exhibiting patient-specific patterns [114], 2) differentiate ictal periods from other times [119], and 3) differentiate the dynamics of two distinct seizure onset types [72]. Though

this evidence suggests a potential role for HFO signal features as possible patient-specific and even seizure-specific temporal biomarkers, it is currently unknown if these measures are at all useful in seizure prediction. A thorough and robust evaluation of HFOs in this context, however, is not possible with current technology because no device exists that is capable of recording high-density intracranial EEG in a chronic ambulatory setting. Still, the identification of a reliable patient-specific pre-seizure biomarker would have great clinical potential.

In this study, we attempt the first comprehensive study on how HFO features – including their rates – change over time. Conceived as a significant expansion of the ideas and analyses presented in Chapter III, this study directly addresses the considerations raised above by evaluating two main objectives: 1) to investigate whether HFO signal features add meaningfully to both the identification of preictal states and seizure prediction performance, and 2) to provide a practical demonstration of the clinical potential HFOs could have in patient-specific seizure prediction – which could be vital to the continued development of this novel temporal biomarker.

Methods

Many of the methods used in this study are the same as those used previously in Chapter III. As such, brief explanations of common methods and data will be provided, but for further detail, please refer to the previous chapter.

Patient population and data: Both this study and the study presented in Chapter III use the same intracranial EEG dataset, which was recorded from the monitoring sessions of 27 patients with refractory epilepsy during the years 2016 - 2018. All associated patient metadata is also the same (found in Table 2), except that patient UMHS-0034 now has a class III surgical outcome (outcome definitions can be found in [120]). All patient metadata for both studies is found in Table 2.

Data analysis: Depicted in Figure 11, the general block process for the analysis presented in this paper differs from the previous study of Chapter III in only a few key steps. Whereas the analysis of the previous study used only HFO rate information to create predictive models (denoted hereafter as the 'HFO rate analysis') the predictive features in this study encompass both rate-based and signal-based HFO information (together referred to later as the 'full HFO analysis'). Further differences between the studies are described below.

TABLE 3: List of full HFO features. A list of HFO features descriptions and their abbreviations, as well as their meta-feature names, is provided below. Note that these features are computed independently for SOZ and OUT channel groups, which results in a designation of either XXX-SOZ or XXX-OUT for the meta-features listed in the last column.

	Feature description:	Abbreviation:	Meta-feature operator: Mean / Median / Std. Dev. / Skewness / Kurtosis
Rate-based features:	Number of HFO detections	**nDets**	**Mean-nDets, …**
	Inter-detection interval	**detIdi**	**Mean-detIdi, …**
Signal-based features:	HFO Duration	**dur**	**Mean-dur, …**
	HFO Amplitude * (dB)	**ampdB**	**Mean-ampdB, …**
	Skewness of waveform	**skew**	**Mean-skew, …**
	Kurtosis of waveform	**kurt**	**Mean-kurt, …**
Mean / Std. dev. / Skewness / Kurtosis::	First derivative of signal (rectified)	**meanD1, …**	**Mean-meanD1, …**
Mean / Std. dev. / Skewness / Kurtosis::	Second derivative of signal (rectified)	**meanD2, …**	**Mean-meanD2, …**

Mean / Std. dev. / Skewness / Kurtosis::	Teager-Kaiser energy	**meanTKE**	**Mean-meanTKE, …**
Power spectrum estimate:	Amplitude at peak frequency	**psePkAmp**	**Mean-psePkAmp, …**
Power spectrum estimate:	Peak frequency	**psePkFreq**	**Mean-psePkFreq, …**
Mean / Std. dev. / Skewness / Kurtosis::	Power spectrum estimate	**meanPSE**	**Mean-meanPSE**
Power spectrum estimate:	Frequency at 25% energy	**freqPSEnrgQ1**	**Mean-freqPSEnrgQ1**
Power spectrum estimate:	Frequency at 50% energy	**freqPSEnrgQ2**	**Mean-freqPSEnrgQ2**
Power spectrum estimate:	Frequency at 75% energy	**freqPSEnrgQ3**	**Mean-freqPSEnrgQ3**

Automated HFO detection and HFO feature computation: To begin, the set of HFOs used in this study was identical to those used in the previous study, which facilitated later comparison of results between the two. Prior to automated detection, a

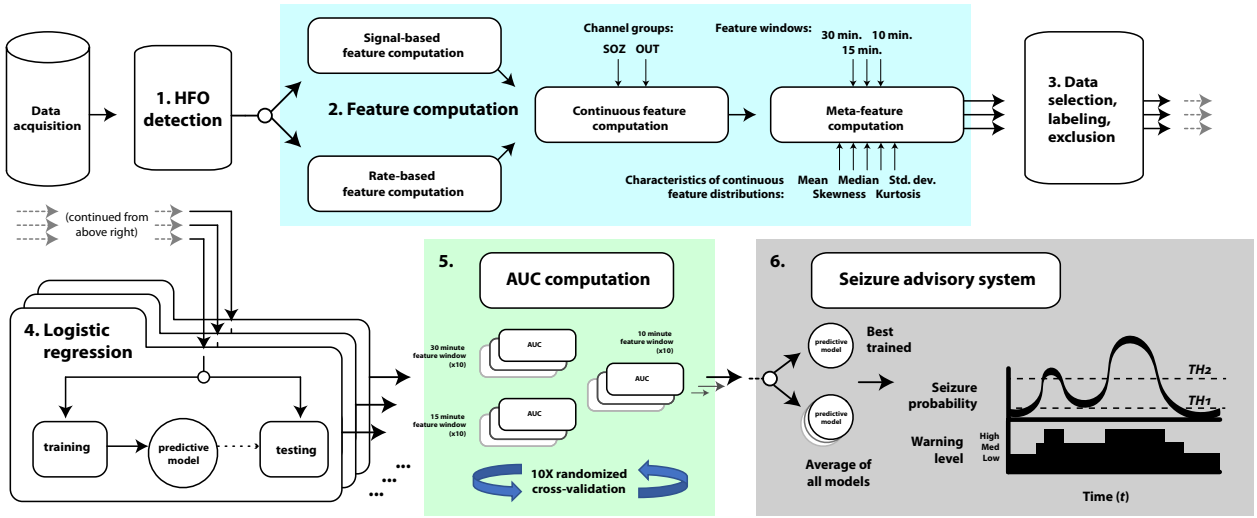

FIGURE 11: Schematic of overall data analysis workflow. After data acquisition, 1) HFO detection is performed. 2) HFO features are computed from newly detected HFOs: first, rate-based and feature-based computations take place, then continuous representations of these features are computed. Finally, meta-features are computed for each of the three different feature windows using the statistical operators of mean, median, standard deviation, skewness and kurtosis. These meta-features are also described in Table 3 above. After 3) data selection, labeling and exclusion, 4) logistic regression is performed with training and testing data over a series of 10 randomized cross-validation runs. 5) Performance comparisons of resulting logistic regression models are made with AUC computations. 6) Using the predictive models of step 4, the models with the best training AUC as well as an average model made from all cross-validation folds are used in the seizure advisory system. This system functions by first using these models to determine the probability of a seizure in time, and by applying two iteratively-determined thresholds TH1 and TH2, the warning level in time (which can be low, medium or high) is then determined from the seizure probability.

set of 25 signal features was assembled that we hypothesized would capture important temporal variations in HFOs. Listed in Table 3 above**,** these signal features encompass linear and non-linear measurements of individual HFO waveforms in both the time and frequency domains. The signal features of all HFO waveforms were computed at the time of HFO detection (a process indicated in Figure 11 as 'signal-based feature detection'); prior to this, each waveform was normalized to have unit amplitude, except for the computation of HFO amplitude ('*ampDb'* in Table 3), which was not normalized. The computation of the two additional rate-based HFO features is described in the next section.

Computation of continuous HFO signal features: HFOs are discrete events, but seizure prediction algorithms in practice are evaluated continuously in time. Thus, it was useful to transform HFO features into approximations of a continuous process. In this study, the computation of these 'continuous' features was identical to the method detailed in the HFO rate analysis of Chapter III. The only difference was the number of inputs to this computation; in the previous paper only HFO rate was used, whereas here the 25 signal-based HFO features were used (shown in Table 3), as well as two rate-based features (the number of detections '*nDets'*, and the inter-detection interval '*detIdi'*). As such, aggregating HFOs by channel group (either SOZ or OUT channel groups) resulted in 54 continuous features.

Computation of windowed meta-features: We evaluated temporal fluctuations in continuous HFO features by characterizing changes in their larger distributions in time across three different temporal horizons, or 'feature windows'; these were 30, 15 and 10

minutes in width. Though this computation was identical to methods of the previous study, here we used a slightly different set of statistical operators to compute the final 'meta-features' that would serve as inputs to logistic regression. Shown in Table 3, these operators consisted of the mean, median, standard deviation, skewness and kurtosis; they were applied individually to all continuous HFO features across the three independent feature windows. Overall, this computation resulted in a total of 270 'meta-features' per feature window.

Feature labeling and exclusion: In general, the data labeling and exclusion processes employed in this paper are identical to those detailed in the second paper: interictal labels were assigned in consecutive and non-overlapping intervals, while preictal observations consisted of the last feature window instance just prior to seizure onset.

Logistic regression model: All meta-features were used as predictor variables in logistic regression. The LASSO logistic regression procedures used to train the predictive models of this study were identical to those used in the rates analysis of Chapter III; for details on this specific implementation, please refer to the corresponding methods section of Chapter III.

Assessing predictive performance: We assessed the predictive performance of models trained by LASSO regression by averaging their corresponding test AUC values over all cross-validation runs. We also applied a bootstrap permutations test to all individual test AUC, which provided validation that observed values were statistically

different than the performance of a random predictor. Similarly to the test AUC, these bootstrap *p*-values were also averaged over all cross-validation runs using the harmonic mean, which has been recommended for averaging the *p*-values of multiple dependent tests [117]. As in Chapter III, we used nominal statistical criteria to determine 'responder' patients: these individuals had at least one feature window with average test AUC ≥ 0.60 and corresponding average bootstrap *p*-values < 0.05. Finally, we used non-parametric statistical tests to compare responder and non-responder prediction performance. We also used these tests to compare the predictive performance of this study's full-feature HFO analysis with the prior HFO rate analysis of Chapter III. Except when indicated, all statistical tests were corrected appropriately for multiple comparisons with the Bonferroni correction.

Implementation of seizure advisory system: The functionality of our advisory system is based on notable past work [66], [69]. Those systems were designed to alert the user of seizure risk through time by displaying one of three discrete warning levels: low, medium, and high risk. In a general sense, the system operates in a continuous fashion through time, and for a given moment, it assesses the probability that a seizure will occur before its next update ($P_{sz}(t)$). It then determines the corresponding warning level ($W_{sz}(t)$) to be displayed by applying two thresholds (Th_1, Th_2) to the current seizure probability. In prospective or pseudo-prospective prediction (which describes prospective prediction methods using previously obtained data) these thresholds are determined and updated dynamically in time, as new information about previous seizures is obtained. If the number of seizure observations are limited – as they are in

this study – these thresholds are determined only once after all data are collected, though this can potentially weaken an outcome's significance by introducing overfitting. Despite this, we chose the latter method of threshold evaluation for this study because of limited available data, and because of the original aim of demonstrating (and not proving) the potential of HFOs in practical seizure prediction.

