Towards automatic classification of pathological epileptic tissue with interictal high frequency oscillations

Michael D. Nunez[a,b,*], Krit Charupanit[a,c], Indranil Sen-Gupta[d], Beth A. Lopour[a,1,*], Jack J. Lin[e,1]

[a] *Department of Biomedical Engineering, University of California, Irvine CA, USA*
[b] *Department of Psychiatry and Biobehavioral Sciences, University of California, Los Angeles CA, USA*
[c] *Institute of Biomedical Engineering, Faculty of Medicine, Prince of Songkla University, Songkhla 90112, Thailand*
[d] *Neurology, University of California Irvine Medical Center, Orange CA, USA*
[e] *Department of Neurology, University of California, Irvine CA, USA*

Abstract

High frequency oscillations (HFOs) recorded by intracranial electrodes have generated excitement for their potential to help localize epileptic tissue for surgical resection (Frauscher et al., 2017). However, previous research has shown that the number of HFOs per minute (i.e. the HFO "rate") is not stable over the duration of intracranial recordings. The rate of HFOs increases during periods of slow-wave sleep (von Ellenrieder et al., 2017), and HFOs that are predictive of epileptic tissue may occur in oscillatory patterns (Motoi et al., 2018). We sought to further understand how between-seizure (i.e. "interictal") HFO dynamics predict the seizure onset zone (SOZ). Using long-term intracranial EEG from 16 subjects, we fit Poisson and Negative Binomial mixture models that describe HFO dynamics and include the ability to switch between two discrete brain states. Oscillatory dynamics of HFO occurrences were found to be predictive of SOZ and were more consistently predictive than HFO rate. Using concurrent scalp-EEG in two patients, we show that the model-found brain states corresponded to (1) non-REM

*Corresponding author
Email addresses: mdnunez1@uci.edu (Michael D. Nunez), blopour@uci.edu (Beth A. Lopour)
[1] These authors contributed equally to this work.

(NREM) sleep and (2) awake and rapid eye movement (REM) sleep. This work suggests that unsupervised approaches for classification of epileptic tissue without sleep-staging can be developed using mixture modeling of HFO dynamics.

Keywords: Epilepsy, surgery, seizure onset zone (SOZ), Hierarchical Bayesian methods, epileptogenic zone, high-frequency oscillations (HFOs)

1. Introduction

Epilepsy is quite prevalent across the globe. For example, 1.2% of the population of the United States in 2015 were reported to have epilepsy (Zack and Kobau, 2017). Of this multitude of people with epilepsy, about 30% to 40% have seizures that cannot be controlled by antiseizure medication (Kwan and Brodie, 2000; Engel, 2018). In such cases of drug-resistant epilepsy, for which seizures greatly decrease the patient's quality of life, surgical interventions such as resection of seizure-generating tissue and implantation of responsive neurostimulators (RNS) are relatively safe procedures (Engel, 2018). The goal of epilepsy surgery is to identify and treat the epileptogenic zone (EZ), typically defined as the minimum amount of tissue that must be surgically removed or stimulated to achieve a seizure free outcome. However, the EZ is a theoretical construct, and no biomarkers exist that can accurately and consistently identify the EZ (Ryvlin et al., 2014). Instead, intracranial electroencephalography (iEEG) is used to localize the seizure onset zone (SOZ), as this is currently thought to be the best approximation of the EZ (Kovac et al., 2017). This information, along with other clinical data (such as patient history, brain imaging to identify structural abnormalities, seizure semiology, etc), is used by clinicians to select brain tissue for treatment. While surgery often results in a reduction of seizures, many patients will not be seizure free, indicating that there is a need for more accurate methods of identifying the EZ (Noachtar and Borggraefe, 2009). Such improvements would allow more patients to benefit from this procedure, with fewer side effects from the surgery and better outcomes (especially those with epilepsy outside of the temporal lobe with normal MRIs; Cohen-Gadol et al., 2006; Noe et al., 2013).

High frequency oscillations (HFOs) have shown promise as a novel marker of the EZ. Specifically, increased incidence (i.e. increased "rate" per minute) of transitory HFOs (Bragin et al., 2002) is thought to be indicative of the

EZ (Jacobs et al., 2008, 2010). HFOs are described as "transitory," as they are typically temporally isolated events that last less than 200 ms with 3 or more cycles (i.e. 6 positive and negative local peaks in the waveform; Staba et al., 2002; Jacobs et al., 2008; Charupanit and Lopour, 2017). HFOs are typically subcategorized as ripple band (80 - 250 Hz) and fast ripple band (250 - 500 Hz) events. These waveforms are thought to be generated by synchronous population firing and/or synchronous postsynaptic activity in the brain, although there is an abundance of possible neural mechanisms and cortical circuits that could generate HFOs (Köhling and Staley, 2011; Staba and Bragin, 2011; Jefferys et al., 2012). HFOs are typically recorded with intracranial macro or micro electrodes. Scalp recordings may also contain high frequency oscillations (Zelmann et al., 2014; von Ellenrieder et al., 2014; Kobayashi et al., 2015; Gotman, 2018); however the overlap with contaminant muscle activity pervasive in scalp electroencephalography (EEG) is a concern when interpreting HFOs in scalp recordings (Nunez et al., 2016; Fitzgibbon et al., 2016; Bernardo et al., 2018).

Research has further sought to differentiate *pathological* HFOs, occurring in the EZ, from *physiological* HFOs, which can occur across the brain due to normal processes. The difficulty in differentiating pathological HFOs from normal brain activity has been a barrier to the use of HFOs in modern clinical practice (Jacobs et al., 2018; Fedele et al., 2019). For example, even though high rates of HFOs are typically thought to be indicative of the SOZ, baseline rates of HFOs outside the SOZ vary across different regions of the cortex (Guragain et al., 2018; Frauscher et al., 2018). Pathological and physiological HFOs are also affected by the sleep state of the patient, and HFO rates during slow wave sleep (i.e. non-rapid eye movement; NREM sleep) are thought to be more differentiating of pathological versus physiological brain activity (Dümpelmann et al., 2015; von Ellenrieder et al., 2016, 2017). Fast ripples are generally more localized to SOZ than ripples, but they occur less frequently and may not be recorded in all patients (Köhling and Staley, 2011; Roehri et al., 2018). Roehri et al. (2018) show that HFOs are not more predictive of SOZ than pathological epileptiform discharges, although the co-occurrence of both is more predictive. Gliske et al. (2018) found that ripples during NREM sleep are only predictive of SOZ in some patients, and that HFO sources were highly variable over time. This led Gliske et al. (2018) to make the argument that long recordings over multiple days must be performed in order to accurately measure interictal, ripple-band HFO dynamics.

Coupling of ripple-band HFOs to slow waves has been observed during

preictal and seizure periods (Weiss et al., 2013; Ibrahim et al., 2014; Guirgis et al., 2015). Moreover, pathological, interictal HFOs may be modulated by high amplitude, low frequency background activity, especially during sleep (Kerber et al., 2014; Frauscher et al., 2015; von Ellenrieder et al., 2016; Motoi et al., 2018). This suggests that the temporal dynamics of HFOs and the precise timing of their occurrence may be an important marker of epileptogenic tissue. However, this characteristic of high frequency activity remains relatively unexplored compared to the simple counting of HFOs per minute.

In this study, we present new techniques and a data analysis pipeline for HFOs based on long-term, spontaneous iEEG that does not require sleep staging of the data. That is, periods of wakefulness and different sleep stages do not need to be identified prior to implementing HFO detection and SOZ prediction algorithms. This was achieved by assuming that HFO rates are generated by the simple mixture of only two brain states. We show that the oscillatory dynamics of HFOs are more consistently predictive of the SOZ across patients than the rates themselves. This prediction was more accurate when based on HFO dynamics during the empirically-found brain state that was reflective of NREM sleep.

2. Materials & Methods

2.1. Patients and iEEG recordings

Patients who had intractable epilepsy and were candidates for resective surgery had intracranial electrodes implanted at the University of California Irvine Medical Center in order better localize the SOZ. We analyzed iEEG data from patients ($N = 16$, 8 female, 36 ± 15 years of age, see **Table 1**) who were implanted with either subdural electrocorticography (ECoG) grids, depth macroelectrodes and/or stereotactic EEG (SEEG). The electrode types and locations were chosen by the clinicians for diagnostic and surgical evaluation.

Long-term iEEG was recorded for all subjects with high sample rates (minimum 2000 Hz, maximum 5000 Hz) in order to capture HFOs in the ripple band with high accuracy. Note that standard clinical sampling rates of 500 Hz and below may not be sufficient to capture ripples due to aliasing of digital signals. It is recommended that a sample rate of at least $250 * 2.5 = 625$ Hz be used to capture ripple-band HFOs; the 2.5 multiplier is Engineer's Nyquist given by Bendat and Piersol (2011). Specific channels

were classified by board-certified epileptologists as being within the SOZ if those channels had time courses indicative of seizure onset before propagation to other channels during any seizure captured via iEEG.

Channels were localized via coregistration of pre- and post-implantation magnetic resonance imaging (MRI) and/or post-implantation computed tomography (CT) as described by Stolk et al. (2017); Zheng et al. (2017); Helfrich et al. (2018); Stevenson et al. (2018). Each intracranial channel was classified as out-of-brain, within white matter, or within grey matter. If the location was on the boundary of the grey and white matter, it was labeled as white matter. If the location was near the edge of the brain, it was labeled as being outside the brain. We did not disaggregate by grey matter regions (hippocampus, amygdala, insular regions, neocortical regions, etc.), although other studies have described differing HFO rates between these regions (Blanco et al., 2011; Wang et al., 2017; Frauscher et al., 2018). Whenever possible, a channel within each strip or grid that was located within white matter was used as a reference. If no such information existed or was unclear from the localization, the channels were referenced to the average of all the channels on the grid or depth strip. The importance of coregistration was assessed with 6 of the 16 patients in which localization information was unavailable. That is, we tested the robustness of our procedure to the presence of noisy channels.

In two patients, scalp EEG, heart rate (electrocardiography; EKG), and eye movements (electrooculography, EOG) were concurrently recorded in order to extract sleep-stage information over time in offline analysis. The data was then sleep staged using the software from Greer and Saletin (2011). Thirty-second epochs of data were classified as NREM slow wave sleep, REM, wakefulness, or artifact. This sleep staging was then compared to HFO model-found brain states, discussed later.

