Neural Fragility as an EEG Marker of the Seizure Onset Zone

Adam Li^{1,9}, Chester Huynh^{1,9}, Zachary Fitzgerald², Iahn Cajigas³, Damian Brusko³, Jonathan Jagid³, Angel Claudio⁴, Andres Kanner⁴, Jennifer Hopp⁵, Stephanie Chen⁵, Jennifer Haagensen⁵, Emily Johnson⁶, William Anderson⁶, Nathan Crone^{6,9}, Sara Inati⁷, Kareem Zaghoul⁷, Juan Bulacio², Jorge Gonzalez-Martinez^{8,*}, and Sridevi V. Sarma^{1,9,*}

¹Department of Biomedical Engineering, Johns Hopkins University, Baltimore, United States

²Epilepsy Center, Cleveland Clinic, Cleveland, United States

³Department of Neurosurgery, University of Miami Miller School of Medicine, Miami, United States

⁴Department of Neurology, University of Miami Miller School of Medicine, Miami, United States

⁵Neurology, University of Maryland Medical Center, Baltimore, United States

⁶Neurology, Johns Hopkins Hospital, Baltimore, United States

⁷Surgical Neurology Branch, NINDS, NIH, Bethesda MD

⁸Neurosurgery, and Epilepsy Center, University of Pittsburgh Medical Center, Pittsburgh, United States

⁹Institute for Computational Medicine, Johns Hopkins University, Baltimore, United States

*Co-senior authors

February 3, 2021

Abstract

Over 15 million epilepsy patients worldwide do not respond to drugs. Successful surgical treatment requires complete removal, or disconnection of the seizure onset zone (SOZ), brain region(s) where seizures originate. Unfortunately, surgical success rates vary between 30%-70% because no clinically validated biological marker of the SOZ exists. We develop and retrospectively validate a new EEG marker - *neural fragility* - in a retrospective analysis of 91 patients by using neural fragility of the annotated SOZ as a metric to predict surgical outcomes. Fragility predicts 43/47 surgical failures with an overall prediction accuracy of 76%, compared to the accuracy of clinicians being 48% (successful outcomes). In failed outcomes, we identify fragile regions that were untreated. When compared to 20 EEG features proposed as SOZ markers, fragility outperformed in predictive power and interpretability suggesting neural fragility as an EEG biomarker of the SOZ.

1 Introduction

Over 15 million epilepsy patients worldwide and 1 million in the US suffer from drug-resistant epilepsy (DRE) [1, 2]. DRE is defined as continued seizures despite adequate trials of two tolerated appropriately chosen anti-epileptic drugs [3]. DRE patients have an increased risk of sudden death and are frequently hospitalized, burdened by epilepsy-related disabilities, and the cost of their care is a significant contributor to the \$16 billion dollars spent annually in the US treating epilepsy patients [4]. Approximately 50% of DRE patients have focal DRE, where a specific brain region or regions, termed the epileptogenic zone (EZ), is necessary and sufficient for initiating seizures and whose removal (or disconnection) results in complete abolition of seizures [5, 6]. The EZ encompasses the clinically identified seizure onset zone (SOZ) and early propagation zone (EPZ). The brain regions associated with the *SOZ* demonstrate the earliest electrophysiological changes during a seizure event, and in general precede the clinical onset of seizures; and the EPZ regions are involved at the time of the earliest clinical (semiological) manifestations during a seizure event. Successful surgical and neuromodulatory treatments can stop seizures altogether or allow them to be controlled with medications [7], but outcomes for both treatments critically depend on accurate localization of the *SOZ*.

Localizing the *SOZ* also relies on accurate placement of electrodes such that they cover the EZ, and the ability to identify abnormalities in the intracranial EEG (iEEG) channels that may correlate to the *SOZ* with the naked eye. Unfortunately, even the most experienced clinicians are challenged because epilepsy is fundamentally a network disease, which cannot be entirely defined by the current methods of localization. Abnormal connections across several channels may constitute a more effective marker of the *SOZ* [8]. Localization thus lends itself to a data-driven network-based computational approach and several electroencephalogram (EEG) algorithms have been proposed to localize the *SOZ* from recordings. Many entail investigations of the spectral power in each iEEG channel including high frequency oscillations [9], but these approaches do not consider network properties of the brain because they treat each EEG channel independently. Others have proposed graph-based analysis of iEEG [10, 11, 12, 13, 14, 15], but these approaches fail to identify internal network properties that cause seizures to occur in the first place.