Working within the constraints of this latter threshold evaluation method, we still wanted to portray a comparison between what could be considered the advisory system's best possible outcome and how it might actually perform in real-world settings. To show precedent in the literature for such a comparison, an analogous approach was used recently in [69].

In this study, the approach with the 'best possible outcome', (referred to as the average model) was evaluated with knowledge of all data and seizures. Specifically, the probability of a seizure in time ($P_{sz}(t)$) was derived from the output of a single 'average predictive model' per feature window; this was created simply by averaging the coefficients of individual models over all 10 cross-validation runs. As each cross-validation fold randomly selected training and testing data from the full temporal extent of a dataset, this 'average' model was thus likely trained on all seizures and data.

The model selected to show more realistic real-world performance was simply the model with the highest training AUC per feature window, (referred to as the best-trained model hereafter). This choice of training AUC did not bias the advisory outcome with knowledge of how the model actually tested on held-out seizures and data, and so,

these models were effectively trained on only 2/3 of all data and seizures. Except for associated warning thresholds, which were optimized over all data, these advisory outcomes include performance for held-out seizures, which more closely resembles pseudo-prospective seizure forecasting.

In the approaches of both average models and best-trained models, the model output used to form the seizure probabilities over the entire dataset was created by inputting all associated meta-feature data, which, for the purposes of the advisory system, were evaluated continuously in time in 1-minute intervals. These data were also subject to the same exclusion procedures as described before, though their computation did not require interictal and preictal labels. Once we determined the seizure likelihood in time as above, we iteratively optimized the warning thresholds (Th_1, Th_2) over a large range of values first by determining various performance metrics for each threshold pair (described further below). Optimal threshold pairs were then identified as those that satisfied the four following conditions:

C_1: maximize the number of seizures correctly identified in high warning

C_2: maximize the percentage of total time spent in low warning

C_3: the number of seizures identified per warning level must monotonically decrease going from high to low warning

C_4: the percentage of total time spent per warning level must monotonically decrease going from low to high warning.

Once the thresholds Th_1 and Th_2 were determined as above, the warning level at time $t-$ which would last for a duration of one minute until the next probability update – was given by the following:

$$W_{sz}(t, Th_1, Th_2) = \begin{cases} 0\ (\text{low}), & P_{sz}(t) < Th_1 \\ 1\ (\text{med}), & Th_1 \le P_{sz}(t) < Th_2 \\ 2\ (\text{high}) & Th_2 \le P_{sz}(t) \end{cases}.$$

We used a number of performance metrics to determine the warning thresholds as above, and to make final performance comparisons between patients and the two different advisory approaches. These included the number of seizures identified per warning level, where, for each seizure, the value of the warning level just prior to seizure onset was recorded; these values per warning level were then summed over the dataset to produce the total number of seizures identified per level. The percentage of total time per warning described in conditions C_2 and C_4 above was calculated without redacted data. Lastly, we also compared the average prediction horizon for seizures correctly identified in high warning (which would constitute a true positive detection), which we defined as the duration of time spent in high warning just prior to seizure onset averaged over only all true positive detections.

Results

Out of all patient models (810 in number), 387 successfully converged under LASSO regression (47.8%). Models that failed to converge with HFO features – either signal-based or rate-based – were conservatively assigned test AUC values of 0.5, and bootstrap test *p* values of 1. Of the 387 models that converged to a solution involving HFO signal features, there was significant variability in their predictive performance (HFO features: AUC test minimum - 0.032, maximum - 0.989, median - 0.651, interquartile range: 0.294).

Average test AUC: We identified models with minimally significant predictive performance by applying the following criteria: AUC test ≥ 0.6 and bootstrap *p* value < 0.05. This was also applied to average test AUC values as well as average bootstrap *p* values, which were averaged using their harmonic mean. These average test AUC values for each patient and feature window are shown in the bar plot of Figure 12. With the statistical criteria above, we identified 13 'responder' patients overall with significant predictive performance in at least one feature window; utilizing a 95% binomial confidence interval, this means that we could reasonably expect between 29.3% - 67.0% of a larger patient population to be responders.

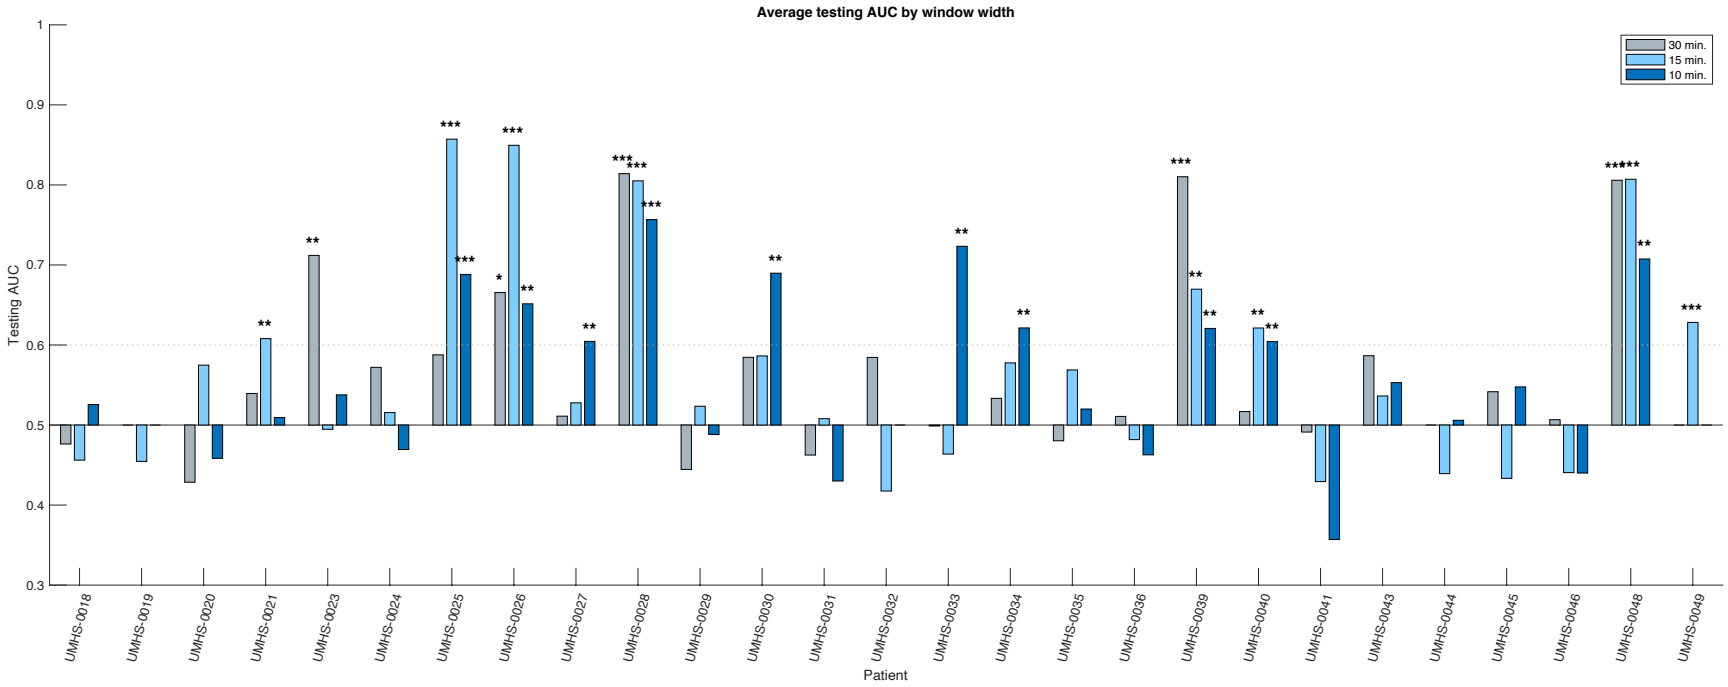

*FIGURE 12: Bar chart of average test AUC values by patient and feature window. 13 individual responder patients have significant predictive performance (average test AUC ≥ 0.6, significant average bootstrap test p-value < 0.05) in one or more feature windows. The statistical significance of the bootstrap test per feature window is indicated with asterisks: *, **, ***; P < 0.05, < 0.01, < 0.001 respectively. Note that the significance is based upon how likely that patient's data could produce the given AUC by random chance, not whether the magnitude of the AUC itself is high.*

Differences across feature windows: Considering the three different feature windows, there were 10 patients with significant average AUC in the 10 minute window, while 8 and 5 responders respectively had significant average AUC in the 15 and 30 minute feature windows. There were no statistically significant differences among feature windows that would suggest a particular window was more frequently significant than any other (Chi-square test: $P = 0.31$).

Comparing the performance of each feature window by the average AUC of significant responder models, the thirty minute window appeared to perform better than the 10 minute window (median AUC in 30 minute window: 0.806, 10 minute window: 0.670), however, this difference was not significant after correcting for multiple comparisons (2-sided rank-sum test, $P = 0.0280$). There were, however, significant differences in the significance of the bootstrap test *p*-values across different windows (Kruskal-Wallis: $P< 0.05$): *p*-values in the 15 minute window were significantly different (lower) from those in the 10 minute window (2-sided rank-sum test, $P<0.05$).

Comparing predictive performance of HFO features versus rates: We compared the predictive performance of this study's full HFO models with the results of models from the study in Chapter III. Considering either the entire patient cohort or even just the group of responder patients separately, there were no statistically significant differences or improvements in the magnitude (average test AUC) of prediction performance between the HFO rates analysis and the full HFO features analysis presented in this paper. Still, this study saw some improvement over Chapter III's

results in the bootstrap *p*-values that were used to assess statistical certainty of these average test AUC. This is described below.

To analyze how the full HFO features analysis could have improved the statistical certainty of better-than-chance AUC, we made a number of comparisons between the associated average *p*-values of the bootstrap test between the HFO rates analysis and the full HFO analysis presented here. Considering the entire cohort, median *p* values for the full HFO features analysis were qualitatively lower in the 15 minute feature window than for the rates analysis (Full HFO analysis: median bootstrap *p* value = 0.0025, HFO rate analysis: median bootstrap *p* value = 0.0190). We tested whether the full HFO features analysis could improve the statistical significance of patient-wise predictive performance when compared to previous HFO rates models. Applying a 1-sided sign rank test, bootstrap *p*-values for the full HFO analysis in the 15 minute feature window were nearly less than those for the HFO rates analysis, but this was not significant after correcting for multiple comparisons (Signed-rank right-tailed test: p=0.0235). Considering the performance of responder patients only (grouped independently by study), bootstrap *p*-values for the 15 minute feature window were again qualitatively lower than that of the rates analysis (Responders: Full HFO analysis: median bootstrap *p*-value = 0.00089, Responders: HFO rates analysis: median bootstrap *p*-value = 0.00138), yet statistical comparisons of these values were also not significant (Rank-sum right-tailed test: p=0.0652). Similar comparisons in other feature windows were not close to reaching statistical significance.