2.2. Automatic detection of high frequency oscillations during overnight recordings

Automatic detection of HFOs is now widely used, and the results of automatic detectors are comparable to that of visual detection (e.g. see Jacobs et al., 2018). We detected HFOs automatically in each channel of iEEG over the duration of each patient's recording using the HFO detection software developed by Charupanit and Lopour (2017). This algorithm finds oscillations that are significantly larger than the amplitude noise floor in the 80 to 250 Hz frequency band. By iteratively generating a Poisson distribution

Patient	**Age**	**G**	**Diagnosis**	**Surgery**	**Engel**	**SOZ**	**nSOZ**	**Hours**
1	50	M	RTLE	RTLo	IB	4	40	13
2	23	F	LTLE	LTLo	IIIA	4	44	25
3	34	M	RTLE	RTLo	IA	3	40	15
4	21	M	RTLE	RTLo[1]	IIB	13	111	16
5*	21	F	RFLE	*RNS*	?	9	111	12
6*	28	F	RFLE	RFLe, PRFLo	?	4	92	5
7*	28	M	LTLE	LTLo	IIIA	13	113	9
8*	44	M	RFLE	RFLo	IA	5	123	12
9	57	F	LTLE	LTLo	IA	8	59	9
10	34	M	BTLE	*RNS*	IIA	14	40	9
11	69	F	RTLE	RTLo	IIIA	2	39	9
12*	18	M	LTLE	LTLo	IIA	11	161	7
13	22	F	LTLE	*RNS*	IIIA	2	75	5
14	53	F	LTLE	LTLo	IA	1	69	5
15	50	F	BTLE	*RNS*	IIIA	11	49	10
16*	27	M	BTLE	*RNS*	IVB	11	109	10

Table 1: Clinical information for each patient including age, gender (G), epilepsy diagnosis, surgery performed, Engel outcome, number of SOZ channels used in model fitting (SOZ), number of non-SOZ channels used in model fitting (nSOZ), and consecutive interictal hours used in the model fitting. * denotes patients for whom all channels were used to fit data to the mixture models, while we included only grey matter localized channels for the other patients. Abbreviations: P = partial, L = left, R = right, B = bilateral, F = frontal, T = temporal, L = lobe, E = epilepsy, Lo = lobectomy, Le = lesionectomy, *RNS* = implanted responsive neurostimulator. [1]This patient also had hypothalamic hamartoma ablation performed. Patients 5 and 6 had unknown Engel Outcomes.

of oscillation ("peak") amplitudes, the detector can identify events with at least 4 consecutive high amplitude oscillations that exist in the tail of the rectified amplitude distribution (i.e. 8 rectified peaks). Specifically, we defined the threshold as peak amplitudes above the 95.8% percentile (i.e. an alpha parameter of $\alpha = 0.042$, which was recommended by Charupanit and Lopour (2017)). Note that the estimation of the noise floor is adaptive and will change for each channel. We also allowed the noise floor to change every 5 minutes *within* each channel to account for nonstationarities, such as changes in state of vigilance and sleep stage. To ensure that HFO rates were not affected by the occurrence of seizures, we analyzed only interictal HFOs that occurred at least 1 hour away from clinically-identified seizures. The resulting dataset had at least 5 hours of iEEG per patient, with a maximum of 25 hours for one patient, and a mean and standard deviation of 10.69 ± 4.90 hours across $N = 16$ patients (see **Table 1**). An example of the changing rate of HFOs from two channels within one patient is given in **Figure 1**.

2.3. Assuming mixtures of HFO distributions

Based on previous research (Dümpelmann et al., 2015; von Ellenrieder et al., 2016, 2017), we assumed that HFO rates would be a function of the state of vigilance and sleep stage. We also suspected that HFO rates might change based on the cognitive brain state of the patient. For these two reasons we allowed our statistical model to identify the states inherent in the data without us explicitly defining each one.

Mathematically, we assumed that HFO counts per second for each channel were distributed from either Poisson or Negative Binomial mixture models, which are typically used to model count data. Operationally, each channel was associated with multiple distributions of HFO occurrences (one distribution per state), with the representative state changing over time. We enforced the restriction that a change in state caused the distributions from all channels to change at the same time. Thus, we assumed that the brain's state of vigilance or sleep changed each channel's HFO dynamics simultaneously (although the distributions for each channel could change in different ways). We constrained the number of possible brain states to only 2-4 per channel, switching at most every five minutes during overnight recordings. After initial results, we constrained the number of Poisson and Negative Binomial mixtures to be *two* per channel based on our empirical findings.

To automatically obtain HFO dynamics across channels and brain states within each patient, we used hierarchical Bayesian methods with JAGS, a

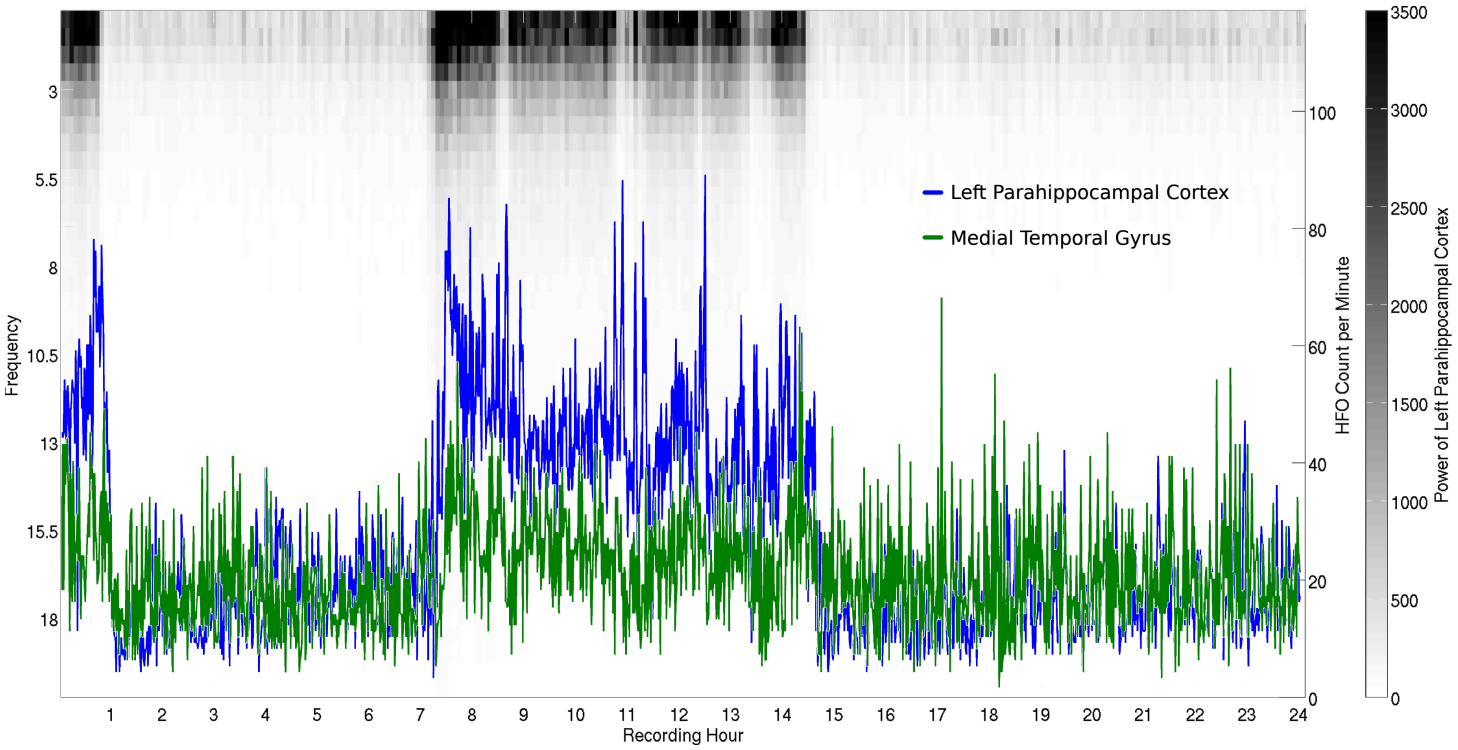

Figure 1: Modulation of HFO rate over 24 hours in one example patient. HFOs were automatically detected using the algorithm by Charupanit and Lopour (2017). Automated HFO counts per minute are shown from an iEEG channel in the left parahippocampal cortex (blue line) and an iEEG channel in the medial temporal gyrus (green line), with HFO counts per minute denoted by the right y-axis. These counts are overlayed on a grey-scale color map of low frequency band power (1-20 Hz) from the same left parahippocampal iEEG channel during the same 24 hour time period. Darker shading indicates higher power in the frequency band denoted by the left y-axis. HFO counts (blue and green lines) are modulated by the sleep-wake cycle (von Ellenrieder et al., 2017), as evidenced by their correlation with delta (1-4 Hz) frequency power (grey-scale color map). HFOs are typically analyzed during slow-wave sleep. However, with the method presented here, the data does not need to be visually sleep-staged prior to classification of SOZ and non-SOZ channels.

program that can easily sample from hierarchical models using Markov Chain Monte Carlo (MCMC) (Plummer, 2003) via the pyjags Python package (Miasko, 2017). We used a hierarchical model to restrict the distributions of HFO rates across brain regions within each brain state. These hierarchical parameters were included because we expect greater variation across brain states within each channel than we do across channels within the same brain state. However, even if the channels have very different rates within the same brain state (for example, because some regions like the hippocampus have increased occurrences of HFOs), the model with hierarchical parameters per brain state can fit data more accurately compared to non-hierarchical models due to a statistical phenomenon known as "shrinkage". Hierarchical parameters for HFO dynamics will "shrink" towards the mean parameters, which will improve estimation in the presence of outliers (Gelman et al., 2013; Boehm et al., 2018). Hierarchical models also allow for better estimation based on small amounts of data (Gelman et al., 2013).

The exact model for each patient, which we will refer to as Model 1, is given by the following equations:

$$
\begin{aligned}
(\text{HFO count})_{te} &\sim \mathit{Poisson}(\lambda_{ke}) && (1)\\
\mu_k &\sim \mathit{Normal}(1, .5^2) && (2)\\
\sigma_k &\sim \mathit{Gamma}(1, 1) && (3)\\
\lambda_{ke} &\sim \mathit{Normal}(\mu_k, \sigma_k^2) \in (0, \infty) && (4)\\
k_w &\sim \mathit{Categorical}(\underline{\pi}) && (5)\\
\underline{\pi} &\sim \mathit{Dirichlet}(1, 1) && (6)
\end{aligned}
$$

Such that HFO count per second t is described by a Poisson process with rate λ for each brain state k and channel e. Each rate λ is drawn from a hierarchical distribution with a mean rate μ for all channels in that brain state. The brain state k for every 5 minute window w is drawn from a Categorical distribution with a vector of probabilities $\underline{\pi}$ totalling 100% for all brain states ($\underline{\pi}$ is a vector of size $k = 2$ in the presented data for just two brain states).