We propose a new EEG marker of the *SOZ*, neural fragility, of brain regions (conceptually described in Figure 1). To create the fragility marker, we first build a personalized dynamic model of the brain network from observed iEEG signals. The model is generative in that it can accurately reconstruct the patient’s iEEG recordings [16, 17]. We then calculate which network nodes are imbalanced, meaning they have more excitatory or less inhibitory influence on the network and thus can trigger seizures. Neural fragility measures the degree to which a node is imbalanced [18]. To evaluate neural fragility as a marker of the *SOZ*, we conducted a retrospective study using iEEG data from 91 patients treated across 5 epilepsy centers: Johns Hopkins Hospital, National Institute of Health, Cleveland Clinic, University of Maryland Medical Center and Jackson Memorial Hospital, University of Miami. In the study population, all DRE patients underwent invasive iEEG monitoring followed by surgical resection or laser ablation of the *SOZ* (44 success and 47 failure outcome). We demonstrate that neural fragility is higher/lower in electrode contacts within clinically annotated SOZs for success/failure patients. In addition, we compare fragility of iEEG nodes to 6 frequency-based and 14 graph theoretic features in a 10-fold nested-cross validation. Neural fragility has an area under the curve (AUC) discrimination score of 0.88 ± 0.064 , which is 13% better compared to the next best feature. In addition, it has a high degree of interpretability, which we demonstrate by computing an interpretability ratio suggesting that the spatiotemporal heatmaps of neural fragility are a robust iEEG biomarker of the *SOZ* that can incorporate seamlessly into the clinical workflow.

Figure 1: Neural fragility intuition (Top) iEEG traces in between seizures (left) and during a seizure (right). **(Bottom)** network schematic showing change in connectivity (right) in fragile node that causes seizure. Describes qualitatively the concept of neural fragility in the context of a dynamical iEEG network, with nodes representing excitatory (E) and inhibitory (I) population of neurons. From a dynamical systems point of view, such imbalance arises from a few fragile nodes causing instability of the network in the form of over-excitation, or under-inhibition. We define fragility of a network node to be the minimum perturbation applied to the node's weights on its neighbors before rendering the network unstable [18, 16]. In systems theory, stable systems return to a baseline condition when a node is perturbed. In contrast, unstable systems can oscillate and grow when a node is perturbed. Thus, this would correspond to a column perturbation of an A matrix fitted to the iEEG data. In the context of epilepsy, a fragile node is one that requires a smaller perturbation to lead to seizure activity.

Figure 2: Neural fragility in a 2-node network To build quantitative intuition on what neural fragility means in the context of a dynamical iEEG system, we construct a 2-node EEG network example with an excitatory (E) and inhibitory (I) population of neurons. For a qualitative description, see Supplementary Figure 1. Although iEEG electrodes are not excitatory or inhibitory, their dynamical impact on each other in a network can be characterized as excitatory and inhibitory (positive and negative weights, respectively). $x_I(t)$ and $x_E(t)$ are the EEG activity of the I and E neuronal population respectively. 'A' is a linear network model quantifying how each population affects the rest over time. Δ (i.e. the fragility), is the amount of change added to a node's connections. The fragility of the node is quantified as the minimal amount of change necessary to cause seizure-like phenomena. (a) shows a stable network without a perturbation added, such that the network responses due to an impulse at I result in a transient that reverts to baseline. (b) shows a perturbation added, but the network is still stable with a slightly larger transient when an impulse is applied to node I . Then (c) shows enough of a perturbation is added, such that the network becomes unstable; an impulse applied at node I results in oscillatory activity that does not quickly return to baseline. The magnitude of the Δ added in (c) is the fragility of node I (i.e. $\sqrt{8}$).