Important features: The choice of a machine learning tool like LASSO regression - rather than a black-box approach like deep learning - not only provides robust classification ability, but also yields an understanding of which features contribute the most to an observed trend, which potentially affords insight into underlying motivating mechanisms. Given the increased number of predictors in this study (270 in total) relative to the same number of seizures available as before, one of our primary concerns was model overfitting. To verify that the regularization of LASSO was functioning as intended, we analyzed on average how many variables were eliminated with LASSO to produce its final sparse models.

Considering first all models that converged to a solution using HFO signal features (387 in total), the number of non-zero coefficients (variables) used in these models was on average significantly less than the total possible number of variables available (Number of non-zero coefficients in all convergent models: 30-minute: 6.72 / 270, 15-minute: 6.96 / 270, 10-minute: 7.90 / 270) – this beneficially resulted in an average of 97.3% of all possible predictors being eliminated with LASSO in these models. Further, the average number of variables in significant responder models was approximately one non-zero coefficient less (Responders: number of non-zero coefficients: 30-minute feature window: 6.08 / 270, 15-minute: 5.84 / 270, 10-minute: 6.71 / 270).

One way to understand the importance of a given feature is to compare how frequently it was used across all models, that is, how often it was non-zero in models that converged. By this measure, the top 20 most important features to significant responder models are shown in Figure 13. As depicted in this figure, among the most

important features by frequency were those related to HFO rate (i.e. meta-features based on the number of detections – '*nDets*' , or the inter-detection interval – '*detIdi*'). In terms of HFO signal features, meta-features based on frequency-domain measures like the amplitude at peak frequency in the power spectrum estimate ('*psePkAmp*') were among the most frequent. As before with the rates analysis, features in OUT channels were among the most important by their frequency, especially for the 30 minute feature window, where, with the exception of '*Std-nDets-SOZ*' in 3rd place, the first eleven highest ranked features were computed from OUT channels.

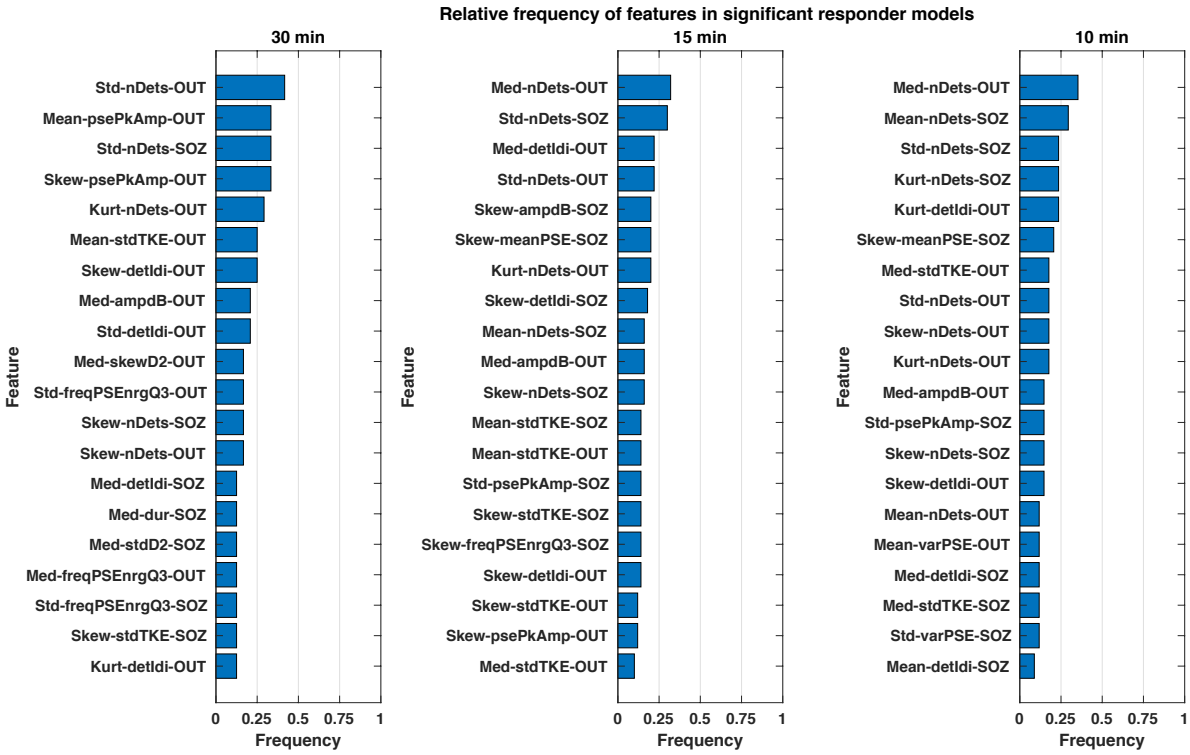

FIGURE 13: Relative frequency of features in significant responder models. The meta-features of HFO rate most important to discerning the preictal HFO response in responders are ranked in descending order (top to bottom) according to how often their respective model coefficients were non-zero for a given feature window. Overall, rate-based features such as the number of detections were seemingly most important; features in OUT channels were also significant, which is a novel finding that suggests channel location might not be vitally important in determining a preictal response in these responders.

We also assessed feature importance by directly comparing the magnitude of the coefficient values themselves. To make this possible, we first scaled '*nDet*' features by the appropriate number of channels in SOZ or OUT groups by patient. Next, we used the median absolute deviation of all coefficients across the cohort (per feature window) to rescale individual coefficient values to unitless quantities. Shown in Figure 14, we compared these scaled values in significant responder models, and ranked their overall importance by feature window with their absolute value. As before with feature frequency, the number of HFO detections (nDets) was ranked highly, while signal-based

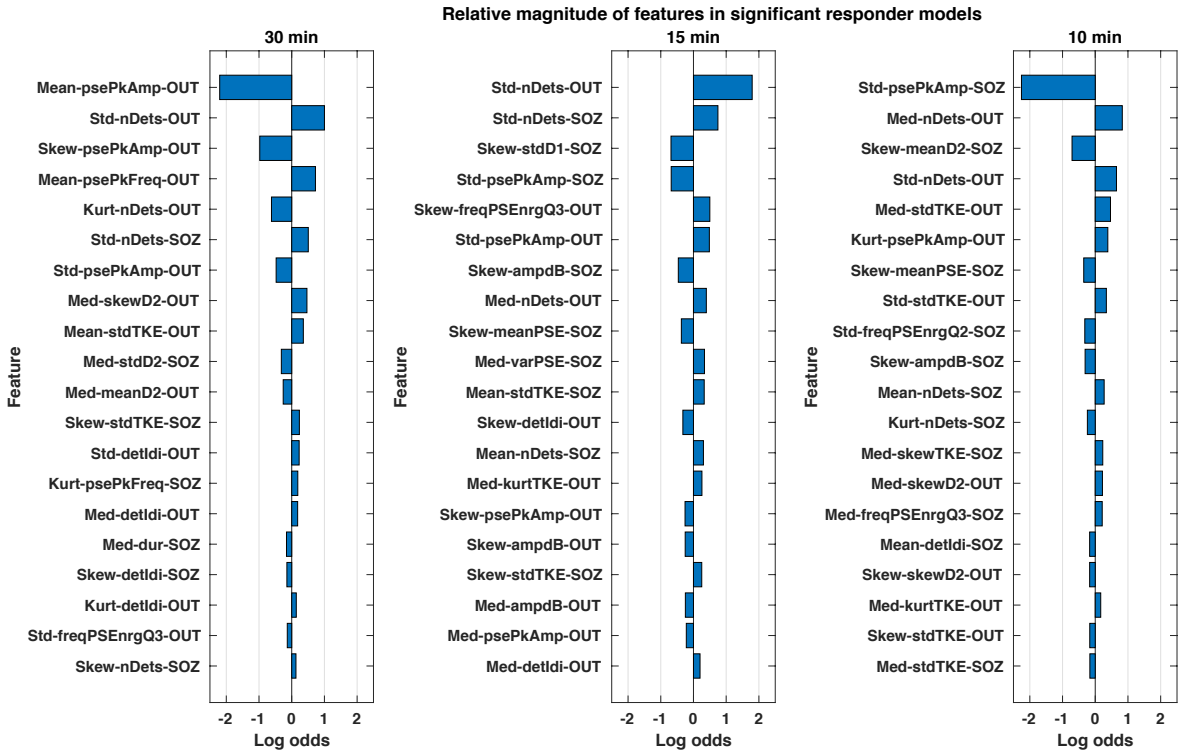

FIGURE 14: Relative magnitude of important features in significant responder models. The magnitude of meta-features of HFOs most important to discerning the preictal HFO response in responders are ranked in descending order of their absolute value (top to bottom). One signal feature above all others was most important: the amplitude at peak frequency of the power spectrum estimate (psePkAmp). Rate-based meta-features were also important, as before with relative feature frequency. Overall, OUT-channel meta-features were significantly important again, especially dominating higher rankings of the 30-minute feature window.

meta-features involving the amplitude at peak frequency of the power spectrum estimate of HFO waveforms ('*psePkAmp*') in particular were the highest-ranked signal feature in general. The importance of features in OUT channels was further reinforced, especially again for the 30 minute feature window, where OUT channel meta-features were top-ranked in the first five positions.

Seizure advisory system results: The purpose of our seizure advisory system implementation was to translate the somewhat abstract conclusions of the average AUC results to a more meaningful and practical demonstration of their potential relevance to a physically realizable system. Similar to [69], we implemented the advisory system in two different ways so that it would give 1) an idea of possible predictive performance, versus 2) the actual predictive performance of the system when tested on held-out data. In Figure 15, a detailed example of the system's output for patient UMHS-0028 for both