2.4. Temporal dynamics within brain states

In order to account for temporal dynamics of HFO rates within one brain state, we can relax the restriction that HFO counts must follow a strictly

Poisson process. It has previously been shown that pathological, interictal HFOs can be modulated by high amplitude, low frequency background activity during sleep (Kerber et al., 2014; Frauscher et al., 2015; von Ellenrieder et al., 2016; Motoi et al., 2018). While we did not measure phase-amplitude coupling of HFOs to slow rhythms directly, we estimated the variance of the HFOs over time in hierarchical models. "Overdispersion" occurs when there is greater variability in the data (e.g. HFO count data) than is expected by a Poisson process. "Underdispersion" occurs when there is smaller variability in the data than is expected by a Poisson process. Overdispersion can be due to the clumping of HFOs in close temporal proximity to one another, such that there are bursts of HFOs occurring in time. Underdispersion could be due to HFOs being time locked to the peaks or troughs of slower rhythms, such that the HFOs occur at approximately constant temporal intervals. Thus we might expect overdispersion or underdispersion to be predictive of SOZ based on the previous literature. We estimated parameters in hierarchical models that provide inference as to whether HFOs occur in oscillatory patterns or are clumped together.

To account for over- and underdispersion, a hierarchical model was fit to each patient's HFO count data using the Negative Binomial distribution (Cook, 2009). The Negative Binomial distribution has two common parameterizations. In the parameterization we used within the JAGS program, the Negative Binomial distribution gives a number of "failures" (e.g. number of HFOs within a given second) before η "successes" where θ is the probability of a success (e.g. the probability that no HFOs occur in a given second) (Plummer, 2003; Cook, 2009). Note that η is not restricted to integers. To aid prediction of SOZ, we transformed the two parameters of the Negative Binomial distribution θ and η to create three parameters: (1) the rate of HFOs λ, as in Model 1 which assumed Poisson count distributions, (2) the coefficient of variation (CV, γ), defined as the ratio of the standard deviation of counts over the mean count rate, and (3) a "clumping coefficient" (CC, ζ), defined as the inverse of η. Note that as η goes to positive infinity, the clumping coefficient ζ goes to zero and the Negative Binomial distribution approaches a Poisson distribution (Cook, 2009). Large clumping coefficients ζ indicate that HFOs are more likely to occur immediately following other HFOs. CVs much greater than 1 ($\gamma >> 1$) also indicate temporal clumping of HFOs, i.e. a CV greater than one indicates that the variance in rate is more than the mean rate. CVs much less than 1 ($\gamma << 1$) indicate oscillatory dynamics, i.e. a CV less than one indicates that the variance in rate is less

than the mean rate. The CC and CV were included in the model to gauge the predictive ability of both parameters to inform the location of the SOZ.

The model for each patient that accounted for over- and underdispersion, which we will refer to as Model 2, is given by the following equations:

$$
\begin{aligned}
(\text{HFO count})_{te} &\sim \mathit{NegBinomial}(\theta_{ke}, \eta_{ke}) && (7)\\
\theta_{ke} &= \eta_{ke}/(\eta_{ke} + \lambda_{ke}) && (8)\\
\zeta_{ke} &= 1/\eta_{ke} && (9)\\
\gamma_{ke} &= 1/\sqrt{\eta_{ke}(1-\theta_{ke})} && (10)\\
\mu_k &\sim \mathit{Normal}(1, .5^2) && (11)\\
\sigma_k &\sim \mathit{Gamma}(1, 1) && (12)\\
\lambda_{ke} &\sim \mathit{Normal}(\mu_k, \sigma_k^2) \in (0, \infty) && (13)\\
\eta_{ke} &\sim \mathit{Uniform}(0, 50) && (14)\\
k_w &\sim \mathit{Categorical}(\underline{\pi}) && (15)\\
\underline{\pi} &\sim \mathit{Dirichlet}(1, 1) && (16)
\end{aligned}
$$

2.5. Classification of SOZ and non-SOZ channels

All parameters were estimated by the medians of posterior distributions from model fits. We used the rate parameter λ and oscillatory dynamic parameters, namely the clumping coefficient ζ and the coefficient of variation γ, to classify channels as SOZ or non-SOZ. We built Receiver Operating Characteristic (ROC) curves by varying the cutoff values of λ, ζ, and γ, where the SOZ was indicated by a high value of λ and low values of ζ and γ. One curve was created for each patient in each automatically-derived brain state. Brain states were sorted by the average amount of standardized delta (1-4 Hz) power across all iEEG records used in the models, and we will refer to them as Brain State 1 (brain state with highest mean delta power across channels) and Brain State 2 (brain state with the lowest mean delta power across channels), see **Figure 2**. ROC curves show the trade-off between true positives (channels identified as SOZ by clinicians that are also automatically labeled as candidate SOZ channels) versus false positives. Clinicians and researchers may find ROC curves useful because of the possible trade-offs between resecting or not-resecting some identified SOZ tissue. These ROC curves show the overall accuracy of our prediction over a continuum of possible cutoff values.

2.6. Comparison to surgical outcome

The ability of our automated method to identify the SOZ was compared to the surgical outcome of the patients. It has been hypothesized that appropriately measured interictal HFO dynamics may be a better marker of the EZ than the current standard measure, which is the site of seizure onset. If this hypothesis is true, surgery based on the SOZ (along with structural abnormalities etc.) may lead to worse surgical outcomes than surgery based on HFO dynamics (along with structural abnormalities etc.). Therefore, if the brain regions identified by HFOs are consistent with the SOZ, and the SOZ is targeted during surgery, we expect the patient to have a good outcome. If HFOs suggest brain regions beyond the SOZ, and the SOZ is targeted during surgery, we expect the patient to have a worse outcome. In other words, surgical outcomes would be worse for patients where HFOs did not predict the SOZ. We used the Engel Outcome Scale (Engel, 1993) to quantify patients' surgical outcome and compare to classification of SOZ using HFOs. Specifically, we compared total area under the ROC curve (AUC) to Engel Outcome.

It should be noted that some patients had lobectomies performed while others had responsive neurostimulators (RNS) implanted (see **Table 1**), depending upon preoperative imaging, the locus of suspected irritative tissue, and postoperative imaging. Also, depending on attendant neurosurgical considerations, the entirety of SOZ tissue was not removed in all patients who underwent surgical resection. We did not include patients in whom an SOZ could not be identified.

3. Data and code availability

Automatically identified HFO data, standardized delta (1-4 Hz) power, channel localization labels, and samples from posterior distributions for both Model 1 and Model 2 are available upon request and at `https://doi.org/10.6084/m9.figshare.12385613`. Patient data (including raw iEEG data) is not widely available to protect the privacy of patients. MATLAB, Python, and JAGS data extraction and analysis code are available at `https://osf.io/3ephr/` and in the following repository `https://github.com/mdnunez/sozhfo` (as of June 2020).

4. Results

4.1. Model-derived brain states correspond to sleep and awake states

The Poisson mixture model (Model 1) and Negative Binomial mixture model (Model 2) automatically separated windows of time into two brain states based on the HFO dynamics in all channels. Brain state labels were thus influenced by the HFO dynamics from all channels simultaneously. We found that the brain state labels of converged Markov chains tracked known sleep/wake dynamics.

In two patients whose data was sleep staged manually using concurrent scalp EEG, the two HFO model-derived brain states blindly separated slow wave sleep (NREM, i.e. sleep stages 1, 2 and 3) from all other states (REM and wakefulness) as shown in the lower two panels of **Figure 2**. In Patient 16, the congruence between the visually sleep staged data and the HFO model-derived brain state from Model 2 was 90.0%, with NREM sleep being correctly identified by the model 91.7% of the time and the other states being correctly identified 83.3% of the time. In Patient 17, the congruence between the sleep staged data and the model-derived brain states was 91.7%, with both NREM and other states being correctly identified 91.7% of the time by the model.

We did not have concurrent EEG, EKG, and EOG in the other patients to evaluate the correspondence between sleep stages and model-derived brain states. However, we could evaluate how well standardized delta power (1-4 Hz) across model electrodes was differentiated by the two model-derived brain states, as a proxy for sleep staged data. In half of the patients (8/16), we found found that delta power was significantly ($p < .001$) different in the two states by both ANOVA and Kruskal-Wallis (non-parametric equivalent to ANOVA) tests at a false positive rate of $\alpha = .001$. An additional patient had a significant Kruskal-Wallis test ($p = .006$) at a false positive rate of $\alpha = .01$ and a small ANOVA p-value ($p = .011$). And one more patient had a significant ANOVA test ($p = .038$) at a false positive rate of $\alpha = .05$. Of the six patients without any indication of significant differences in delta power between the two model-derived brain states, four of these patients did not have localization information to exclude artifactual electrodes before the standardized mean delta calculation. Only two of 10 patients for whom we included localization information did not show evidence for model-derived brain states corresponding to delta power.