Figure 3: Clinical complexity and our experimental paradigm (a) Schematic of the difficulty of different epilepsy etiologies that might arise in DRE patients. Since there is no biomarker for the EZ and it is never observed directly, the network mechanisms that cause seizures are complex. Case clinical complexity ordered by increasing localization difficulty: lesional (1), focal temporal (2), focal extratemporal (3), and multi-focal (4) that are present in the dataset. These four categories simplify the possible epilepsy presentations, but provide a broad categorization of simple to complex cases observed in the clinic. (b) Schematic of our experimental design. (bottom row) Shows a simplified analogous workflow that clinicians take to evaluate their confidence in a proposed *SOZ* localization resulting in a surgery. During invasive monitoring, clinicians identify the *SOZ* from iEEG patterns (e.g. spiking/rhythmic activity). When possible, subsequent surgical resection or laser ablation, generally including the *SOZ* along with a variable extent of additional tissue, is performed. Post-operatively, patients are followed for 12+ months and categorized as either success, or failure as defined in Supplementary Materials and Methods - Dataset collection, resulting in an Engel or ILAE score. (top row) We evaluate various representations of iEEG in the form of spatiotemporal heatmaps, creating a partitioned summary of the clinically annotated *SOZ* around seizure onset, feed them into a Random Forest classifier and compute a probability of success (i.e. a confidence score) in the clinically hypothesized *SOZ*. The probability was then compared with the actual outcome of patients. These predictions can then be stratified based on the clinical covariates, such as the actual surgical outcome. For a feature to be an accurate representation of the underlying epileptic phenomena, the following assumptions are made. As a result of seizure freedom, assume that the clinically hypothesized *SOZ* was sufficient, and the probability of success has a high value. In contrast, if seizures continue, then the *SOZ* was not sufficient and the probability should have a low value.

2 Results

There is no clinically validated biomarker of the *SOZ*, yet it forms an important component for the location of the underlying EZ [19, 20]. This presents a serious challenge for clinicians to accurately localize the *SOZ* and has led to surgical success rates to vary between 30-70% despite large brain regions being removed. With no biomarker available, clinicians are driven to perform extensive evaluations with neuroimaging, clinical testing, and visual inspection of scalp EEG recordings (Figure 3). When non-invasive tests are inconclusive, patients undergo intracranial monitoring, during which intracranial EEG (iEEG) electrodes are either placed directly on the cortex or implanted into the brain (example in Supplemental Figure S1). iEEG provides high temporal resolution data that enables clinicians to visually detect abnormal activity, such as spikes and high frequency bursts, in between seizures (interictal) and during seizures (ictal). Specifically, clinicians attempt to identify the electrodes involved in the *SOZ* and early spread [21]. Surgical resection is then performed on the basis of the hypothesis, unless the region overlaps with eloquent cortex.

We analyze every patient's iEEG using fragility and 20 other baseline features, resulting in spatiotemporal heatmaps for every feature. The baseline features include spectral power in various frequency bands (e.g. delta band 1-4 Hz) and specific graph measures of bivariate correlation measures (e.g. eigenvector centrality and degree of correlation and coherence matrices), which have been previously reported in the literature to correlate to the *SOZ* [11, 12, 13, 10]. We consider all of these as potential EEG features that can represent the data in different ways, potentially useful for *SOZ* localization.

2.1 Description of Neural Fragility

Neural fragility is a paradigm shift in the EEG analytics space. It is a concept based on the conjecture that focal seizures arise from a few fragile nodes, i.e., the *SOZ*, which renders the cortical epileptic network on the brink of instability. When one observes iEEG data during interictal, or preictal periods (Figure 1 left), activity recorded from each channel is noisy and hovers around a baseline value - a "balanced" network. In contrast, when one observes iEEG data during a seizure event (Figure 1 right), activity (i) grows in amplitude, (ii) oscillates, and (iii) spreads in the brain - an "unbalanced" network. From a dynamical systems perspective, the iEEG network has switched from a stable (non-seizure) to an unstable (seizure) network. Biologically, imbalance can occur through any number of mechanisms, such as elevated glutamate [22, 23, 24, 25, 26, 27, 28, 29], genetic disorder impacting synaptic inhibition [30], decreased GABA [31], inclusion of axo-axonic gap junctions [32], loss of inhibitory chandelier cells [33], or axonal sprouting from layer V excitatory pyramidal cells [34]. The similarity between these mechanisms is imbalance in excitatory and inhibitory connections within a neural network, with perturbations leading to recurring seizures. Neural fragility is a way of modeling this phenomena at the iEEG network level.