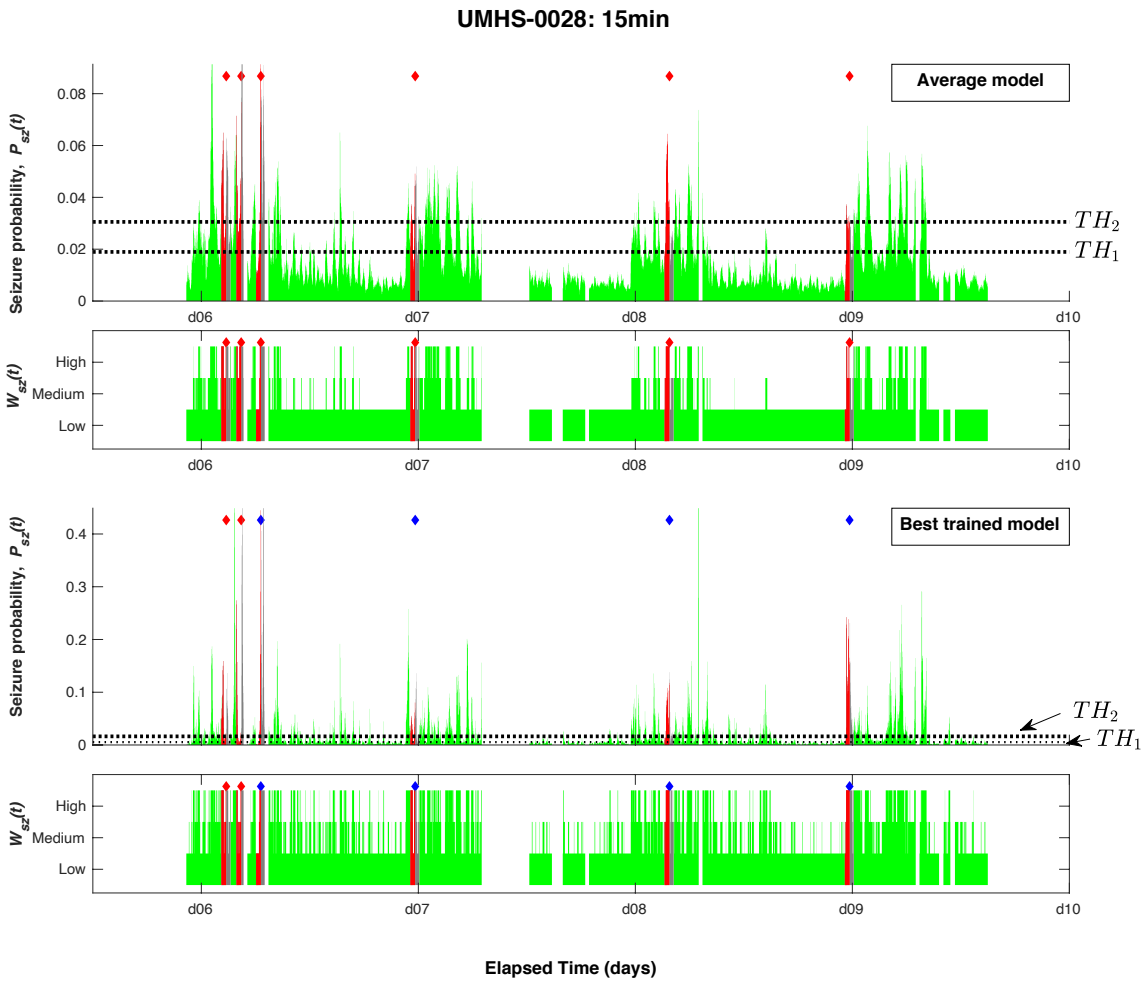

FIGURE 15: Example of seizure advisory system output. A detailed visualization of the seizure advisory system's output is shown for patient UMHS-0028, a responder with significant average AUC in all three feature windows. Seizures are denoted with red diamonds, while preictal periods are visualized with red bars. Average model : The second plot row represents the warning system's output for the average model of UMHS-0028 in the 15-minute feature window, which was created by optimizing the thresholds TH1 and TH2 in the seizure probability plot above this in the first row, where these probabilities were determined by the average model. The same plots are shown in the third and fourth rows, Best-trained model: where seizure probabilities were determined instead by the best trained model. In the third and fourth plot rows specifically, seizures trained on are represented by blue diamonds, while test seizures that were held-out are denoted with red diamonds. In this patient, the average model outperformed the best-trained model; both had similar sensitivity but the average model had higher specificity (fewer false positives).

average and best trained models is visualized, while advisory output for average and best-trained models for all responder models is shown in Figure 16.

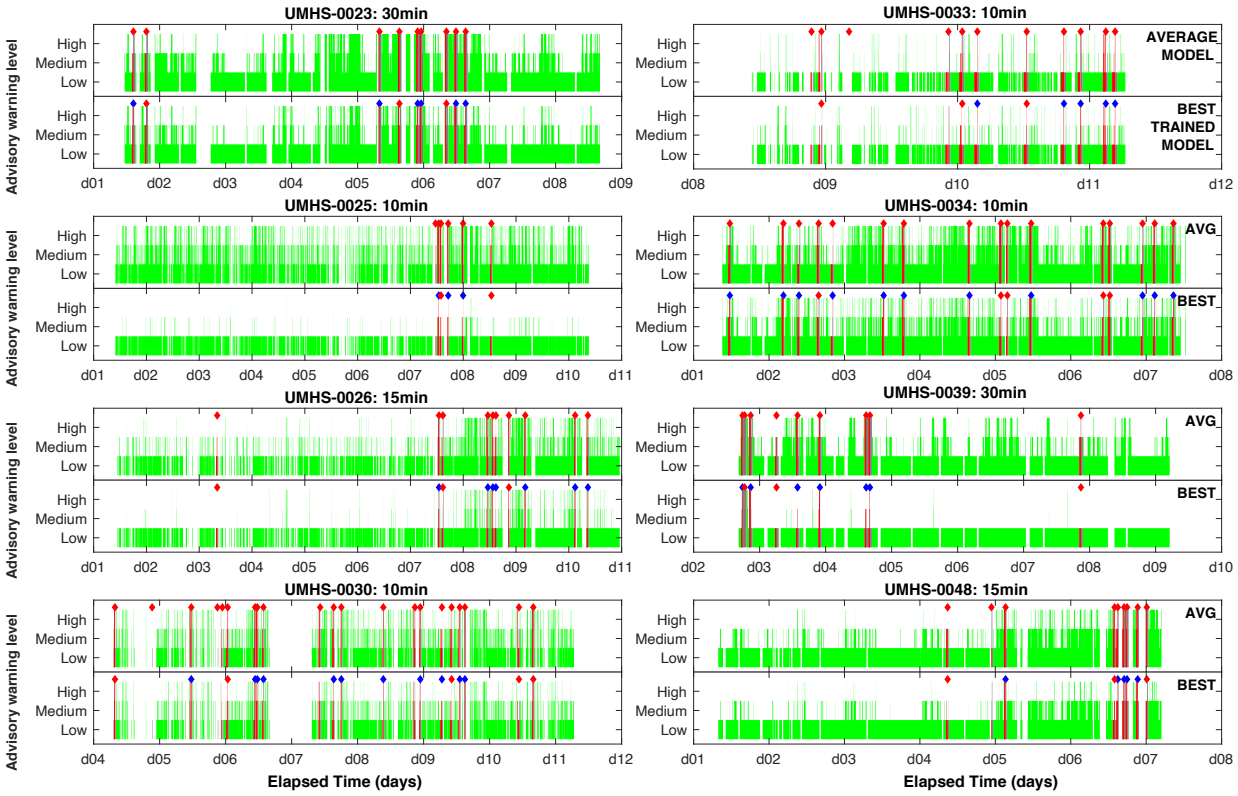

FIGURE 16: Advisory system output for select responder windows. Similar to Figure 15, visualized above are the advisory system's results for several responder patients and various feature windows. Note that these patients include any responder whose average and best models converged to an optimum pair of threshold values; for some responders however, the thresholds did not converge and that is why such results are not presented above. Overall, in some patients the average model performed better, while in others the best-trained model performed just as well, though with higher specificity. There were very few instances, however, of the sensitivity of the best-trained models exceeding the sensitivity of average models. Again, for average models, red diamonds represent only seizures while for best-trained models, blue diamonds represent training seizures, while red diamonds represent test seizures that were held out in training.

In Table 4 below, we compare the advisory performance of average and best-trained models for significant responder models. For average responder models, the median percentage of time spent in high warning ranged from 13.6% (15 minute window) to 17.2% (10 minute feature window). The median performance of the best trained responder models was even better considering this metric – ranging from 3.3% (30min) - 5.8% (15 min). Next, considering the percentage of seizures correctly identified for each warning level, the median values of average models were slightly higher (qualitatively) than the best trained responder models; these values ranged from 72.7% (10min.) to 88.9% (30min) (Best-trained models: 10min.: 56.0% 30min.: 80%). Considering seizures identified in low warning (which are essentially false negatives), the average models performed better than the best trained models, with 0% false negatives for all feature windows. Finally, we consider the average warning time in high level before seizure onset occurs (i.e. its prediction horizon). In average responder

TABLE 4: Responder advisory performance comparisons. Median advisory performance for average and best-trained models in all feature windows is reported below for the three performance metrics compared: the percentage of time spent in a level, the percentage of seizures correctly identified per level, and the average prediction horizon in high warning. In general, the sensitivity of average models was higher than best-trained models, while this was reversed for prediction specificity. Average prediction horizons were significantly longer in average models. (P< 0.05, 30 and 10 min. features windows).

Significant responder models: median performance

Feature Window:

	30min			15min			10min		
Average model:									
	Time in level (%)	Seizures identified (%)	Avg. prediction horizon (min.)	Time in level (%)	Seizures identified (%)	Avg. prediction horizon (min.)	Time in level (%)	Seizures identified (%)	Avg. prediction horizon (min.)
High	13.7%	88.9%	15.2	13.6%	87.3%	11.0	17.2%	72.7%	7.8
Med	25.5%	11.1%	-	25.0%	10.6%	-	30.1%	10.0%	-
Low	55.3%	0.0%	-	62.1%	0.0%	-	52.2%	0.0%	-
Best trained model:									
W_{SZ}	Time in level (%)	Seizures identified (%)	Avg. prediction horizon (min.)	Time in level (%)	Seizures identified (%)	Avg. prediction horizon (min.)	Time in level (%)	Seizures identified (%)	Avg. prediction horizon (min.)
High	3.3%	80.0%	10.6	5.8%	75.0%	8.6	5.0%	56.3%	3.3
Med	18.9%	20.0%	-	28.0%	20.0%	-	21.4%	16.7%	-
Low	77.8%	0.0%	-	58.3%	5.0%	-	75.7%	10.6%	-

models, this ranged from 7.8 minutes (10minute feature window), to 15.2 minutes (30 minute feature window). For the best trained models, these values had a range of 3.3 to 10.6 minutes (10, 30 minute feature windows); these values for average models were significantly longer in duration than those for the best trained models (Rank-sum test: 30 min: $P < 0.05$, 10 min: $P < 0.05$).

Discussion

We analyzed the temporal characteristics of HFO signal features to evaluate their potential contribution to enhanced seizure prediction performance relative to that of HFO rates alone. In general, the data and results presented in this study offer further supporting evidence that HFOs can act as temporal biomarkers of epilepsy.

Comparing HFO rates with features: The number of responder patients increased by three individuals in this study to a total of 13. Despite these additional responders, however, the use of HFO signal features did not result in significant systematic cohort-wide improvements in seizure prediction performance, as reflected in the average test AUC metric. Yet, there were qualitative indications - for the entire patient population and separately for responder patients - that the certainty or confidence of this test performance against random-chance prediction was increased by HFO signal features. While narrowly missing statistical significance, comparisons of this model certainty with that of HFO rates nevertheless demonstrated that HFO signal features have the potential to improve HFO-based seizure prediction. For certain responder individuals

like patient UMHS-0028, the increases in prediction performance resulting from the use of HFO signal features were considerable.

Despite the limitations of this study's short-term intracranial monitoring data, the predictive performance of our HFO-based classification algorithm for certain patient subsets was reasonably within the range of several notable seizure-prediction studies that each used the more ideal chronic Neurovista datasets. The top-ranked classification algorithm in the Kaggle seizure prediction contest of 2016, for instance, achieved an overall AUC of 0.81, which reduced to 0.76 when tested with held-out data [64]. Prediction studies conducted since the 2016 competition have achieved similar performance, including one where the range of average AUC by patient was 0.69 - 0.90 [65]. Given the differences in the data and methods used, it is difficult to directly compare our results with these studies. Still, there were five responder patients with average test AUC values over 0.80 - which suggests that HFO-based seizure prediction performs similarly.