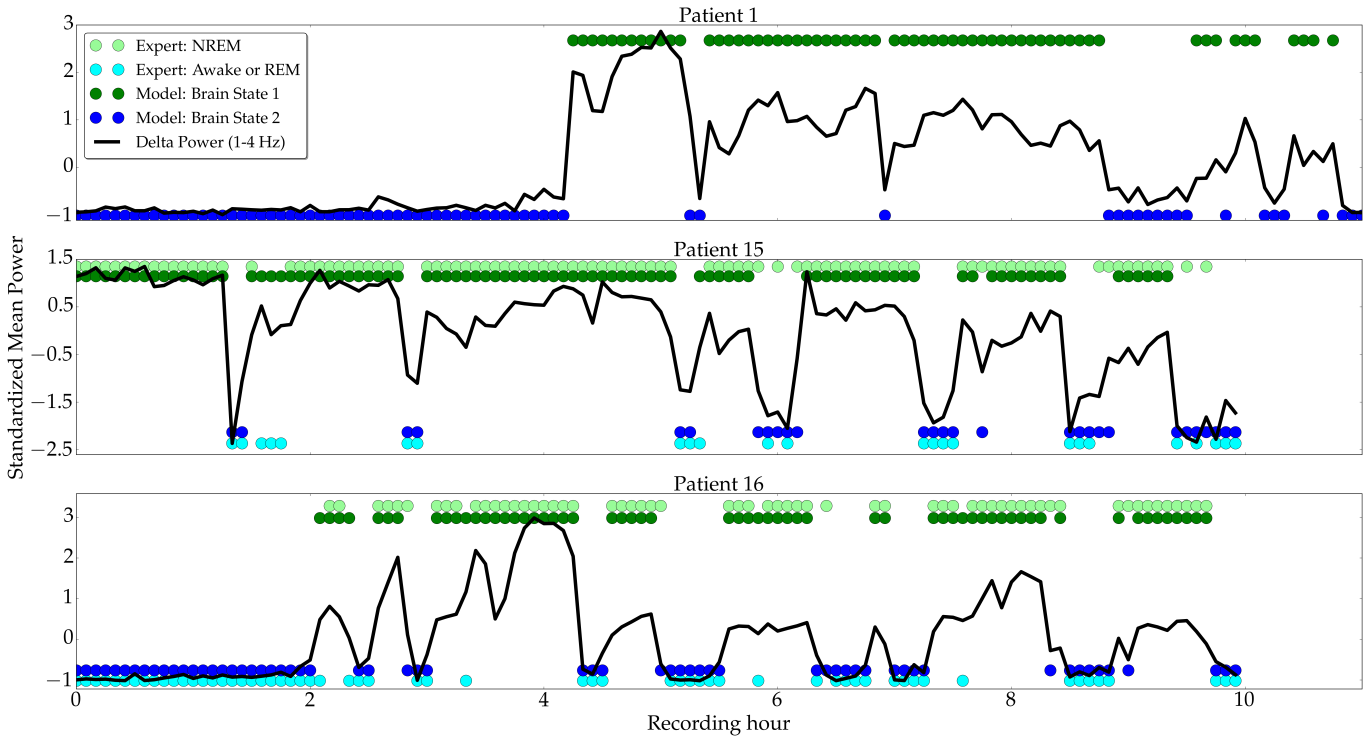

Figure 2: Model-derived brain states correspond to sleep and awake states. Representative examples are shown from Patients 1, 15, and 16 with brain states 1 (green dots) and 2 (blue dots) obtained automatically every 5 minutes from Model 2. They were separated into states 1 and 2 by mean slow-wave delta power (1-4 Hz; standardized mean across channels), with brain state 1 containing more delta power. Black lines give the standardized mean delta power across channels. Brain states obtained automatically from Model 1 were nearly identical to brain states obtained from Model 2. In patients 15 and 16, who had concurrent iEEG and scalp EEG, the HFO model-derived brain states differentiated slow wave sleep (i.e. sleep stages 1, 2 and 3; denoted by light green dots in the upper portion of the bottom two subplots) from all other states (REM sleep and wakefulness; denoted by teal dots in the lower portion of the bottom two subplots) based on a comparison to expert sleep staging.

4.2. HFO oscillatory dynamics predict SOZ

For 15 of 16 patients, small values of the HFO clumping coefficient ζ were predictive of SOZ (AUC $> .62$) in at least one of the two brain states found by Model 2. In 13 out of 16 patients, the clumping coefficients in brain state 1 differentiated all SOZ channels (i.e. 100% true positive rate) for false positive rates less than 62%. Seven of 16 patients had false positive rates of $< 20\%$ for a 100% true positive rate with data from brain state 1. The predictive ability of the CV was similar. The CV differentiated SOZ channels with a 100% true positive rate for false positive rates less than 47% in 13 of 16 patients' brain state 1. Using the CV, 8 of 16 patients had false positive rates of $< 20\%$ for a 100% true positive rate with data from brain state 1. For the clumping coefficients and the CV, data from both brain states 1 and 2 could be used to differentiate between SOZ and non-SOZ channels. The ROC curves and distributions of AUCs for both the clumping coefficient and CV parameter are shown in **Figures 3** and **4**.

The distributions of AUC values show some patients' SOZ channels are highly differentiated by CV parameters and clumping coefficients. Eight of 16 patients had AUC values greater than 0.9 for prediction by the clumping coefficient (with an AUC value of of 1 representing perfect prediction; see Bottom of **Figure 3**). Nine of 16 patients had AUC values greater than 0.9 for prediction via the CV (see Bottom of **Figure 4**).

Importantly, the prediction of SOZ by the clumping coefficient does not clearly depend upon localization of channels using CT and MRI scans to identify grey matter channels and exclude other channels from the analysis. Small clumping coefficients were predictive of SOZ even in those patients for which we did not exclude channels that were outside the brain. In contrast, prediction via the CV was degraded by the inclusion of channels outside of neural tissue, as shown in the Top Right of **Figure 4**.

By combining all channels across all patients, we can also obtain information about the general predictability of SOZ using these parameters. The number of channels used in the models varied by patient (minimum of 41, maximum of 172, and a mean and standard deviation of 87 ± 39 channels across $N = 16$ patients, see **Table 1**), and the number of SOZ channels also varied by patient (minimum of 1, maximum of 14, and a mean and standard deviation of 7 ± 4). However, combining channels across patients (total channel count of $E = 1390$) provides estimates of the cutoff values for these parameters that could be used to evaluate SOZ during interictal periods in NREM sleep. The aggregate ROC curve for the clumping coefficient is shown

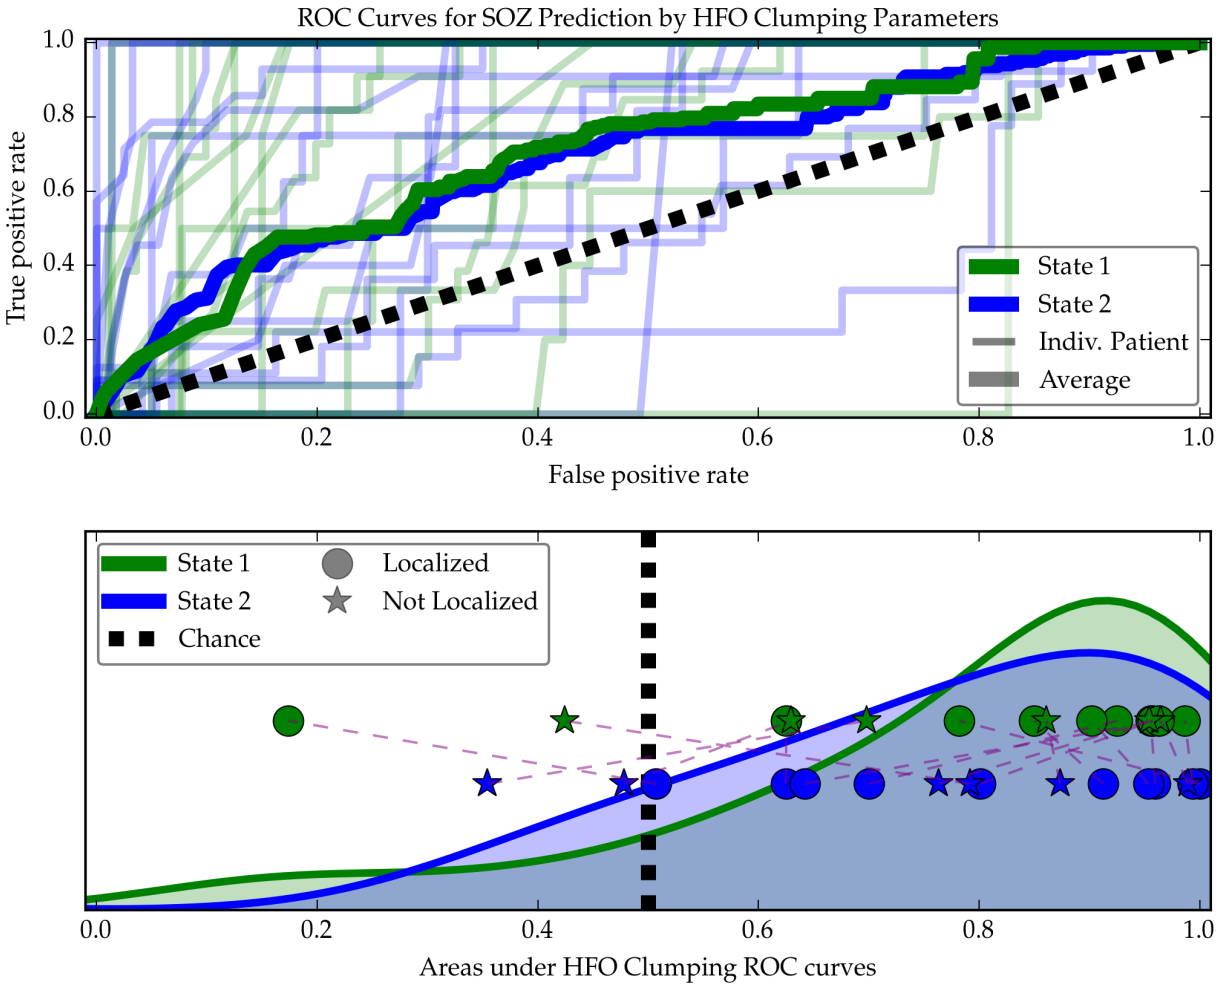

Figure 3: SOZ localization based on clumping coefficient. (Top) ROC curves for all patients ($N = 16$) when using small values of the HFO clumping coefficient ζ to identify SOZ channels in each brain state from Model 2. The brain state with the most delta (1-4 Hz) power in each patient was labeled brain state 1 (green lines) while the other model-found brain state was labeled brain state 2 (blue lines). ROC curves for individual subjects are displayed using fine lines, and the averages are shown in bold. The bold dashed line indicates an ROC at chance prediction. (Bottom) Distributions of AUC values based on clumping coefficients in both brain states for each patient. The exact AUC values are denoted as circles or stars on the x-axis while the shaded distributions are a density approximation from $N = 16$ values for each brain state. Circles denote patients where the analysis was performed exclusively on grey matter channels, while stars denote patients for which all channels were included. Dotted purple lines indicate the pairing of brain state 1 and 2 within a single patient. The bold dashed line indicates an AUC of 0.5.