To demonstrate how fragility is computed from a dynamical model, we consider a 2-node network as shown in Figure 2. In panel a), a stable network is shown where excitation and inhibition are balanced. The network model is provided in the top row and takes a linear form of $\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t)$, where t is a time index (typically one millisecond). When the inhibitory node is stimulated by an impulse, both nodes transiently respond and the EEG returns to baseline (bottom row in Figure 2). In panel b), the inhibitory node's connections are slightly perturbed in a direction that makes the inhibitory node *less* inhibitory (see red changes to its connectivity to the excitatory node). These changes are reflected in the model and diagram. Now, when the inhibitory node is stimulated by an impulse, the EEG responses from each node have a larger transient

response but still return to baseline. Finally, in panel c), the inhibitory node's connections are further perturbed in a direction that makes the inhibitory node *less* inhibitory. Now, when the inhibitory node is stimulated by an impulse, the EEG responses oscillate demonstrating that the network has gone unstable. The fragility of the inhibitory node is thus quantified as square root of 8 which is the norm of the perturbation vector applied to the first column in the network model. Our conjecture is that small changes in connection strengths at *SOZ* nodes cause an imbalance in connectivity between nodes in the network, resulting in susceptible seizures.

To test our conjecture, we estimate a linear time-varying model with a sliding window of A matrices, that characterize a linear dynamical system: $x(t+1) = Ax(t)$ [17, 16]. This model is a generative model representing the linear dynamics between the set of iEEG channels within a small window of time. Of course, this model is always *estimated* from data, and we have shown in [17] that this linear approximation is a valid model of the data in terms of mean-squared error. From this model, we compute the minimum norm perturbation matrix, $\{\Delta\}_{j=1}^N$, over all N channels [16]. The 2-norm of each matrix represent the neural fragility for that channel. Computing the fragility now across sequentially estimated A matrices results in a spatiotemporal neural fragility heatmap. For a full detailed description of neural fragility, see Supplementary Materials and Methods - section Neural fragility of a network.

2.2 Fragility heatmap highlights the clinical *SOZ* in successful patients

To qualitatively assess the usefulness of fragility in localizing electrodes of interest, we first look at specific examples of patients with varying outcomes and epilepsy etiologies (see Figure 3). We analyzed with fragility and demonstrate how it may provide additional information for *SOZ* localization. In Figure 4, we show an example of three different patients with differing surgical treatments, outcomes, Engel class and clinical complexity along with their fragility heatmaps and corresponding raw iEEG data (for full clinical definitions; see Supplementary Excel table).

In Figure 4a, the red electrode labels on the y-axis correspond to the clinical *SOZ* electrodes; the red electrodes are typically a subset of the resected region. It shows the period 10 seconds before and after electrographic seizure onset, indicated by the black dashed line. In Patient_1, the clinical *SOZ* shows a high degree of fragility, even before seizure onset, which is not visible in the raw EEG. The heatmap for Patient_1 also captures the propagation of seizure activity (Supplementary Figure S9). This patient had a successful surgery, and so we can assume the epileptogenic tissue within the clinical resection likely contained the *SOZ* and EPZ regions; it is likely the clinicians correctly localized the EZ. When viewing the raw EEG data in Figure 4b (top), Patient_1 has iEEG signatures that are readily visible around seizure EEG onset (Figure 4b). We see high-frequency and synchronized spiking activity at onset that occurs in electrodes that clinicians annotated as *SOZ*, which correspond to the most fragile electrodes at onset. In addition, the fragility heatmap captures the onset in the **ATT** and **AD** electrodes and early spread of the seizure into the **PD** electrodes. Specifically, **ATT1 (anterior temporal lobe area)** shows high fragility during the entire period before seizure onset (Figure 4). This area was not identified with scalp EEG, or non-invasive neuroimaging. In this patient, the fragility heatmap agreed with clinical visual EEG analysis, identifying the *SOZ*, which was included in the surgery and led to a seizure free patient.

2.3 Fragility heatmap highlights other possibly epileptic regions outside the *SOZ* in failed patients

In Figure 4, Patient_26 and Patient_40 both show distinct regions with high fragility that were *not* in the clinically annotated *SOZ* (or the resected region), and both had recurrent seizures