As in our rates analysis of Chapter III, we also analyzed which HFO features were the most important to identifying preictal periods. In responder models, the most common features tended to be rate-based, while a varied mix of linear and nonlinear signal features in both time and frequency domains were also common. Ranked by the absolute value of their magnitude, the features that contributed significantly to identifying the preictal state in responders were also rate-based, but one signal feature - the amplitude at peak frequency of the power spectrum estimate (abbreviated '*psePkAmp*') - emerged as top-ranked. As in our rates analysis of Chapter III, meta-

features computed across OUT channels were highly ranked, which again suggests that HFOs outside the SOZ might be linked to preictal states – this has limited but intriguing support in the literature [121].

With our implementation of the seizure advisory system, we put the positive predictive performance of responder models into the context of 'pseudo-prospective' seizure prediction. For average and best-trained responder models, the percentage of time spent in high warning was within the performance reported by the one true prospective study (which, for their responder patients, ranged from 30% - 3%) [66]. Next considering the percentage of seizures correctly identified for each warning level, all values reported for average and best trained responder models were within range of the same values reported in a recent pseudo-prospective study using the NeuroVista data [69]. In general description, the average models had better sensitivity, and were better at identifying seizures in high warning when compared to the best trained models. Conversely, the best trained models tended to be more sensitive, with higher percentages of time spent in low warning and less false positives. These same relationships between average and best trained models were also similarly reflected in the performance of the two evaluation methods detailed in [69] – the performance of their advisory system was more sensitive when trained with all data, and it was more specific when the probabilities were evaluated pseudo-prospectively (or, without knowledge of some seizures, as was emulated in our best- trained method).

These encouraging results were achieved with more limited data. In contrast to the study discussed above [69], however, it is important to note here that our advisory

system cannot truly be considered pseudo-prospective, as the predictive models used to generate the probability of a seizure in time were those that had been either cross-validated over the entire dataset (as in the average models) or those that were trained on seizures that could have originated from any particular time within the dataset; and thus the division of data for another held-out test set was unrealistic given that the number of seizure observations was already very small. As such, the results of our advisory system likely represent a 'best possible' outcome (average models) or a more realistic outcome (best trained models) because training data (along with test data) generated the system's output - and the corresponding average *training* AUC from these data was often extremely high (near 0.90 or above) for many patients (data not shown).

Still, this approach of using the full extent of available data in time was recently used in another seizure prediction study based upon the chronic NeuroVista data [69]. This 'best-possible' method was presented as a way to discern an upper bound on predictive performance by patient - which is essentially a model's predictive potential. This bound was then compared with the more realistic performance of a true pseudo-prospective prediction algorithm, which allowed the authors to estimate potential gains in predictive performance that future work could mine with the aid of better methods or more ideal datasets. Here, we emulate these methods, because we have limited data, but still must show the value of HFOs in seizure prediction. Given that the requisite high-density iEEG dataset of sufficient duration does not currently exist, the results of

our seizure advisory system thus stand as a proof-of-concept only and are not intended as validated evidence of actual pseudo-prospective seizure prediction performance.

Another concern with the methods of this study could be that of overfitting, which can result when too many variables are used to characterize too few observations. Though the number of predictors in this study increased to a total of 270, we attempted to account for this increased number of variables by the use of LASSO methods: the resulting reduction in variables was evident considering that only an average of roughly 7 out of 270 variables were used in most predictive models. Future work using HFO signal features could better account for the reduced number of seizure observations relative to high numbers of predictors and interictal baseline data by using more advanced forms of dimensionality reduction, or even by using unsupervised prediction approaches that characterize excursions in the data without pre-defined labels of what 'is' or 'is not' a seizure.

Conclusion: The findings of this paper add further support to the notion that HFOs are temporal biomarkers of seizure onset in refractory epilepsy. These results also indicate that the use of many signal-based features might not justify the computational cost of including them in future device architectures, as HFO-based seizure prediction using only HFO rates performed similarly. Overall, these findings are powerful evidence for the development of chronic datasets that support high-density intracranial EEG recordings, which will enable future work on HFO-based seizure prediction.

Chapter V: Discussion and Conclusion

This thesis represents a significant contribution to our understanding of high-frequency oscillations in refractory epilepsy. Enabled by the use of a large clinical cohort, as well as methodological advances in peri-ictal HFO processing, each of the studies in the preceding chapters successively built and developed a framework of evidence to support the novel idea that high-frequency oscillations can act as temporal biomarkers of epilepsy and seizure onset. Overall, the strength of this temporal association with preictal periods was variable among patients, a finding likely reflective of the diverse etiologies of clinical refractory epilepsy. In certain 'responder' patient subsets, however, this preictal HFO signal was clearly identifiable and consistent. In a first-of-its-kind analysis, this thesis also confirmed that it is possible to use the temporal properties of HFOs in seizure prediction; again for certain 'responder' patients, the performance of HFO seizure prediction algorithms was robust and significant.

This thesis also represents methodological advances related to HFOs and seizure prediction. Throughout the three studies presented, a number of novel approaches were introduced to identify, process, and analyze HFOs near seizures. First, HFOs that resembled muscle activity occurring near the surface of the scalp were redacted as artifacts; this increased HFO specificity, especially near or during seizures, given that ictal muscle contractions could significantly impact automated detection. These three studies are the first to use this technique to improve HFO detection near seizures.

Second, the problem of calculating a baseline threshold for HFO detection near seizures was addressed by modifying how data were fed to the automated detector, which prevented high-amplitude seizure activity from arbitrarily influencing this threshold for detection. Lastly, in order to analyze changes in HFOs with sufficient resolution in time, HFO characteristics (i.e. their rate or signal features) were transformed into novel continuous measures, which also made these quantities more useful in the context of seizure prediction. In terms of seizure prediction, specific limitations in iEEG monitoring data related to its duration or temporal volatility were addressed by randomly partitioning training and testing observations in time, instead of in chronological order. While the use of random data sampling is certainly not novel, its intended purpose for use in seizure prediction with limited hospital monitoring data can be considered so.

In general, the use of these methods would not have been possible without first the creation of a flexible framework of computer code that could accommodate the messy clinical reality of a large epilepsy database. The upfront investment in time and effort for the creation of this framework was significant, but this was necessary in order to overcome the technical challenges of analyzing millions of very short events occurring across 190 days of high-resolution intracranial EEG data, which were recorded from 30 different patients in an epilepsy database with over 50 terabytes of total data.

Overall the central goal of this thesis was to evaluate temporal changes in HFOs to see whether they could differentiate the preictal brain state from other interictal times. The study in Chapter II first assessed whether this preictal effect was at all present in individuals, and if so, how prevalent it was across a large clinical cohort. First at an

individual level, a potential stereotypical response was evaluated by comparing preictal and interictal data trends in aggregate. The extent and prevalence of this preictal effect across the entire cohort was then determined with population-level inference. There were no population-wide differences in preictal and interictal HFO rates when averaged over long periods of time. Using a novel continuous transform of HFO rates, however, we identified significant population-wide linearly increasing trends in preictal HFO rates occurring within epileptic channels that were absent during interictal times. Examining the population distribution of these temporal trends, we further identified a 'responder' subset of 10 patients, each with significantly increased preictal HFO trends relative to others in the cohort.

These findings are supported by earlier preliminary findings as documented in [56] and [55]. Other studies that investigated high-frequency activity before seizures also found evidence of stereotyped preictal changes for certain individuals [74], [99]– which lends support to our identification and separation of the 'responder' patient subset.

In addition to the limitations mentioned in the discussion section of Chapter II, the manner in which data were aggregated for this study's main analysis – which was grouped first by channel group, then by preictal or interictal segments, then at the population level – bears some additional consideration. In statistics, repeated measurements taken from the same statistical subject – either at different points in time or for multiple factors or variables – are considered dependent. Thus, in order to apply population-level statistical inference to these measurements, this dependence must in some way be accounted for. Ideally, the manner in which these measurements covary

is known. Considering the nature of the HFO data in this study – which originated from many different patients with different epileptic foci, pathologies, sleep habits, medication protocols, electrode types and configurations, and possibly different seizure types – it was apparent from the variability of the findings from patient to patient and even between different seizures of the same individual that it would be impossible to fully specify this covariance for all patients.

Still, we utilized two small adjustments in our analysis to ameliorate the influence of these unknown covariates. First, we assumed that channels from the same group (either SOZ or OUT channels) would covary similarly to one another in order to reduce the covariance between individual channels. Whenever possible, we performed paired statistical tests like the signed-rank test – which are better suited to repeated measures designs [122]. These pairwise comparisons within patients were implemented to reduce covariance between individuals when assessing population-level significance. Though we did not attempt it here or in the other studies of this thesis, future work comparing population-level HFO data in aggregate would benefit from fully specifying the covariance between all data observations – though this would potentially require much more data than might be available with hospital monitoring sessions. Another potentially more feasible approach is to use a statistical method – such as the method of generalized estimating equations [123] – that robustly adjusts estimated data by its covariance, even when the structure of this covariance is unknown.

The studies presented in the third and fourth chapters extend the methods and findings of the first study in significant ways. While the first study was centered on

population-level comparisons of HFO trends for its analysis, these studies focus on patient-specific analyses in order to demonstrate their relevance to seizure prediction. To further facilitate their intended context in seizure prediction, they compare preictal and interictal data segments individually, rather than in aggregate, as was used in the first study.

The analysis of the third chapter characterized temporal changes in continuous HFO rates in a similar manner as the first study. Instead of comparing these values using population inference, however, the study in Chapter III used them to create logistic classification models that could be used in seizure prediction. The overall process of training these predictive models and then testing them with new held-out data was cross-validated over ten randomized runs, which allowed for a more realistic and consistent assessment of their ability to differentiate preictal and interictal data. We used a minimal statistical benchmark to differentiate patient models with better-than-chance predictive ability; applying this criteria to the average test AUC over all patients resulted in the identification of 10/27 ‘responder’ patients, who each had a significant result in at least one of the feature windows.

The methodology of the third study in Chapter IV expanded to incorporate additional information about HFO signal features, which we hypothesized could improve seizure prediction, based on the usefulness of HFO features in other studies. Other than the use of these additional features, the overall methodology of the second and third studies was the same. This allowed us to directly test our hypothesis that the use of HFO signal features might aid prediction performance. The results from this study identified a total

of 13 responder patients – a gain of three responders over the prior study. Despite these additional responders, there was no statistical evidence that the use of expanded HFO features resulted in higher average test AUC, though the statistical performance of these models as reflected in their average bootstrap *p* values was improved for some comparisons. Overall however, this is actually a positive result because it suggests that using only HFO rate – which is easier to compute – might be sufficient for accurate seizure prediction; this would be especially useful for the translation of our algorithm into a physical device.