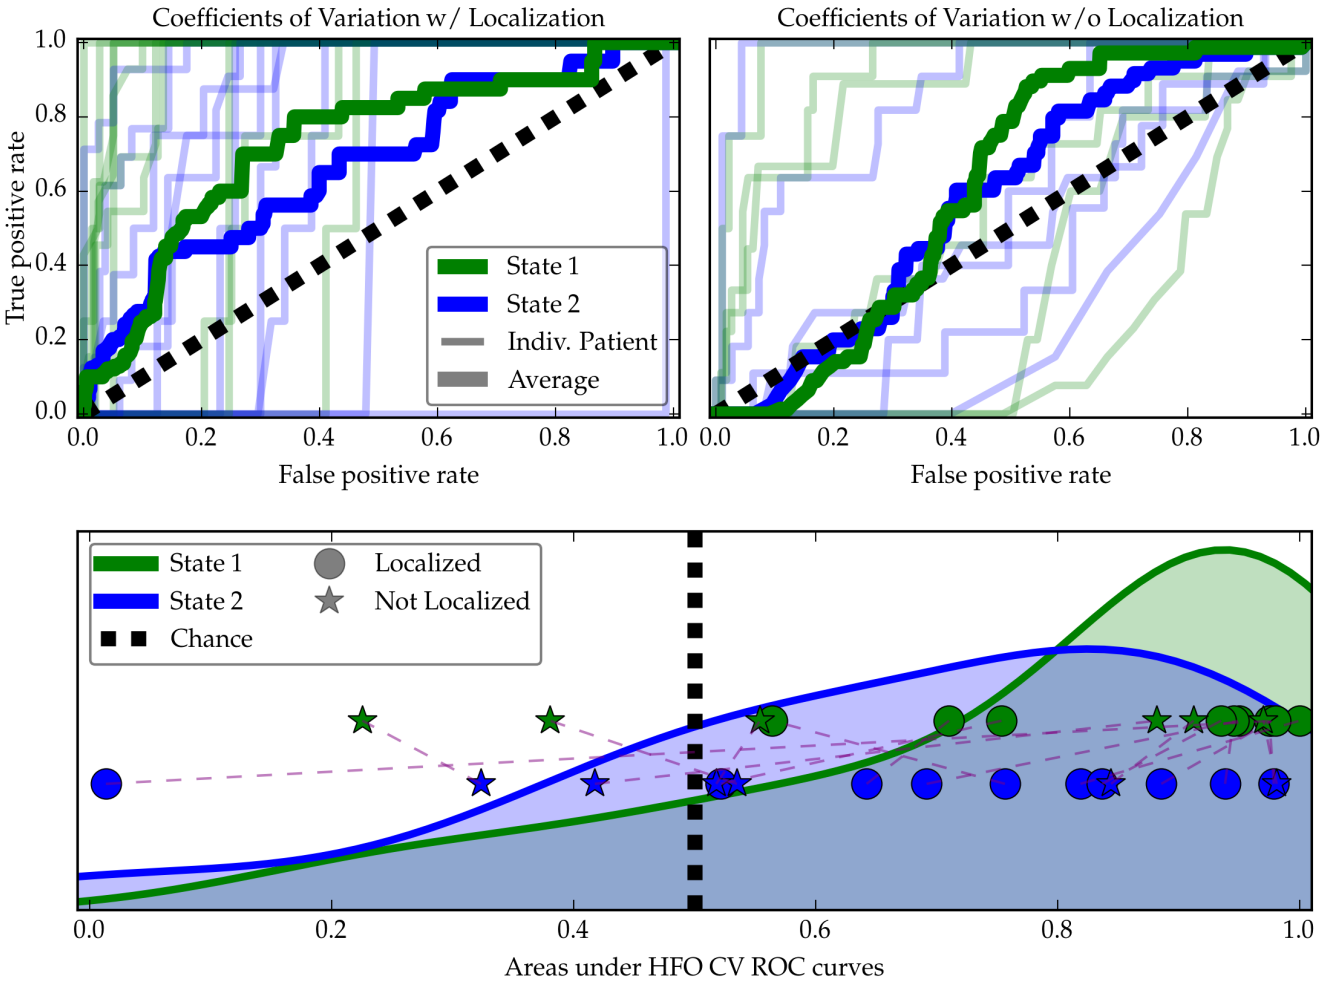

Figure 4: SOZ localization based on CV. (Top Left) ROC curves for patients in which only grey matter channels were used ($n = 10$) when using small values of the HFO coefficient of variation (CV) γ to identify SOZ channels in each brain state from Model 2. The brain state with the most delta (1-4 Hz) power in each patient was labeled brain state 1 (green lines) while the other model-found brain state was labeled brain state 2 (blue lines). ROC curves for individual subjects are displayed using fine lines, and the averages are shown in bold. The bold dashed line indicates an ROC at chance prediction. (Top Right) ROC curves for patients in which all channels were used to fit the models ($n = 6$) when using small values of the HFO coefficient of variation (CV) γ to identify SOZ channels in each brain state from Model 2. (Bottom) Distributions of AUC values based on CV in both brain states for each patient. The exact AUC values are denoted as circles or stars on the x-axis while the shaded distributions are a density approximation from $N = 16$ values for each brain state. Circles denote patients where the analysis was performed exclusively on grey matter channels, while stars denote patients for which all channels were included. Dotted purple lines indicate the pairing of brain state 1 and 2 within a single patient. The bold dashed line indicates an AUC of 0.5.

in **Figure 5**. Across all $N = 16$ patients, when clumping coefficients less than or equal to 1 ($\zeta \leq 1$) were treated as indicative of the SOZ in brain state 1, the false positive rate was only 32% across all channels, with a corresponding true positive rate of 81%. A true positive rate of 100% could be achieved by treating all clumping coefficients less than or equal to 3.52 ($\zeta \leq 3.52$) as indicative of the SOZ in brain state 1, although this would result in a 72% false positive rate.

4.3. Prediction of SOZ using HFO rate depends on electrode localization and brain state

The HFO rate could be used to identify the SOZ channels in different states and patients with success rates similar to the clumping coefficient and CV parameters. Large HFO rate parameters λ were able to identify the SOZ for 12 of 16 patients with a better than chance accuracy in at least one brain state, as shown in the bottom of **Figure 6**. In 7 patients, the SOZ could be identified exactly (100% true positive rate) with a $< 20\%$ false positive rate when using high HFO rates as cutoffs for SOZ. Note that only Model 2 results are reported here, as Model 1 parameter results were nearly identical to those from Model 2.

However the performance differences between state 1 and state 2 were more pronounced for HFO rate than for clumping coefficient or CV. Only seven of 16 patients had HFO rates that were predictive of SOZ in both state 1 and state 2 when considering AUCs of .6 or greater to indicate SOZ predictability. Recall that 12 of 16 patients had clumping coefficients that were predictive of SOZ in both states, and nine of 16 patients had CV values that were predictive of SOZ in both states.

These results also depended on localization of channels using CT and MRI scans to identify grey matter channels and exclude other channels from the analysis. This can be seen by comparing the Top Left and Top Right subplots in **Figure 6**. Three patients who did not have localization information had AUC values smaller than chance in both model-found brain states, indicating that *small* HFO rates were predictive in these three patients. Smaller than chance AUC values in brain state 1 can be observed in the Bottom subplot of **Figure 6** on the left side of the distribution.

*4.4. Assuming two brain states was **necessary** for SOZ prediction*

We found that assuming *at least* two brain states was necessary to obtain SOZ prediction from all three models. Although the rate of HFOs increases

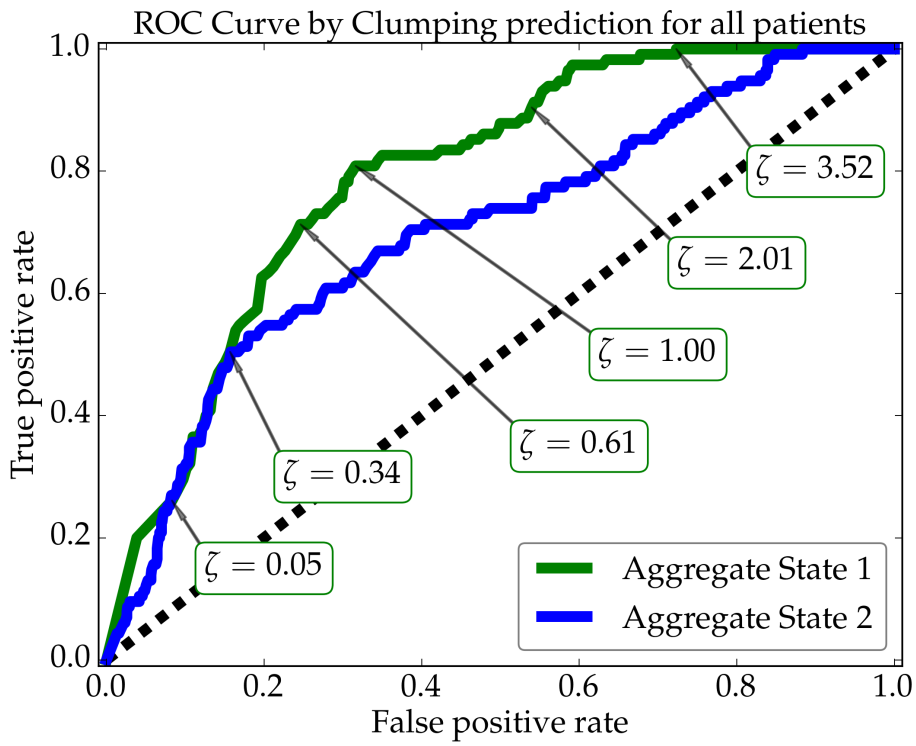

Figure 5: The aggregate ROC curves for all included channels ($E = 1390$) across all patients ($N = 16$) when using small values of the HFO clumping coefficient ζ to identify SOZ channels. The brain state with the most delta (1-4 Hz) power in each patient was labeled brain state 1 (green line) while the other model-found brain state was labeled brain state 2 (blue line). The clumping coefficients of all channels from all patients' model-found brain state 1 (green line) were used to find cutoff values of ζ shown in the text boxes, with clumping coefficients ζ less than or equal to these values resulting in the shown false positive and true positive rates. The blue line shows the resulting ROC curve when the clumping coefficients from all channels from all patients' model-found brain state 2 were used to find cutoff values. The bold dashed line indicates an ROC at chance prediction.