after surgery. Specifically in Patient_26, the **ABT** (anterior basal temporal lobe), **PBT** (posterior basal temporal lobe) and **RTG29-39** (mesial temporal lobe) electrodes were highly fragile, but not annotated as *SOZ*. Patient_26 had a resection performed in the right anterior temporal lobe region. Clinicians identified the **RAD**, **RHD** and **RTG40/48** electrodes as the *SOZ*. Patient_40 had laser ablation performed on the electrode region associated with **Q2**, which was not fragile. From seizure onset, many electrodes exhibit the EEG signatures that are clinically relevant, such as spiking and fast-wave activity [35]. In Patient_40, the **X'** (posterior-cingulate), **U'** (posterior-insula) and **N'/M'/F'** (superior frontal gyrus) were all fragile compared to the **Q2** electrode, which recorded from a lesion in the right periventricular nodule. In the raw EEG data, one can see synchronized spikes and spike-waves in these electrodes, but the patient had seizures continue after resection. In the corresponding fragility heatmap, the **ABT** and the **RTG29-32** electrodes are highly fragile compared to the clinically annotated *SOZ* region. In the raw EEG shown in Figure 4b, it is not apparent that these electrodes would be part of the *SOZ*. Visual analysis of the EEG was insufficient for Patient_26 and Patient_40, which ultimately led to insufficient localizations and then to failed surgical outcomes. Based on the fragility heatmaps, the fragile regions could be hypothesized to be part of the *SOZ*, and possibly candidates for resection.

2.4 Fragility outperforms all iEEG features in predicting surgical outcomes

To test the validity of neural fragility and the baseline features as *SOZ* markers, we investigate each feature's ability to predict surgical outcomes of patients when stratified by the set of *SOZ* contacts and the rest which we denote as the *SOZ* complement, *SOZ^C*. We train a Structured Random Forest model (RF) [36] for each feature that takes in a partitioned spatiotemporal heatmap based on the clinically annotated *SOZ* and generates a probability of success - a confidence score in the clinical hypothesis. We test each feature's model on a held out data set by applying a varying threshold to the model's output and computing a receiver operating characteristic (ROC) curve. The ROC curve plots true positive rates versus false positive rates and the area under the curve (AUC) is a measure of predictive power of the feature. The larger the AUC, the more predictive the feature is and thus the more valid it is as an iEEG marker of the *SOZ*. We also compute the precision (PR) or positive predictive value (PPV) for each feature, which is the proportion of predicted successful or "positive" results that are true positives (actual successful surgeries). In addition, we compute the negative predictive value (NPV) as well. The larger the PPV/NPV, the more predictive the feature is and thus the more valid it is as an iEEG marker of the *SOZ*.

Since each patient's implantation has a varying number of electrodes, we summarize each feature's distribution of the *SOZ* and *SOZ^C* electrodes into quantile statistics over time (for details, see Supplementary Materials and Methods). As an input to the RF models, we compute the quantile statistics over time of the *SOZ* and *SOZ^C* set for each patient and each feature allowing each patient's heatmaps to have the same dimensions, which are input into a machine learning RF model. RF models are attractive because they are non-parametric, are interpretable (they are a set of decision trees performing a consensus procedure), and are able to handle higher dimensional data better compared to other models, such as Logistic Regression [37]. As an output of the trained RF model, each heatmap gets mapped to a probability that the outcome will be a success. An RF model is tuned for each feature through 10-fold cross-validation (CV; see Supplementary Materials and Methods - Experimental design for full details), resulting in a uniform and rigorous benchmark against neural fragility on the same set of patients. The "high-gamma" frequency band feature encompasses what some would consider HFOs (i.e. 90-300

Figure 4: Fragility heatmaps, and corresponding raw EEG traces of successful and failed outcome patients. (a) From top to bottom, Patient_1 (success, NIH treated, CC1, Engel score 1), Patient_26 (failure, JHH treated, CC3, Engel score 4), and Patient_40 (failure, CClinic treated, CC4, Engel score 3) are shown respectively. The color scale represents the amplitude of the normalized fragility metric, with closer to 1 denoting fragile regions and closer to 0 denoting relatively stable regions. (Left) Overlaid average neural fragility value of each electrode in the window of analysis we used. Black dark squares represent a depth electrode that is not shown easily on the brain. Black lines outline where the clinicians labeled SOZ. Note in Patient_26, RAD and RHD electrodes are denoted by the squares with the color showing the average over the entire electrode. (Right) Clinically annotated heatmaps of the implanted ECoG/SEEG electrodes with red y-axis denoting SOZ contacts. The red contacts are also part of the surgical resection in these patients. Data is shown in the turbo colormap. Best seen if viewed in color. (b) Corresponding raw EEG data for each patient with electrodes on y-axis and time on x-axis with the dashed white-lines denoting seizure onset. Each shows 10 seconds before seizure onset marked by epileptologists, and 10 seconds after. EEG was set at a monopolar reference with line noise filtered out. Not all electrodes are visualized in the brain plot because channels that were deemed noisy, or in white matter were not included in analysis (for more information, see section Supplementary Materials and Methods). In addition, only a select set of channels are chosen for the heatmap and time-series for sake of visualization on a page and to demonstrate select channels that demonstrated different fragility values. Each EEG snapshot is shown at a fixed scale for that specific snapshot that was best for visualization, ranging from 200 uV to 2000 uV.