There is no direct evidence in the literature to support the findings of the second and third papers, because to our knowledge there are no studies that use HFOs in seizure prediction. Considering both studies together however, the average test AUC of responder patients was well within the range of results reported in many other prediction studies in the literature, [63]-[65] and many of these studies had the advantage of the much longer Neurovista dataset, while our studies did not.

Specifically considering the seizure advisory system of the third paper, overall, the performance of the average model in responders was quite high, and was within range of other prediction studies using the advisory system method [66], [69]. The performance of our advisory system for more realistic models that were not trained on all seizures was comparatively more variable and reduced. Still, some patients achieved excellent results despite their models not being trained on all seizures. Considered altogether, the reasonable performance of the seizure advisory system for some

patients further demonstrates the potential that HFOs could have in seizure prediction – provided that the right data is available for such continued work.

The studies presented in the second and third chapters do have some inherent limitations. The primary limitation for both is the required use of relatively short datasets that are characterized by many variable and non-stationary exogenous factors. These data are in no way representative of the normal ambulatory setting that a seizure prediction device would eventually operate in [12]. As such, the number of seizures captured per patient is generally far less than would be minimally required for patient-specific and seizure-specific prediction. This results in a significant imbalance between the number of preictal and interictal observations – and this imbalance is made worse when data are evaluated continuously through time, as is the case with these studies. This imbalance can result in overtrained models that perfectly identify certain seizures while effectively ignoring others. In both studies, however, we attempted to address potential overfitting by minimizing cross-validation error during model training.

The limited number of seizures recorded per patient is also problematic given the large number of predictive features used in both studies, especially in the third. This imbalance of (many) predictors and (few) observations can also result in model overfitting. This consideration was a primary reason for our use of LASSO regression in both studies; given that approximately 97% of variables were eliminated with LASSO in convergent models of the full HFO analysis of Chapter IV, the imbalance of variables to observations in this case was significantly reduced. Even with this reduction, however, the possibility for model overfitting existed – which is why only testing performance (i.e.

model performance on held-out data) was reported for both studies, which was also averaged across all ten randomized cross-validation runs to better inform the consistency of this performance in patients.

In both studies, the last major caveat to consider is also derived from limited number of recorded seizures per patient. The AUC performance metric has been a standard tool used by many noteworthy prediction studies [63]–[65], [68] to describe the overall predictive performance of their seizure prediction models. While the use of this single value as a performance metric has many advantages, the imbalance of preictal and interictal data can also complicate the interpretation of associated AUC by overstating the specificity of a model [62], which, for these studies, is how often interictal segments were correctly identified. Additional limitations to our use of AUC in these studies also occur because the ratio of interictal to preictal data was not constant across patients and feature windows; this was a design choice that was necessary because of limited data in many patients. Finally, given the eventual clinical implications of online seizure prediction or forecasting algorithms to the patients themselves, it is unlikely that an abstract performance metric like AUC would hold any meaning to patients or clinicians alike when trying to optimize an algorithm's efficacy. This is in part the reason for concluding the third and final study with the seizure advisory demonstration, for if seizure prediction and forecasting algorithms and their results are not interpretable, there is likely little chance they will be useful to patients and clinicians – which could further stunt their widespread adoption.

The studies presented in this thesis form a foundation of preliminary evidence that support HFOs as a temporal biomarker in epilepsy. Given the limited data that were required for this work, future studies could validate these principal findings in different high-density datasets recorded from many clinically diverse patients. Ideally, this validation would also systematically account for the significant variability of data between different recorded channels, patients, and even between seizures. Using a robust statistical tool like the generalized estimating equations in this regard could better inform the statistical conclusions of population-level HFO analyses (such as those that were conducted in the first study). This could in turn identify certain clinical factors that might significantly predispose an individual to being a preictal HFO 'responder'.

In the field of seizure prediction, future studies using HFOs in this manner will likely require high-density intracranial EEG datasets that are ambulatory and much longer, containing many more seizures. If such data existed, future work could rigorously investigate HFO-based seizure prediction in the context of patient-specific and even seizure-specific algorithms. Such work could use sophisticated non-linear machine learning classifiers to finally tease out differences in preictal HFOs by seizure type. Also possible in this future work could be the use of unsupervised machine learning techniques which have also been used in health diagnostic studies [124], [125]; these methods are not as biased as supervised methods such as logistic regression because they abandon notions of preictal and interictal periods and instead characterize outliers in the data regardless of what or when such excursions are or occur.

Conclusion

The work of this thesis has identified and developed a novel pre-seizure biomarker that has significant clinical potential for patients with drug-resistant epilepsy. This work confirms that high-frequency oscillations are temporal biomarkers of seizure onset, and shows further that HFOs can be used in seizure prediction, especially for certain individuals. While there are likely many challenges in the winding road toward seizure freedom for many, it is hoped that the work of this thesis eventually contributes in some small way to the gathering global effort to reduce and eliminate the influence of seizures on the everyday lives of those suffering from epilepsy.

References

[1] J. A. Hobson, “REM sleep and dreaming: Towards a theory of protoconsciousness,” *Nat. Rev. Neurosci.*, vol. 10, no. 11, pp. 803–814, 2009.

[2] E. F. Pace-Schott and J. A. Hobson, “The neurobiology of sleep: genetics, cellular physiology and subcortical networks.,” *Nat. Rev. Neurosci.*, vol. 3, no. 8, pp. 591–605, Aug. 2002.

[3] S. Weinstein, “Seizures and epilepsy: An overview,” *Epilepsy Intersect. Neurosci. Biol. Math. Eng. Phys.*, pp. 65–77, 2016.

[4] D. Minecan, A. Natarajan, M. Marzec, and B. Malow, “Relationship of epileptic seizures to sleep stage and sleep depth,” *Sleep*, vol. 25, no. 8, pp. 899–904, 2002.

[5] N. Hitiris, R. Mohanraj, J. Norrie, and M. J. Brodie, “Mortality in epilepsy,” *Epilepsy Behav.*, vol. 10, no. 3, pp. 363–376, 2007.

[6] P. T. Fernandes, D. A. Snape, R. G. Beran, and A. Jacoby, “Epilepsy stigma: What do we know and where next?,” *Epilepsy Behav.*, vol. 22, no. 1, pp. 55–62, 2011.

[7] M. Bishop and C. A. Allen, “The impact of epilepsy on quality of life: A qualitative analysis,” *Epilepsy Behav.*, vol. 4, no. 3, pp. 226–233, 2003.

[8] E. Dwyer, “Stigma and epilepsy.,” *Trans. Stud. Coll. Physicians Phila.*, vol. 13, no. 4, pp. 387–410, 1991.

[9] R. Nickel *et al.*, “Quality of life issues and occupational performance of persons with epilepsy,” *Arq. Neuropsiquiatr.*, vol. 70, no. 2, pp. 140–144, 2012.

[10] R. Mahrer-Imhof *et al.*, “Quality of life in adult patients with epilepsy and their family members,” *Seizure*, vol. 22, no. 2, pp. 128–135, 2013.

[11] R. S. Fisher, “Epilepsy from the Patient’s Perspective: Review of Results of a Community-Based Survey,” *Epilepsy Behav.*, vol. 1, no. 4, pp. 9–14, 2000.

[12] D. R. Freestone, P. J. Karoly, and M. J. Cook, “A forward-looking review of seizure prediction,” *Curr. Opin. Neurol.*, vol. 30, no. 2, pp. 167–173, 2017.

[13] V. K. Jirsa, W. C. Stacey, P. P. Quilichini, A. I. Ivanov, and C. Bernard, “On the nature of seizure dynamics.,” *Brain a J. Neurol.*, vol. 137, pp. 2210–2230, Jun. 2014.

[14] L. Parker, M. Padilla, Y. Du, K. Dong, and M. A. Tanouye, “Drosophila as a model for epilepsy: bss is a gain-of-function mutation in the para sodium channel gene that leads to seizures,” *Genetics*, vol. 187, no. 2, pp. 523–534, 2011.

[15] S. J. Cho *et al.*, “Zebrafish as an animal model in epilepsy studies with multichannel EEG recordings,” *Sci. Rep.*, vol. 7, no. 1, pp. 1–10, 2017.

[16] J. Gotman, “A few thoughts on ‘ What is a seizure ?,’” vol. 22, no. Suppl 1, pp. 2011–2014, 2013.

[17] P. N. Banerjee, D. Filippi, and W. A. Hauser, “The descriptive epidemiology of epilepsy - a review,” *Epilepsy Res.*, vol. 85, no. 1, pp. 31–45, 2009.

[18] et al. Fiest, Kirsten M., “Prevalence and incidence of epilepsy: A systematic review and meta-analysis,” *Neurology*, vol. 88, no. 3, pp. 296–303, 2016.

[19] W. A. Hauser, S. S. Rich, J. Lee, J. F. Annegers, and V. E. Anderson, “Risk of Recurrent Seizures After Two Unprovoked Seizures,” *N. Engl. J. Med.*, vol. 338, no. 7, pp. 429–434, 1998.

[20] S. Sharma and V. Dixit, “Epilepsy – A Comprehensive Review,” *Int. J. Pharma Res. Rev.*, vol. 2, no. 12, pp. 61–80, 2013.

[21] R. J. Staba, M. Stead, and G. A. Worrell, “Electrophysiological Biomarkers of Epilepsy,” *Neurotherapeutics*, vol. 11, no. 2, pp. 334–346, 2014.

[22] R. L. Macdonald and K. M. Kelly, “Antiepileptic Drug Mechanisms of Action,” *Epilepsia*, vol. 34, pp. S1–S8, 1993.

[23] P. Kwan and M. J. Brodie, “Early Identification of Refractory Epilepsy,” *N. Engl. J. Med.*, vol. 342, no. 5, pp. 314–319, 2000.

[24] M. J. Morrell and C. Halpern, “Responsive Direct Brain Stimulation for Epilepsy,” *Neurosurg. Clin. N. Am.*, vol. 27, no. 1, pp. 111–121, 2016.

[25] G. K. Bergey *et al.*, “Long-term treatment with responsive brain stimulation in adults with refractory partial seizures,” *Neurology*, vol. 84, no. 8, pp. 810–817, 2015.

[26] N. V. Klinger and S. Mittal, “Clinical efficacy of deep brain stimulation for the treatment of medically refractory epilepsy,” *Clin. Neurol. Neurosurg.*, vol. 140, no. 2016, pp. 11–25, 2016.

[27] S. W.C. and L. B., “Technology Insight: Neuroengineering and epilepsy - Designing devices for seizure control,” *Nature Clinical Practice Neurology*, vol. 4, no. 4. pp. 190–201, 2008.

[28] G. P. Thomas and B. C. Jobst, “Critical review of the responsive neurostimulator system for epilepsy,” *Med. Devices Evid. Res.*, vol. 8, pp. 405–411, 2015.

[29] M. Mohan *et al.*, “The long-term outcomes of epilepsy surgery,” *PLoS One*, vol. 13, no. 5, pp. 1–16, 2018.

[30] M. Zijlmans, P. Jiruska, R. Zelmann, F. Leijten, J. G. R. Jefferys, and J. Gotman, “High frequency oscillations as a new biomarker in epilepsy,” *Ann. Neurol.*, vol. 71, no. 2, pp. 169–178, 2012.