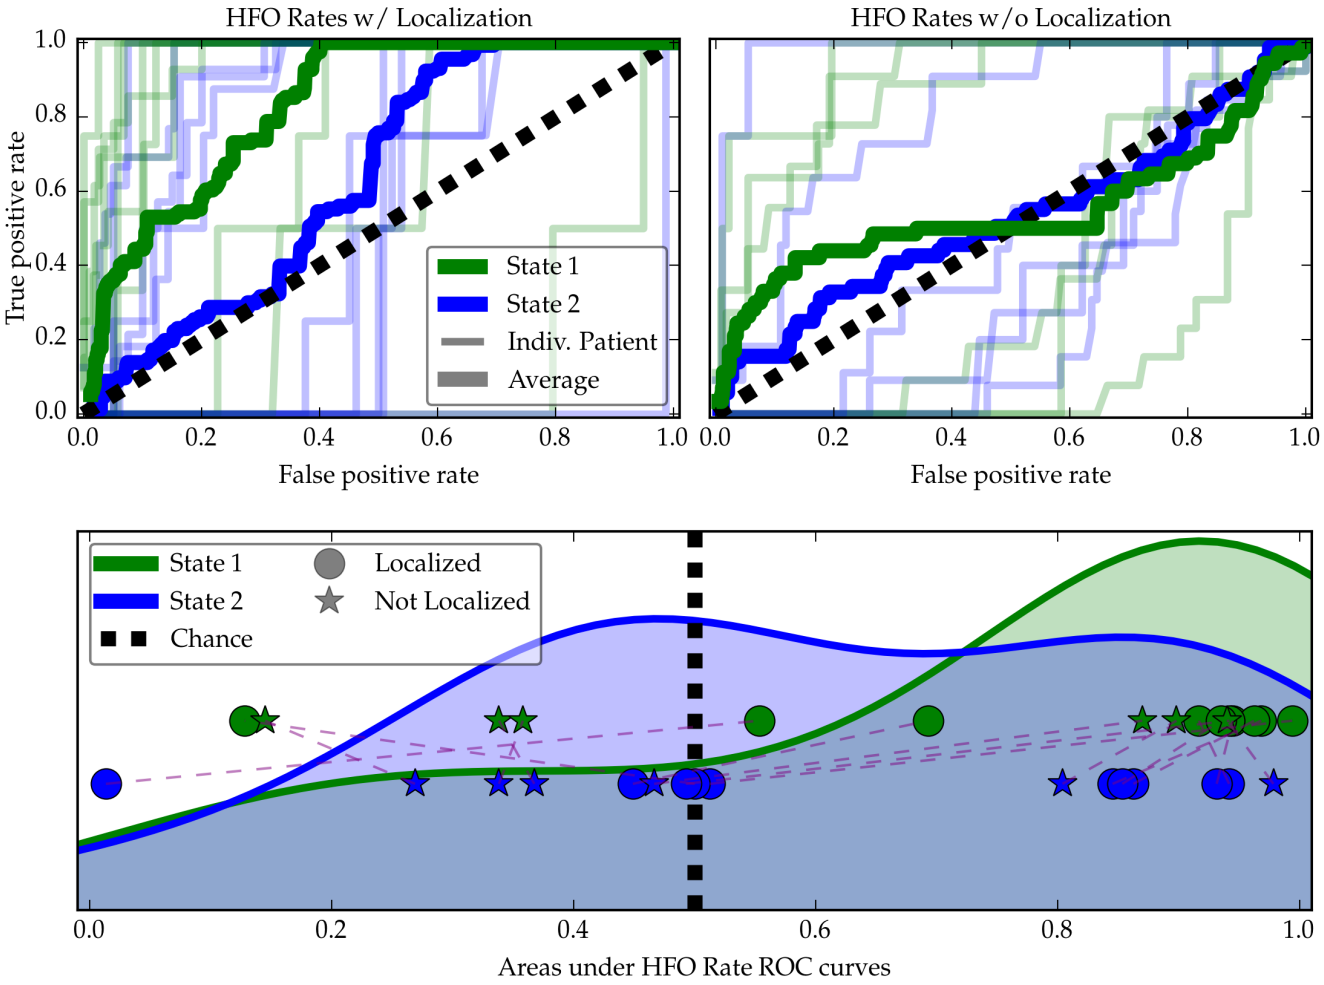

Figure 6: SOZ localization based on HFO rate. (Top Left) ROC curves for patients in which only grey matter channels were used ($n = 10$) when using high HFO rates λ to identify SOZ channels in each brain state from Model 2. The brain state with the most delta (1-4 Hz) power in each patient was labeled brain state 1 (green lines) while the other model-found brain state was labeled brain state 2 (blue lines). ROC curves for individual subjects are displayed using fine lines, and the averages are shown in bold. The bold dashed line indicates an ROC at chance prediction. (Top Right) ROC curves for patients in which all channels were used to fit the models ($n = 6$) when using high HFO rates to identify SOZ channels in each brain state from Model 2. (Bottom) Distributions of AUC values based on HFO rate in both brain states for each patient. The exact AUC values are denoted as circles or stars on the x-axis while the shaded distributions are a density approximation from $N = 16$ values for each brain state. Circles denote patients where the analysis was done exclusively on grey matter channels, while stars denote patients for which all channels were included. Dotted purple lines indicate the pairing of brain state 1 and 2 within a single patient. The bold dashed line indicates an AUC of 0.5.

during NREM sleep (von Ellenrieder et al., 2017), the classification accuracy based on the two brain states could be similar if the relative rates between channels remain the same. To test whether the model-derived brain states captured independent information about HFOs, we calculated the Pearson correlation between the two brain states for HFO rates λ, clumping coefficients ζ, and coefficients of variation γ. In the majority of patients, the correlation coefficients of these measures indicated that HFO dynamics were different across the two brain states, although there was a large range of correlation values. This meant that, in a subset of patients, both brain states contained useful information for SOZ prediction using HFO dynamics. The Pearson correlations between brain states for all patients were as follows: $\rho_\zeta = 0.45 \pm 0.25$ (mean $\pm$ standard deviation) for clumping coefficients, $\rho_\lambda = 0.59 \pm 0.24$ for HFO rates, and $\rho_\gamma = 0.68 \pm 0.24$ for CV. Furthermore, the distributions of AUC values were different for the two states, as can be seen by noting the path of the purple dotted lines denoting the pairing of brain states 1 and 2 in the bottom subplots of **Figures 3**, **4**, and **6**.

*4.5. Two brain states were **sufficient** for prediction but do not describe all HFO dynamics*

Each model was fit using Markov Chain Monte Carlo in JAGS with six chains of 5,200 samples each. This was performed in parallel with 200 burn-in samples and a thinning parameter of 10. This procedure resulted in $(5{,}200 - 200)/10 = 500$ posterior samples from each chain for each parameter. For those patients whose data did not converge to two states, two-state models were enforced by (1) assuming models with two brain states when performing the parameter fitting procedure on each patients' data and (2) keeping the most likely two-state models by removing non-converging Markov Chains in order to achieve convergence across all kept chains. Non-convergence was judged by looking for chain time courses that did not converge to a one-peaked posterior distribution or did not match the majority of other chains. We also calculated the Gelman-Rubin statistics, $\hat{R}$, for each parameter which compare the estimated between-chain and within-chain variances (Gelman and Rubin, 1992). Note that removing chains is an unorthodox method in Bayesian analysis, and *does not strictly guarantee model convergence.* The posterior samples from each chain of 500 samples were combined to form one posterior sample between 500 and 3,000 samples for each parameter in each model.

While removing non-converging Markov Chains to enforce two brain states in each model was unorthodox, this procedure enabled prediction of the SOZ, especially when using the clumping coefficient and coefficient of variation as shown in **Figures 3** and **4**. Thus, two state models are sufficient to *predict SOZ*. However the non-convergence results imply that a two-state model is *not* sufficient to *describe all HFO dynamics*. Six patients had Model 2 converge in all six chains, four patients had Model 2 converge in five of six chains, five patients had Model 2 converge in half the chains, while one patient's model only "converged" with one chain.

4.6. Comparison of SOZ prediction to surgical outcome

We used the Engel Outcome Scale (Engel, 1993) to quantify patients' surgical outcome. The Engel Outcome was compared to the classification of SOZ using three HFO parameters obtained from the mixture models (namely rate λ, coefficient of variation γ, and clumping coefficient ζ). Specifically, we compared total area under the ROC curve to Engel Outcome. We did not find a clear relationship between Engel outcome and SOZ prediction using any of the three parameters (see **Figure 7**). As a reminder, the entirety of SOZ tissue was not removed in all patients and some patients were treated via RNS (refer to **Table 1**). The diagnoses and treatments also varied across patients, some of whom had frontal lobe epilepsy or bilateral temporal lobe epilepsy. Such scenarios might explain a lack of correspondence between the Engel Outcome Scale and AUC. However, we repeated the analysis after removal of the three patients with frontal lobe epilepsy to gauge their impact on the results, and the results were nearly identical.

5. Discussion

5.1. Towards automatic classification of SOZ with interictal HFOs

Here we show that mixture modeling of HFO dynamics can help clinicians extract SOZ information during interictal periods in an unsupervised manner. Such tools may be valuable for localization of epileptogenic tissue during surgical planning. We found that some patients exhibited low correlations of HFO parameters between brain states, indicating that determining the patient's brain state during the time of the recording is critical.

In particular for HFO analysis, it may be desirable to identify periods of NREM sleep. In these patients, fitting mixture models is an effective way of obtaining information about HFO dynamics without the need for

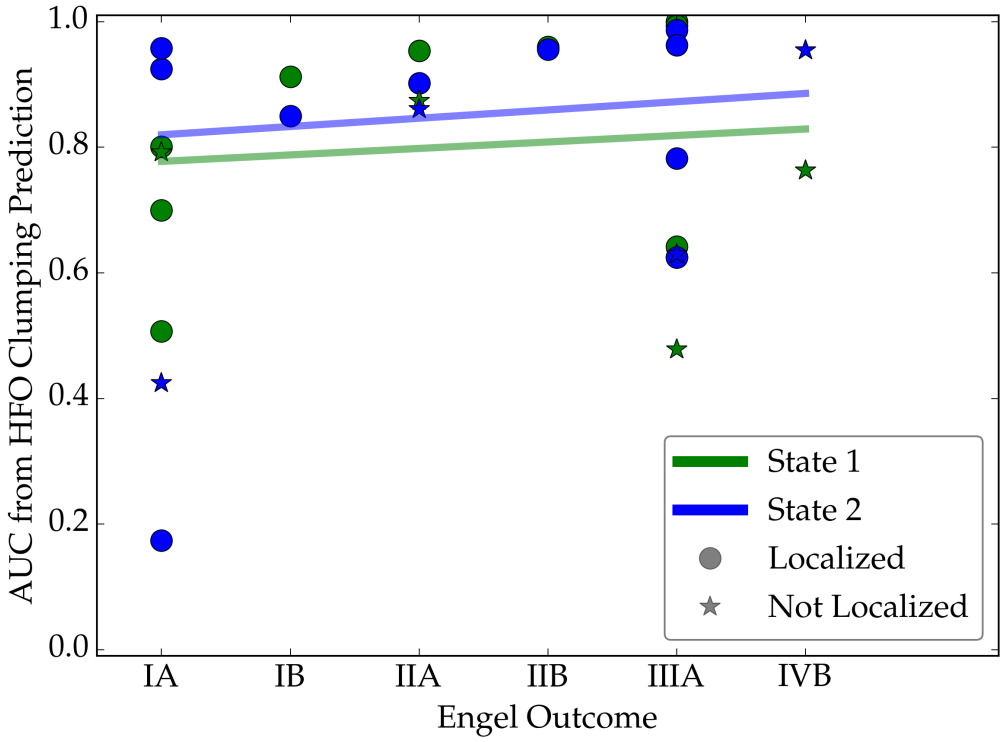

Figure 7: There was no clear relationship between Engel outcome and SOZ classification accuracy using the HFO clumping coefficient. The brain state with the most delta (1-4 Hz) power in each patient was labeled brain state 1 (green symbols) while the other model-found brain state was labeled brain state 2 (blue symbols). AUC values are denoted as circles or stars. Circles denote results based only grey matter channels, while stars denote results in which all channels were used. The lines denote least squares fits of the ordered Engel classifications versus AUC values for brain state 1 (green line) and brain state 2 (blue line) across patients.

concurrent EEG and manual sleep staging. Using the techniques presented here, there was no need to sleep stage the data (such as in von Ellenrieder et al., 2017) because the Poisson- and Negative Binomial-mixture models identified changing HFO dynamics automatically. Approximate sleep-stages were automatically obtained as a result of the distribution demixing. In patients with high correlations across model-fit brain states, the consistency of the parameters may make it unnecessary to fit a brain state model or sleep stage the data. However, we do not yet have the ability to identify which patients will require fitting brain state models.