Hz) activity.

In terms of AUC (measures discrimination), neural fragility performs the best (Figure 5a). Compared to the top 3 leading baseline features, neural fragility outperforms the beta band (15-30 Hz) power by more than 13%. The AUC of fragility is the highest with a value of 0.88 ± 0.064 , compared to the next best representation, the beta frequency band, with a value of 0.82 ± 0.040 . From the ROC (and PR) curves, we observe that the fragility consistently has a higher sensitivity for the same false positive rate compared to the 20 other feature representations (Supplemental Figure S11). In terms of effect size, neural fragility improves over the beta frequency band with a large effect size of 0.97 Cohen's D with 9 out of 10 folds improving the AUC (Paired Wilcoxon Rank-sum PValue = 0.027 (see Supplemental Figure S11). As a result of the 10-fold CV, we can compare the AUC on the same set of patients in a paired effect size and statistical test improving AUC on 9 out of the 10 CV folds.

In terms of PR (measures PPV), neural fragility also performs the best (Figure 5b). In terms of average precision, which weighs the predictive power in terms of the total number of patients, neural fragility also obtains an average precision of 0.83 ± 0.076 , which is $>5\%$ better than the next best feature. In addition, compared to the clinical surgical success rate of 47%, fragility had a $76\% \pm 6\%$ accuracy in predicting surgical outcome, a PPV of 0.903 ± 0.103 and a NPV of 0.872 ± 0.136 . In terms of PR, neural fragility improves over the beta band power with a medium effect size of 0.57 Cohen's D with 8 out of 10 folds improving the PR (PValue = 0.08).

When we compare the differences in the success probabilities predicted by the models stratified by surgical outcomes, we observe that neural fragility has the highest effect size of 1.507 (1.233-1.763 - 95% confidence interval) (Figure 5c) and lowest pvalue of $6.748e-31$ (Figure 5d). In addition to having good discrimination, we compute how well-calibrated the model is, having the success probabilities values reflect true risk strata in the patient population. We quantify how well calibrated the success probability distributions are over the held-out test set. In Supplementary Figure S10, we show that the RF model trained on neural fragility produces well-calibrated success probability values. A perfectly calibrated model would have a 20% confidence in a set of patients with exactly 20% success outcomes. The Brier-loss (a measure of calibration; 0 being perfect) was 0.162 ± 0.036 , which was a 6.7% improvement compared to the next best feature.

2.5 Fragility correlates with clinical complexity and treatment outcomes

Neural fragility correlates with treatment outcomes and clinical complexity. Successful outcomes at 12+ months post-op are defined as seizure free (Engel class I and ILAE scores of 1 and 2) and failure outcomes are defined as seizure recurrence (Engel classes 2-4). In addition, we can categorize patients by their clinical complexity (CC) as follows: (1) lesional, (2) focal temporal, (3) focal extratemporal, and (4) multi-focal (Figure 3) [21]. We stratify the distribution of the success probabilities based on Engel score and show a decreasing trend in success probability as Engel score increases (Figure 6b). The effect sizes when comparing against Engel class I were 1.067 for Engel class II ($P = 4.438e-50$), 1.387 for Engel class III ($P = 2.148e-69$), and 1.800 for Engel class IV ($P = 4.476e-74$). Although the AUC indirectly implies such a difference would exist between Engel score I (success) and Engel score II-IV (failures), it is reassuring to see that the level of confidence decreases as the severity of the failure increases; Engel score IV patients experience no changes in their seizures, suggesting the true epileptic tissue was not localized and resected, while Engel score II patients experience some changes, suggesting there was portions of the true epileptic tissue that may have been resected. We also compare the success probability distributions with respect to the ILAE score, which is another stratification of the surgical outcomes (Figure 6c). There is a similar trend as seen in the Engel score - as ILAE score