[31] K. Kobayashi *et al.*, “Detection of seizure-associated high-frequency oscillations above 500 Hz,” *Epilepsy Res.*, vol. 88, no. 2–3, pp. 139–144, 2010.

[32] G. Buzsaki, “The Hippocampo-Neocortical Dialogue,” pp. 81–92, 1996.

[33] T. Nagasawa, C. Juhász, R. Rothermel, K. Hoechstetter, S. Sood, and E. Asano,

"Spontaneous and visually driven high-frequency oscillations in the occipital cortex: Intracranial recording in epileptic patients," *Hum. Brain Mapp.*, vol. 33, no. 3, pp. 569–583, 2012.

[34] A. Bragin, J. Engel Jr, and R. J. Staba, "High-frequency oscillations in epileptic brain," *Curr. Opin. Neurobiol.*, vol. 23, no. 2, pp. 151–156, 2010.

[35] J. Jacobs *et al.*, "High-frequency oscillations (HFOs) in clinical epilepsy," *Prog. Neurobiol.*, vol. 98, no. 3, pp. 302–315, 2012.

[36] H. Fujiwara *et al.*, "Resection of ictal high-frequency oscillations leads to favorable surgical outcome in pediatric epilepsy," *Epilepsia*, vol. 53, no. 9, pp. 1607–1617, 2012.

[37] C. Haegelen *et al.*, "High-frequency oscillations, extent of surgical resection, and surgical outcome in drug-resistant focal epilepsy," *Epilepsia*, vol. 54, no. 5, pp. 848–857, 2013.

[38] J. R. Cho *et al.*, "Resection of individually identified high-rate high-frequency oscillations region is associated with favorable outcome in neocortical epilepsy," *Epilepsia*, vol. 55, no. 11, pp. 1872–1883, 2014.

[39] Y. Höller *et al.*, "High-frequency oscillations in epilepsy and surgical outcome. A meta-analysis," *Front. Hum. Neurosci.*, vol. 9, no. OCTOBER, pp. 1–14, 2015.

[40] T. Fedele *et al.*, "Resection of high frequency oscillations predicts seizure outcome in the individual patient," *Sci. Rep.*, vol. 7, no. 1, pp. 1–10, 2017.

[41] A. Bragin, J. Engel, C. L. Wilson, I. Fried, and G. W. Mathern, "Hippocampal and entorhinal cortex high-frequency oscillations (100-500 Hz) in human epileptic brain and in kainic acid-treated rats with chronic seizures," *Epilepsia*, vol. 40, no. 2, pp. 127–137, 1999.

[42] G. a. Worrell, K. Jerbi, K. Kobayashi, J. M. Lina, R. Zelmann, and M. Le Van Quyen, "Recording and analysis techniques for high-frequency oscillations," *Prog. Neurobiol.*, vol. 98, no. 3, pp. 265–278, 2012.

[43] B. Frauscher *et al.*, "High-frequency oscillations: The state of clinical research," *Epilepsia*, vol. 58, no. 8, pp. 1316–1329, 2017.

[44] R. J. Staba *et al.*, "Quantitative Analysis of High-Frequency Oscillations (80 – 500 Hz) Recorded i," 2002.

[45] J. Engel and F. L. da Silva, "High-frequency oscillations - Where we are and where we need to go," *Prog. Neurobiol.*, vol. 98, no. 3, pp. 316–318, 2012.

[46] L. Menendez de la Prida, R. J. Staba, and J. A. Dian, "Conundrums of high-frequency oscillations (80-800 Hz) in the epileptic brain," *J. Clin. Neurophysiol.*, vol. 32, no. 3, pp. 207–219, 2015.

[47] D. King-Stephens, "The Ambiguous Nature of Fast Ripples in Epilepsy Surgery," *Epilepsy Curr.*, vol. 19, no. 2, pp. 91–92, 2019.

[48] J. Jacobs and M. Zijlmans, "HFO to Measure Seizure Propensity and Improve Prognostication in Patients With Epilepsy," *Epilepsy Curr.*, 2020.

[49] P. Jiruska *et al.*, "Update on the mechanisms and roles of high-frequency oscillations in

seizures and epileptic disorders," *Epilepsia*, vol. 58, no. 8, pp. 1330–1339, 2017.

[50] M. Zijlmans *et al.*, "How to record high-frequency oscillations in epilepsy: A practical guideline," *Epilepsia*, vol. 58, no. 8, pp. 1305–1315, 2017.

[51] S. V Gliske, Z. T. Irwin, C. Chestek, and W. C. Stacey, "Clinical Neurophysiology Effect of sampling rate and filter settings on High Frequency Oscillation detections," *Clin. Neurophysiol.*, vol. 127, no. 9, pp. 3042–3050, 2016.

[52] R. J. Staba, C. L. Wilson, A. Bragin, and I. Fried, "Quantitative analysis of high-frequency oscillations (80-500 Hz) recorded in human epileptic hippocampus and entorhinal cortex," *J. Neurophysiol.*, vol. 88, no. 4, pp. 1743–1752, 2002.

[53] G. Buzsáki *et al.*, "High-frequency network oscillation in the hippocampus," *Science (80-.).*, vol. 256, no. 5059, pp. 1025–1027, 1992.

[54] R. Zelmann, F. Mari, J. Jacobs, M. Zijlmans, F. Dubeau, and J. Gotman, "A comparison between detectors of high frequency oscillations," *Clin. Neurophysiol.*, vol. 123, no. 1, pp. 106–116, 2012.

[55] J. Jacobs *et al.*, "High frequency oscillations (80-500 Hz) in the preictal period in patients with focal seizures," *Epilepsia*, vol. 50, no. 7, pp. 1780–1792, 2009.

[56] A. Pearce, D. Wulsin, J. A. Blanco, A. Krieger, B. Litt, and W. C. Stacey, "Temporal changes of neocortical high-frequency oscillations in epilepsy.," *J. Neurophysiol.*, vol. 110, no. 5, pp. 1167–1179, Sep. 2013.

[57] K. Lehnertz *et al.*, "Seizure prediction by nonlinear EEG analysis," *IEEE Eng. Med. Biol. Mag.*, vol. 22, no. 1, pp. 57–63, 2003.

[58] F. Mormann *et al.*, "On the predictability of epileptic seizures," *Clin. Neurophysiol.*, vol. 116, no. 3, pp. 569–587, 2005.

[59] F. Mormann, R. G. Andrzejak, C. E. Elger, and K. Lehnertz, "Seizure prediction: The long and winding road," *Brain*, vol. 130, no. 2, pp. 314–333, 2007.

[60] D. E. Snyder, J. Echauz, D. B. Grimes, and B. Litt, "The statistics of a practical seizure warning system," *J. Neural Eng.*, vol. 5, no. 4, pp. 392–401, 2008.

[61] L. Kuhlmann *et al.*, "Patient-specific bivariate-synchrony-based seizure prediction for short prediction horizons," *Epilepsy Res.*, vol. 91, no. 2–3, pp. 214–231, 2010.

[62] L. Kuhlmann, K. Lehnertz, M. P. Richardson, B. Schelter, and H. P. Zaveri, "Seizure prediction — ready for a new era," *Nat. Rev. Neurol.*, vol. 14, no. 10, pp. 618–630, 2018.

[63] B. H. Brinkmann *et al.*, "Crowdsourcing reproducible seizure forecasting in human and canine epilepsy," *Brain*, vol. 139, no. 6, pp. 1713–1722, 2016.

[64] L. Kuhlmann *et al.*, "Epilepsyecosystem.org: Crowd-sourcing reproducible seizure prediction with long-term human intracranial EEG," *Brain*, vol. 141, no. 9, pp. 2619–2630, 2018.

[65] P. J. Karoly *et al.*, "The circadian profile of epilepsy improves seizure forecasting," *Brain*, vol. 140, no. 8, pp. 2169–2182, 2017.

[66] M. J. Cook *et al.*, “Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: A first-in-man study,” *Lancet Neurol.*, vol. 12, no. 6, pp. 563–571, 2013.

[67] N. D. Truong *et al.*, “Convolutional neural networks for seizure prediction using intracranial and scalp electroencephalogram,” *Neural Networks*, vol. 105, pp. 104–111, 2018.

[68] C. Reuben *et al.*, “Ensembling crowdsourced seizure prediction algorithms using long-term human intracranial EEG,” *Epilepsia*, vol. 61, no. 2, pp. e7–e12, 2020.

[69] M. I. Maturana *et al.*, “Critical slowing down as a biomarker for seizure susceptibility,” *Nat. Commun.*, vol. 11, no. 1, pp. 1–12, 2020.

[70] O. Stojanović, L. Kuhlmann, and G. Pipa, “Predicting epileptic seizures using nonnegative matrix factorization,” *PLoS One*, vol. 15, no. 2, pp. 1–13, 2020.

[71] M. Eberlein *et al.*, “Convolutional Neural Networks for Epileptic Seizure Prediction,” *Proc. - 2018 IEEE Int. Conf. Bioinforma. Biomed. BIBM 2018*, pp. 2577–2582, 2019.

[72] M. Levesque, P. Salami, J. Gotman, and M. Avoli, “Two Seizure-Onset Types Reveal Specific Patterns of High-Frequency Oscillations in a Model of Temporal Lobe Epilepsy,” *J. Neurosci.*, vol. 32, no. 38, pp. 13264–13272, 2012.

[73] V. I. Dzhala and K. J. Staley, “Transition from Interictal to Ictal Activity in Limbic Networks In Vitro,” *J. Neurosci.*, vol. 23, no. 21, pp. 1–8, 2003.

[74] G. A. Worrell, L. Parish, S. D. Cranstoun, R. Jonas, G. Baltuch, and B. Litt, “High-frequency oscillations and seizure generation in neocortical epilepsy,” *Brain*, vol. 127, no. 7, pp. 1496–1506, 2004.

[75] P. Jiruska *et al.*, “High-Frequency Network Activity, Global Increase in Neuronal Activity, and Synchrony Expansion Precede Epileptic Seizures In Vitro,” *J. Neurosci.*, vol. 30, no. 16, pp. 5690–5701, 2010.

[76] Y. Sato, S. M. Doesburg, S. M. Wong, C. Boelman, A. Ochi, and H. Otsubo, “Preictal surrender of post-spike slow waves to spike-related high-frequency oscillations (80-200 Hz) is associated with seizure initiation,” *Epilepsia*, vol. 55, no. 9, pp. 1399–1405, 2014.

[77] U. Malinowska, G. K. Bergey, J. Harezlak, and C. C. Jouny, “Identification of seizure onset zone and preictal state based on characteristics of high frequency oscillations,” *Clin. Neurophysiol.*, vol. 126, no. 8, pp. 1505–1513, 2015.

[78] P. Salami, M. Lévesque, and M. Avoli, “High frequency oscillations can pinpoint seizures progressing to status epilepticus,” *Exp. Neurol.*, vol. 280, pp. 24–29, 2016.