The similarity of HFO rates in REM sleep compared to HFO rates during wakefulness has previously been shown (Staba et al., 2004). In two patients we found that the generators of HFOs during REM and wakefulness are often similar within each channel. However, HFOs provided more accurate SOZ localization during brain state 1, the brain state with larger delta power (1-4 Hz). In at least half the patients, this model-derived brain state likely reflects NREM sleep in patients due the separation of delta power across the two states. We confirmed that model-derived brain state 1 did reflect NREM sleep in two of these patients whose data was sleep staged. This fact further emphasizes the need to differentiate NREM from REM and awake prior to analysis, to only analyze HFO rates coincident with high delta (1-4 Hz) power as a proxy for NREM sleep, or to use automatic iEEG sleep-staging (Reed et al., 2017) and/or SOZ classification with HFO mixture modeling discussed in this paper. It is also possible that the mixture modeling captures interictal HFO dynamics *independent of sleep stage* that are predictive of SOZ. However this possibility should be explored further in other datasets.

5.2. HFO clumping is a more reliable predictor than HFO rate

HFOs that *do not* occur in bursts or clumps (perhaps due to a Poisson process or at regular intervals) are more consistently predictive of SOZ than a high rate of HFOs. The clumping coefficient and coefficient of variation (CV) were measured using hierarchical mixture models. Specifically, we found two clumping coefficients and CVs per iEEG channel using a model of two brain states obtained from a mixture of Negative Binomial distributions of HFO occurrences. Small clumping coefficients were predictive of SOZ in most patients, usually in the brain state corresponding to large delta power and possibly reflective of NREM sleep. Only two patients' clumping coefficients had AUC smaller than .62 using clumping coefficients from brain

state 1. This result is consistent with previous research that shows pathological HFOs are modulated by high amplitude, low frequency background, especially during sleep (Kerber et al., 2014; Frauscher et al., 2015; von Ellenrieder et al., 2016; Motoi et al., 2018). For most patients, the clumping coefficient was predictive in *both states* (12 of 16 had AUC $> .6$ in both states), suggesting that the oscillatory dynamics of HFOs may be a more reliable predictor of epileptogenic tissue than high HFO rates when the sleep state of the patient is unknown.

HFO rate was also informative of the SOZ, but it was less consistently predictive across patients than HFO clumping coefficients and coefficients of variation. In over half of the patients, HFO rates were predictive of SOZ (AUC $> .6$) only during one brain state. Furthermore, precise localization was necessary to remove channels not located in grey matter. These channels were likely subject to many false positive HFO detections. This implies that searching for channels with high HFO rates to identify the SOZ during interictal periods might lead to incorrect assessments of the SOZ if all implanted channels are considered. Our results also suggest that if HFO rates are to be used in prediction of the SOZ, they should be assessed during NREM sleep. This supports previous findings in the field (Dümpelmann et al., 2015; von Ellenrieder et al., 2016, 2017).

5.3. Limitations of this study

Differences across patients and channels (such as differences in electrode size, differences in brain shape and volume conduction, differences in disease state, etc.) may all play a role in the potential of HFOs to predict pathological tissue. However, the methods presented in this paper were promising for the limited number of patients explored in this study. The importance of validating the predictive nature of these methods in additional patients is obvious and cannot be overstated. Moreover, future studies should include a more detailed analysis of electrodes within the resected volume in order to make a quantitative comparison to surgical outcome. This is a more valuable test of clinical utility, as it evaluates the ability of the quantitative method to identify the epileptogenic zone, rather than the seizure onset zone (for which standard clinical criteria already exist).

Our choice of automatic detection algorithm and detection parameters will also have a significant impact on the results. We chose a simple algorithm, due to the large amount of data to be analyzed, but implementation

of a more complex algorithm with post-processing steps to reject false positive detections may improve the specificity of the detection and classification of SOZ channels. For example, if the HFOs occur in a regular, oscillatory pattern, the temporal dynamics may appear more random or clumped with the addition of false positives due to artifacts.

In our analysis, we treated all detected events equally, without attempting to separate pathological and physiological HFOs. It is possible that these two types of HFOs have similar rates, but different temporal dynamics, in which case our proposed method could help distinguish between them. However, here we could only classify events as being outside the SOZ, which would include both physiological HFOs and artifacts. Therefore, this question must be more explicitly studied with cognitive paradigms to elicit physiological HFOs or analysis focused on the fast ripple frequency band (250-500 Hz).

The potential for applying these methods to scalp EEG data is uncertain, although there is some theoretical evidence that this could be possible (von Ellenrieder et al., 2014). However the presence of pervasive muscle artifact in scalp EEG makes the practical application of scalp-HFO detectors difficult (Nunez et al., 2016). It has been found experimentally that signals originating in the cortex have a magnitude as much as 200 times lower than the electrical signal from muscle activity (Fitzgibbon et al., 2016). This is especially challenging in areas of interest to many epileptologists, such as the temporal lobe, due to the proximity of neck and facial musculature.

There may also be differences in pathological HFO dynamics between intracranial depth electrodes and cortical surface electrodes. These two types of sensors record from different amounts of cortical depth and volume, and intrinsic differences in neural behavior between different spatial scales could exist (Nunez et al., 2019). We might even expect differences in neural behavior between iEEG electrodes of different diameters at similar locations within the same patient due to these reasons (Nunez et al., 2019). There have been conflicting reports on the effect of electrode size on the ability to measure HFOs (Worrell et al., 2008; Châtillon et al., 2013). Lastly, we would expect depth iEEG electrodes to be more likely to be contaminated by noise, as the channels at the ends are sometimes outside the brain. In this study, we collapsed across all types of intracranial electrodes.

5.4. Future improvements to algorithmic implementation

Faster methods of fitting Poisson and Negative Binomial mixture models are necessary for these methods to be applied in a clinical setting. In

this study, we wished to fit hierarchical models in order to understand the relationship between channels and patients. However, in future studies, simple algorithms to fit mixture models of Negative Binomial distributions and other distributions, such as presented by (Nagode, 2015), may be sufficient.

Some models presented here did not converge as judged by the Gelman-Rubin statistic, $\hat{R}$, although the median posterior parameters were still informative for SOZ classification. This seemed to be due to the non-convergence of specific HFO rates and oscillatory dynamics for subsets of channels in some patients. This could be caused by artifacts being introduced to the HFO rates by the automatic detector or due to actual physiological or pathological deviations from that channel's rate in that brain state. It could be that the adaptive noise floor, which changed every 5 minutes within each channel using our automatic HFO detector (Charupanit and Lopour, 2017), injected artifactual HFO dynamics into the models. It is possible that fitting an HFO detector and Poisson / Negative Binomial hierarchical models *concurrently* would alleviate this convergence issue.

Model convergence is usually a bare minimum for hierarchical Bayesian model building. However because the outcome of this study was SOZ classification and the non-converged models were still able to classify SOZ, the results are still clinically relevant. Models that allow "noise" in the HFO dynamics to occur with some limited frequency could alleviate this issue. This could facilitate model convergence and may even yield better classification of the SOZ. In pilot analyses of models with three or four brain states, we were unable to fit these mixture models in JAGS with sufficient convergence of chains. Thus, the resulting posterior distributions of HFO parameters were difficult to interpret. We are unsure if the data would be better described by a model with more brain states. Future work should seek to expand the number of brain states while allowing for artifactual HFO dynamics.

Acknowledgements

This work was supported by a grant from the American Epilepsy Society.

Conflict of interest statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

Bendat, J. S. and Piersol, A. G. (2011). *Random data: analysis and measurement procedures*, volume 729. John Wiley & Sons.

Bernardo, D., Nariai, H., Hussain, S. A., Sankar, R., Salamon, N., Krueger, D. A., Sahin, M., Northrup, H., Bebin, E. M., Wu, J. Y., et al. (2018). Visual and semi-automatic non-invasive detection of interictal fast ripples: A potential biomarker of epilepsy in children with tuberous sclerosis complex. *Clinical Neurophysiology*, 129(7):1458–1466.

Blanco, J. A., Stead, M., Krieger, A., Stacey, W., Maus, D., Marsh, E., Viventi, J., Lee, K. H., Marsh, R., Litt, B., et al. (2011). Data mining neocortical high-frequency oscillations in epilepsy and controls. *Brain*, 134(10):2948–2959.

Boehm, U., Marsman, M., Matzke, D., and Wagenmakers, E.-J. (2018). On the importance of avoiding shortcuts in applying cognitive models to hierarchical data. *Behavior research methods*, 50(4):1614–1631.

Bragin, A., Wilson, C. L., Staba, R. J., Reddick, M., Fried, I., and Engel Jr, J. (2002). Interictal high-frequency oscillations (80–500hz) in the human epileptic brain: Entorhinal cortex. *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, 52(4):407–415.

Charupanit, K. and Lopour, B. (2017). A simple statistical method for the automatic detection of ripples in human intracranial EEG. *Brain topography*, 30(6):724–738.

Châtillon, C., Zelmann, R., Hall, J., Olivier, A., Dubeau, F., and Gotman, J. (2013). Influence of contact size on the detection of hfos in human intracerebral eeg recordings. *Clinical Neurophysiology*, 124(8):1541–1546.

Cohen-Gadol, A. A., Wilhelmi, B. G., Collignon, F., White, J. B., Britton, J. W., Cambier, D. M., Christianson, T. J., Marsh, W. R., Meyer, F. B., and Cascino, G. D. (2006). Long-term outcome of epilepsy surgery among 399 patients with nonlesional seizure foci including mesial temporal lobe sclerosis. *Journal of neurosurgery*, 104(4):513–524.

Cook, J. D. (2009). Notes on the negative binomial distribution.

Dümpelmann, M., Jacobs, J., and Schulze-Bonhage, A. (2015). Temporal and spatial characteristics of high frequency oscillations as a new biomarker in epilepsy. *Epilepsia*, 56(2):197–206.

Engel, J. (1993). Outcome with respect to epileptic seizures. In Engel, J., editor, *Surgical treatment of the epilepsies*, pages 609–621. Raven Press, New York,NY, 2 edition.

Engel, J. (2018). The current place of epilepsy surgery. *Current Opinion in Neurology*, 31(2):192–197.

Fedele, T., Ramantani, G., and Sarnthein, J. (2019). High frequency oscillations as markers of epileptogenic tissue-end of the party? *Clinical neurophysiology: official journal of the International Federation of Clinical Neurophysiology*, 130(5):624–626.

Fitzgibbon, S., DeLosAngeles, D., Lewis, T., Powers, D., Grummett, T., Whitham, E., Ward, L., Willoughby, J., and Pope, K. (2016). Automatic determination of EMG-contaminated components and validation of independent component analysis using EEG during pharmacologic paralysis. *Clinical Neurophysiology*, 127(3):1781–1793.

Frauscher, B., Bartolomei, F., Kobayashi, K., Cimbalnik, J., van t Klooster, M. A., Rampp, S., Otsubo, H., Höller, Y., Wu, J. Y., Asano, E., et al. (2017). High-frequency oscillations: The state of clinical research. *Epilepsia*, 58(8):1316–1329.