Figure 5: Area under the curve and average precision performance. Specific results for neural fragility are marked in red for each of the panels (a-d). **(a)** Discrimination plot (measured with AUC) shows the relative performance of benchmark feature representations compared to that achieved with neural fragility. **(b)** A similar average-PR curve shows the relative positive predictive value of all features compared with fragility. Average precision is the analogous area under the curve for the PR curve. **(c)** A summary of the Cohen's D effect size measurements between the success and failed outcome CS distributions across all features. The effect size of neural fragility is significantly greater than that of the beta band ($\alpha = 0.05$) **(d)** The corresponding PValues of the effect size differences between success and failed outcomes, computed via the Mann-Whitney U-test. Note that the samples in (c) are not shown for visual sake; data was approximately bell-curve shaped and a box plot adequately summarizes the main descriptive statistics of the distribution.

increases, the success probability decreases.

We also analyze the success probability with respect to clinical measures of the epilepsy severity, such as the clinical complexity (CC) of the patient, which is a categorization of the etiology of the disease. CC is determined by what type of seizures the patient exhibits rather than the severity of the seizures after-surgery (for more details see Supplementary Materials and Methods - Dataset collection). CC is a factor that is determined before surgery and one we expect would correlate with failure rate. The higher the CC, the more difficult localization of the *SOZ* is and hence possibly the more seizure recurrences after surgery. In Figure 6a, we observed that CC1 and CC2 (i.e., lesional and temporal lobe) patients had a similar distribution of success probabilities (Cohen's $D = 0.020$; MW U-test PValue = 0.799; LQRT PValue = 0.76), while CC3 and CC4 had significantly lower distributions (CC3: [Cohen's $D = 0.779$; MW U-test PValue = 2.808e-19; LQRT PValue = 0.00]; CC4: [Cohen's $D = 1.138$; MW U-test PValue = 7.388e-26; LQRT PValue = 0.00]). This trend is not optimized for directly in our discrimination task, but it aligns with clinical outcomes. CC1 and CC2 are comparable, which agrees with current data suggesting that lesional and temporal lobe epilepsy have the highest rates of surgical success [38, 39]. Extratemporal (CC4) and multi-focal (CC4) patients tend to have lower success rates due to insufficient localizations and thus the neural fragility confidence in those *SOZ* localizations should be low [38, 40]. Finally, we examine the success probability differences based on gender, handedness, onset age and surgery age to show that there are no relevant differences (Supplemental Figure S13).

Figure 6: Neural fragility of patients stratified by clinical covariates. (a) Distribution of success probability values per patient stratified by clinical complexity (CC; see Supplementary Materials and Methods - Dataset collection), where lesional (1) and temporal lobe epilepsy (2) patients have similar distributions because they are generally the "easier" patients to treat, whereas extratemporal (3) and multi-focal (4) have lower general probabilities because they are "harder" patients to treat. It is important to note that the classification experiment posed did not explicitly optimize this separation between clinical case complexities. (b) The distribution of the probability values per patient stratified by Engel score. Due to the AUC being high for fragility, it is expected that Engel I has high CS, while Engel II-IV have lower success probability. However, the relative downward trend in the success probabilities from Engel II-IV indicated that neural fragility is present in the clinical *SOZ* in varying degrees from Engel II-IV, suggesting that it correlates with the underlying severity of failed outcomes. Engel IV has the lowest average predicted probability of success as expected. (c) A similar distribution for another measure of surgical outcome, the ILAE score, where 1 are considered success and 2-6 are considered failure. Here, ILAE 2-5 follow a decreasing trend with ILAE-6 having the lowest average predicted probability of success.

3 Discussion

We demonstrate that neural fragility, a networked-dynamical systems based representation of iEEG is a strong candidate for a highly interpretable iEEG biomarker of the *SOZ*. We compared neural fragility to 20 other popular features using data from $n=91$ patients treated across five centers with varying clinical complexity. Neural fragility performed the best in terms of AUC, PR and interpretability, bench-marking proposed iEEG features in a uniform fashion over the same sets of patients using the best possible parameters of a RF model to compute a confidence measure (i.e. probability of success) of the clinically annotated *SOZ*. We spent over four years to successfully collect and annotate this heterogeneous dataset. To facilitate reproducibility and further investigation of iEEG features, we made this dataset BIDS-compliant and publicly available through the OpenNeuro website (for details see Supplementary Materials and Methods - Code and data availability).