[79] A. Bragin, C. L. Wilson, J. Almajano, I. Mody, and J. Engel, “High-frequency oscillations after status epilepticus: Epileptogenesis and seizure genesis,” *Epilepsia*, vol. 45, no. 9, pp. 1017–1023, 2004.

[80] R. D. Traub *et al.*, “A possible role for gap junctions in generation of very fast EEG oscillations preceding the onset of, and perhaps initiating, seizures,” *Epilepsia*, vol. 42, no. 2, pp. 153–170, 2001.

[81] W. Stacey, M. Le Van Quyen, F. Mormann, and A. Schulze-Bonhage, “What is the present-day EEG evidence for a preictal state?,” *Epilepsy Res.*, vol. 97, no. 3, pp. 243–51, Dec. 2011.

[82] B. Litt *et al.*, “Epileptic seizures may begin hours in advance of clinical onset: A report of five patients,” *Neuron*, vol. 30, no. 1, pp. 51–64, 2001.

[83] J. A. Blanco *et al.*, “Data mining neocortical high-frequency oscillations in epilepsy and controls,” *Brain*, vol. 134, no. 10, pp. 2948–2959, 2011.

[84] A. Matsumoto *et al.*, “Pathological and physiological high-frequency oscillations in focal human epilepsy,” *J. Neurophysiol.*, vol. 110, no. 8, pp. 1958–1964, 2013.

[85] S. Ren, S. V. Gliske, D. Brang, and W. C. Stacey, “Redaction of false high frequency oscillations due to muscle artifact improves specificity to epileptic tissue,” *Clin. Neurophysiol.*, vol. 130, no. 6, pp. 976–985, 2019.

[86] S. V. Gliske, Z. T. Irwin, K. A. Davis, K. Sahaya, C. Chestek, and W. C. Stacey, “Universal automated high frequency oscillation detector for real-time, long term EEG,” *Clin. Neurophysiol.*, vol. 127, no. 2, pp. 1057–1066, 2016.

[87] N. S. Abend *et al.*, “Interobserver reproducibility of EEG interpretation in critically Ill children,” *J. Clin. Neurophysiol.*, vol. 28, no. 3, pp. 333–334, 2011.

[88] S. R. Benbadis, W. C. Lafrance, G. D. Papandonatos, K. Korabathina, K. Lin, and H. C. Kraemer, “Interrater reliability of EEG-video monitoring,” *Neurology*, vol. 73, no. 11, pp. 843–846, 2009.

[89] O. Aalen, “Nonparametric Inference for a Family of Counting Processes,” *Ann. Stat.*, vol. 6, no. 4, pp. 701–726, 2007.

[90] H.-G. Muller and J.-L. Wang, “Hazard Rate Estimation under Random Censoring with Varying Kernels and Bandwidths,” *Biometrics*, vol. 50, no. 1, pp. 61–76, 1994.

[91] H.-G. Muller and U. Stadtmuller, “Variable Bandwidth Kernel Estimators of Regression Curves,” *Ann. Stat.*, vol. 15, no. 1, pp. 182–201, 1987.

[92] M. Sedigh-Sarvestani *et al.*, “Rapid Eye Movement Sleep and Hippocampal Theta Oscillations Precede Seizure Onset in the Tetanus Toxin Model of Temporal Lobe Epilepsy,” *J. Neurosci.*, vol. 34, no. 4, pp. 1105–1114, 2014.

[93] H. Luna-Munguia, P. Starski, W. Chen, S. Gliske, and W. C. Stacey, “Control of in vivo ictogenesis via endogenous synaptic pathways,” *Sci. Rep.*, vol. 7, no. 1, pp. 1–13, 2017.

[94] S. Wiebe, W. T. Blume, J. P. Girvin, and M. Eliasziw, “A randomized, controlled trial of surgery for temporal-lobe epilepsy,” *N. Engl. J. Med.*, vol. 345, no. 5, pp. 311–318, 2001.

[95] K. Noe *et al.*, “Long-term outcomes after nonlesional extratemporal lobe epilepsy surgery,” *JAMA Neurol.*, vol. 70, no. 8, pp. 1003–1008, 2013.

[96] L. Andrade-Valença *et al.*, “Interictal high frequency oscillations (HFOs) in patients with focal epilepsy and normal MRI,” *Clin. Neurophysiol.*, vol. 123, no. 2012, pp. 100–105, 2012.

[97] M. Zijlmans, J. Jacobs, Y. U. Kahn, R. Zelmann, F. Dubeau, and J. Gotman, “Ictal and

interictal high frequency oscillations in patients with focal epilepsy," *Clin. Neurophysiol.*, vol. 122, no. 4, pp. 664–671, 2011.

[98] P. N. Modur, S. Zhang, and T. W. Vitaz, "Ictal high-frequency oscillations in neocortical epilepsy: Implications for seizure localization and surgical resection," *Epilepsia*, vol. 52, no. 10, pp. 1792–1801, 2011.

[99] C. Alvarado-Rojas *et al.*, "Slow modulations of high-frequency activity (40-140 Hz) discriminate preictal changes in human focal epilepsy," *Sci. Rep.*, vol. 4, pp. 1–9, 2014.

[100] A. Bragin, A. Azizyan, J. Almajano, C. L. Wilson, and J. Engel, "Analysis of chronic seizure onsets after intrahippocampal kainic acid injection in freely moving rats," *Epilepsia*, vol. 46, no. 10, pp. 1592–1598, 2005.

[101] M. O. Baud *et al.*, "Multi-day rhythms modulate seizure risk in epilepsy," *Nat. Commun.*, vol. 9, no. 1, pp. 1–10, 2018.

[102] F. H. Lopes da Silva, W. Blanes, J. Parra, D. N. Velis, S. N. Kalitzin, and P. Suffczynski, "Dynamical Diseases of Brain Systems: Different Routes to Epileptic Seizures," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 5, pp. 540–548, 2003.

[103] M. L. Saggio *et al.*, "A taxonomy of seizure dynamotypes," *Elife*, vol. 9, Jul. 2020.

[104] K. Gadhoumi, J. M. Lina, F. Mormann, and J. Gotman, "Seizure prediction for therapeutic devices: A review," *J. Neurosci. Methods*, vol. 260, no. 029, pp. 270–282, 2016.

[105] D. R. Freestone *et al.*, "Seizure Prediction: Science Fiction or Soon to Become Reality?," *Curr. Neurol. Neurosci. Rep.*, vol. 15, no. 11, 2015.

[106] C. Alexandre Teixeira *et al.*, "Epileptic seizure predictors based on computational intelligence techniques: A comparative study with 278 patients," *Comput. Methods Programs Biomed.*, vol. 114, no. 3, pp. 324–336, 2014.

[107] M. A. van 't Klooster *et al.*, "Tailoring epilepsy surgery with fast ripples in the intraoperative electrocorticogram," *Ann. Neurol.*, vol. 81, no. 5, pp. 664–676, 2017.

[108] J. a Blanco *et al.*, "Unsupervised classification of high-frequency oscillations in human neocortical epilepsy and control patients.," *J. Neurophysiol.*, vol. 104, no. 5, pp. 2900–2912, 2010.

[109] J. M. Scott, S. Ren, S. V. Gliske, and W. C. Stacey, "Preictal variability of high-frequency oscillation rates in refractory epilepsy," *Epilepsia*, 2020.

[110] S. V Gliske *et al.*, "Distinguishing false and true positive detections of high frequency oscillations," *J. Neural Eng.*, vol. 17, no. 5, p. 056005, Sep. 2020.

[111] P. Mirowski, D. Madhavan, Y. LeCun, and R. Kuzniecky, "Classification of patterns of EEG synchronization for seizure prediction," *Clin. Neurophysiol.*, vol. 120, no. 11, pp. 1927–1940, 2009.

[112] H. Ung *et al.*, "Intracranial EEG fluctuates over months after implanting electrodes in human brain," *J. Neural Eng.*, vol. 14, no. 5, pp. 1–25, 2017.

[113] M. Zijlmans, J. Jacobs, R. Zelmann, F. Dubeau, and J. Gotman, "High-frequency oscillations mirror disease activity in patients with epilepsy.," *Neurology*, vol. 72, no. 11,

pp. 979–86, Mar. 2009.

[114] S. V Gliske *et al.*, “Variability in the location of high frequency oscillations during prolonged intracranial EEG recordings,” *Nat. Commun.*, vol. 9, no. 1, 2018.

[115] R. Tibshirani, “Regression shrinkage and selection via the lasso: A retrospective,” *J. R. Stat. Soc. Ser. B Stat. Methodol.*, vol. 73, no. 3, pp. 273–282, 2011.

[116] C. W. Lu, K. A. Malaga, K. L. Chou, C. A. Chestek, and P. G. Patil, “High density microelectrode recording predicts span of therapeutic tissue activation volumes in subthalamic deep brain stimulation for Parkinson disease,” *Brain Stimul.*, vol. 13, no. 2, pp. 412–419, 2020.

[117] D. J. Wilson, “Erratum: The harmonic mean p-value for combining dependent tests (Proceedings of the National Academy of Sciences of the United States of America (2019) 166 (1195-1200) DOI: 10.1073/pnas.1814092116),” *Proc. Natl. Acad. Sci. U. S. A.*, vol. 116, no. 43, p. 21948, 2019.

[118] K. A. Davis *et al.*, “The effect of increased intracranial EEG sampling rates in clinical practice,” *Clin. Neurophysiol.*, vol. 129, no. 2, pp. 360–367, 2018.

[119] S. V. Gliske, W. C. Stacey, K. R. Moon, and A. O. Hero, “The intrinsic value of HFO features as a biomarker of epileptic activity,” *ICASSP, IEEE Int. Conf. Acoust. Speech Signal Process. - Proc.*, vol. 2016-May, pp. 6290–6294, 2016.

[120] H. G. Wieser *et al.*, “Proposal for a new classification of outcome with respect to epileptic seizures following epilepsy surgery,” *Epilepsia*, vol. 42, no. 2, pp. 282–286, 2001.

[121] H. P. Zaveri, S. M. Pincus, I. I. Goncharova, R. B. Duckrow, D. D. Spencer, and S. S. Spencer, “Localization-related epilepsy exhibits significant connectivity away from the seizure-onset area,” *Neuroreport*, vol. 20, no. 9, pp. 891–895, 2009.

[122] B. Rosner, R. J. Glynn, and M. L. T. Lee, “The Wilcoxon signed rank test for paired comparisons of clustered data,” *Biometrics*, vol. 62, no. 1, pp. 185–192, 2006.

[123] J. A. Hanley, A. Negassa, M. D. d. B. Edwardes, and J. E. Forrester, “Statistical analysis of correlated data using generalized estimating equations: An orientation,” *Am. J. Epidemiol.*, vol. 157, no. 4, pp. 364–375, 2003.

[124] O. Salem, Y. Liu, A. Mehaoua, and R. Boutaba, “Online Anomaly Detection in Wireless Body Area Networks for Reliable Healthcare Monitoring,” *IEEE J. Biomed. Heal. Informatics*, vol. 18, no. 5, pp. 1541–1551, 2014.

[125] X. She *et al.*, “Adaptive multi-channel event segmentation and feature extraction for monitoring health outcomes,” 2020.