Frauscher, B., von Ellenrieder, N., Ferrari-Marinho, T., Avoli, M., Dubeau, F., and Gotman, J. (2015). Facilitation of epileptic activity during sleep is mediated by high amplitude slow waves. *Brain*, 138(6):1629–1641.

Frauscher, B., von Ellenrieder, N., Zelmann, R., Rogers, C., Nguyen, D. K., Kahane, P., Dubeau, F., and Gotman, J. (2018). High-frequency oscillations in the normal human brain. *Annals of neurology*, 84(3):374–385.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and Hall/CRC, 3 edition.

Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, pages 457–472.

Gliske, S. V., Irwin, Z. T., Chestek, C., Hegeman, G. L., Brinkmann, B., Sagher, O., Garton, H. J., Worrell, G. A., and Stacey, W. C. (2018). Variability in the location of high frequency oscillations during prolonged intracranial EEG recordings. *Nature communications*, 9(1):2155.

Gotman, J. (2018). Oh surprise! fast ripples on scalp EEG. *Clinical neurophysiology: official journal of the International Federation of Clinical Neurophysiology*, 129(7):1449.

Greer, S. M. and Saletin, J. M. (2011). sleepsmg [computer software]. `https://sourceforge.net/projects/sleepsmg/`.

Guirgis, M., Chinvarun, Y., Del Campo, M., Carlen, P. L., and Bardakjian, B. L. (2015). Defining regions of interest using cross-frequency coupling in extratemporal lobe epilepsy patients. *Journal of neural engineering*, 12(2):026011.

Guragain, H., Cimbalnik, J., Stead, M., Groppe, D. M., Berry, B. M., Kremen, V., Kenney-Jung, D., Britton, J., Worrell, G. A., and Brinkmann, B. H. (2018). Spatial variation in high-frequency oscillation rates and amplitudes in intracranial EEG. *Neurology*, 90(8):e639–e646.

Helfrich, R. F., Fiebelkorn, I. C., Szczepanski, S. M., Lin, J. J., Parvizi, J., Knight, R. T., and Kastner, S. (2018). Neural mechanisms of sustained attention are rhythmic. *Neuron*, 99(4):854–865.

Ibrahim, G. M., Wong, S. M., Anderson, R. A., Singh-Cadieux, G., Akiyama, T., Ochi, A., Otsubo, H., Okanishi, T., Valiante, T. A., Donner, E., et al. (2014). Dynamic modulation of epileptic high frequency oscillations by the phase of slower cortical rhythms. *Experimental neurology*, 251:30–38.

Jacobs, J., LeVan, P., Chander, R., Hall, J., Dubeau, F., and Gotman, J. (2008). Interictal high-frequency oscillations (80–500 hz) are an indicator of seizure onset areas independent of spikes in the human epileptic brain. *Epilepsia*, 49(11):1893–1907.

Jacobs, J., Wu, J. Y., Perucca, P., Zelmann, R., Mader, M., Dubeau, F., Mathern, G. W., Schulze-Bonhage, A., and Gotman, J. (2018). Removing high-frequency oscillations: a prospective multicenter study on seizure outcome. *Neurology*, 91(11):e1040–e1052.

Jacobs, J., Zijlmans, M., Zelmann, R., Chatillon, C.-É., Hall, J., Olivier, A., Dubeau, F., and Gotman, J. (2010). High-frequency electroencephalographic oscillations correlate with outcome of epilepsy surgery. *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, 67(2):209–220.

Jefferys, J. G., de La Prida, L. M., Wendling, F., Bragin, A., Avoli, M., Timofeev, I., and da Silva, F. H. L. (2012). Mechanisms of physiological and epileptic hfo generation. *Progress in neurobiology*, 98(3):250–264.

Kerber, K., Dümpelmann, M., Schelter, B., Le Van, P., Korinthenberg, R., Schulze-Bonhage, A., and Jacobs, J. (2014). Differentiation of specific ripple patterns helps to identify epileptogenic areas for surgical procedures. *Clinical Neurophysiology*, 125(7):1339–1345.

Kobayashi, K., Akiyama, T., Oka, M., Endoh, F., and Yoshinaga, H. (2015). A storm of fast (40–150 h z) oscillations during hypsarrhythmia in west syndrome. *Annals of neurology*, 77(1):58–67.

Köhling, R. and Staley, K. (2011). Network mechanisms for fast ripple activity in epileptic tissue. *Epilepsy research*, 97(3):318–323.

Kovac, S., Vakharia, V. N., Scott, C., and Diehl, B. (2017). Invasive epilepsy surgery evaluation. *Seizure*, 44:125–136.

Kwan, P. and Brodie, M. J. (2000). Early identification of refractory epilepsy. *New England Journal of Medicine*, 342(5):314–319.

Miasko, T. (2017). pyjags (version 1.2.2) [computer software]. `https://github.com/tmiasko/pyjags`.

Motoi, H., Miyakoshi, M., Abel, T. J., Jeong, J.-W., Nakai, Y., Sugiura, A., Luat, A. F., Agarwal, R., Sood, S., and Asano, E. (2018). Phase-amplitude coupling between interictal high-frequency activity and slow waves in epilepsy surgery. *Epilepsia*, 59(10):1954–1965.

Nagode, M. (2015). Finite mixture modeling via rebmix. *Journal of Algorithms and Optimization*, 3(2):14–28.

Noachtar, S. and Borggraefe, I. (2009). Epilepsy surgery: a critical review. *Epilepsy & Behavior*, 15(1):66–72.

Noe, K., Sulc, V., Wong-Kisiel, L., Wirrell, E., Van Gompel, J. J., Wetjen, N., Britton, J., So, E., Cascino, G. D., Marsh, W. R., et al. (2013). Long-term outcomes after nonlesional extratemporal lobe epilepsy surgery. *JAMA neurology*, 70(8):1003–1008.

Nunez, M. D., Nunez, P. L., and Srinivasan, R. (2016). Electroencephalography (EEG): neurophysics, experimental methods, and signal processing. In Ombao, H., Linquist, M., Thompson, W., and Aston, J., editors, *Handbook of Neuroimaging Data Analysis*, pages 175–197. Chapman & Hall/CRC.

Nunez, P. L., Nunez, M. D., and Srinivasan, R. (2019). Multi-scale neural sources of EEG: Genuine, equivalent, and representative. a tutorial review. *Brain Topography*, 32(2):193–214.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003), Vienna, Austria.*

Reed, C. M., Birch, K. G., Kamiński, J., Sullivan, S., Chung, J. M., Mamelak, A. N., and Rutishauser, U. (2017). Automatic detection of periods of slow wave sleep based on intracranial depth electrode recordings. *Journal of neuroscience methods*, 282:1–8.

Roehri, N., Pizzo, F., Lagarde, S., Lambert, I., Nica, A., McGonigal, A., Giusiano, B., Bartolomei, F., and Bénar, C.-G. (2018). High-frequency oscillations are not better biomarkers of epileptogenic tissues than spikes. *Annals of neurology*, 83(1):84–97.

Ryvlin, P., Cross, J. H., and Rheims, S. (2014). Epilepsy surgery in children and adults. *The Lancet Neurology*, 13(11):1114–1126.

Staba, R. J. and Bragin, A. (2011). High-frequency oscillations and other electrophysiological biomarkers of epilepsy: underlying mechanisms. *Biomarkers in medicine*, 5(5):545–556.

Staba, R. J., Wilson, C. L., Bragin, A., Fried, I., and Engel Jr, J. (2002). Quantitative analysis of high-frequency oscillations (80–500 hz) recorded in human epileptic hippocampus and entorhinal cortex. *Journal of neurophysiology*, 88(4):1743–1752.

Staba, R. J., Wilson, C. L., Bragin, A., Jhung, D., Fried, I., and Engel Jr, J. (2004). High-frequency oscillations recorded in human medial temporal lobe during sleep. *Annals of neurology*, 56(1):108–115.

Stevenson, R. F., Zheng, J., Mnatsakanyan, L., Vadera, S., Knight, R. T., Lin, J. J., and Yassa, M. A. (2018). Hippocampal ca1 gamma power predicts the precision of spatial memory judgments. *Proceedings of the National Academy of Sciences*, 115(40):10148–10153.

Stolk, A., Griffin, S., van der Meij, R., Dewar, C., Saez, I., Lin, J. J., Piantoni, G., Schoffelen, J.-M., Knight, R. T., and Oostenveld, R. (2017). Integrated analysis of anatomical and electrophysiological human intracranial data. *bioRxiv*, page 230912.

von Ellenrieder, N., Beltrachini, L., Perucca, P., and Gotman, J. (2014). Size of cortical generators of epileptic interictal events and visibility on scalp EEG. *Neuroimage*, 94:47–54.

von Ellenrieder, N., Dubeau, F., Gotman, J., and Frauscher, B. (2017). Physiological and pathological high-frequency oscillations have distinct sleep-homeostatic properties. *NeuroImage: Clinical*, 14:566–573.

von Ellenrieder, N., Frauscher, B., Dubeau, F., and Gotman, J. (2016). Interaction with slow waves during sleep improves discrimination of physiologic and pathologic high-frequency oscillations (80–500 hz). *Epilepsia*, 57(6):869–878.

Wang, S., So, N. K., Jin, B., Wang, I. Z., Bulacio, J. C., Enatsu, R., Dai, S., Chen, Z., Gonzalez-Martinez, J., and Najm, I. M. (2017). Interictal ripples nested in epileptiform discharge help to identify the epileptogenic zone in neocortical epilepsy. *Clinical Neurophysiology*, 128(6):945–951.

Weiss, S. A., Banks, G. P., McKhann Jr, G. M., Goodman, R. R., Emerson, R. G., Trevelyan, A. J., and Schevon, C. A. (2013). Ictal high frequency oscillations distinguish two types of seizure territories in humans. *Brain*, 136(12):3796–3808.

Worrell, G. A., Gardner, A. B., Stead, S. M., Hu, S., Goerss, S., Cascino, G. J., Meyer, F. B., Marsh, R., and Litt, B. (2008). High-frequency oscillations in human temporal lobe: simultaneous microwire and clinical macroelectrode recordings. *Brain*, 131(4):928–937.

Zack, M. M. and Kobau, R. (2017). National and state estimates of the numbers of adults and children with active epilepsy-united states, 2015. *Morbidity and Mortality Weekly Report*, 66(31):821–825.

Zelmann, R., Lina, J., Schulze-Bonhage, A., Gotman, J., and Jacobs, J. (2014). Scalp EEG is not a blur: it can see high frequency oscillations although their generators are small. *Brain topography*, 27(5):683–704.

Zheng, J., Anderson, K. L., Leal, S. L., Shestyuk, A., Gulsen, G., Mnatsakanyan, L., Vadera, S., Hsu, F. P., Yassa, M. A., Knight, R. T., et al. (2017). Amygdala-hippocampal dynamics during salient information processing. *Nature communications*, 8:14413.