Challenges in validating iEEG features as SOZ markers Many features have been proposed as potential biomarkers for the *SOZ*, but none have successfully translated into the clinical workflow [11, 12, 41, 38]. Current limitations for evaluating computational approaches to localization are largely attributed to i) insufficient benchmarking to other iEEG features and ii) a lack in sufficient sampling across epilepsy etiologies.

Even with a seemingly successful feature derived from data, it is important to benchmark against existing approaches to provide a holistic view of the value of the said feature. Without benchmarking, it is easy to become overly optimistic in terms of the performance of a feature, whereas it may be that other iEEG features perform just as well. While other traditional features such as the power in the beta band seem to be informative in *SOZ* localization, neural fragility outperforms in effect size, p-value and interpretability.

Although predicting surgical outcomes in our experimental setup is promising, it will be important to understand why certain localizations are successful and why certain are insufficient. For example, in patients with lesions on MRI scans that correlate with the patient’s EEG and seizure semiology, surgical resection can lead to seizure freedom in approximately 70% of patients [21]. Even in these relatively straightforward cases, localization is not perfect, possibly due to chronic effects of epilepsy such as kindling, which can cause neighboring tissue to become abnormal and epileptogenic [42]. There is a need to sample a heterogeneous and large patient population and derive a feature is invariant to epilepsy type and clinical covariates.

Clinical case complexity and surgical outcomes In our dataset, we saw varying outcomes across the five clinical centers. Difference in seizure outcomes across the clinical centers can be explained by several contributory factors affecting the complexity of the epileptic syndromes. The multifactorial contributory factors can be related to i) the percentage of non-lesional versus lesional cases, ii) multiple surgical interventions in the past, iii) patient selection and iv) group experience. For example, JHH is a tertiary referral center with high complexity cases (non-lesional MRIs, multiples surgical interventions in the past, complex semiology and EEG interpretation). The relative center-by-center outcomes are not reflective of the actual center’s surgical outcome rate, but rather samples from the clinical cases that those clinicians saw at the time of our IRB. The pooled 91 patients though do reflect the approximately 50% surgical success rates seen in DRE patients [43, 44, 45]. By including a multitude of centers, we sought to build a diversified sample of varying clinical complexities and practices, thus lending the evaluation of neural fragility more confidence. An important next step would be the prospective evaluation of neural fragility in the context of different clinical case complexities to determine if surgical treatment can be improved with *SOZ* localization assistance.

Limitations of the most popular iEEG features A google scholar search using the keywords: "localization of seizure onset zone epilepsy intracranial EEG" produces over 24,000 results. This is striking considering no computational tools to assist in SOZ localization are in the clinical workflow today. The majority of proposed features lack consistency in finding objective iEEG quantities that correlate to clinically annotated SOZ because they fail to capture internal properties of the iEEG network which are critical to understand when localizing the SOZ. Proposed algorithms either (i) compute EEG features from individual channels (e.g. spectral power in a given frequency band), thus ignoring dependencies between channels [9, 46] to name a few, or they (ii) apply network-based measures to capture pairwise dependencies in the EEG window of interest. Specifically, correlation or coherence between each pair of EEG channels is computed and organized into an adjacency matrix, on which summary statistics are derived including degree distribution and variants of centrality [11, 12, 13, 10, 15]. Such network-based measures are not based on well formulated hypotheses of the role of the epileptic tissue in the iEEG network, and many different networks (adjacency matrices) can have identical summary statistics resulting in ambiguous interpretations of such measures.

A popular EEG feature that has been proposed as an iEEG marker of the SOZ and reported in over 1000 published studies is High-frequency oscillations (HFOs) ([9, 47, 48, 49, 50, 51, 52, 41] to name a few). HFOs are spontaneous events occurring on individual EEG channels that distinctively stand out from the background signal and are divided into three categories: ripples (80–250 Hz), fast ripples (250–500 Hz), and very-fast ripples (>500 Hz) [53]. Retrospective studies suggested that resected brain regions that generate high rates of HFOs may lead to good post-surgical outcome (e.g., [54, 55]). Although they found significant effects for resected areas that either presented a high number of ripples or fast ripples, effect sizes were small and only a few studies fulfilled their selection criteria [55]. Furthermore, several studies have also questioned the reproducibility and reliability of HFOs as a marker [56, 54, 57, 48] and there are also physiologic, non-epileptic HFOs. Their existence poses a challenge, as disentangling them from clinically relevant pathological HFOs still is an unsolved issue [58, 57].

Similar inconclusive results hold in completed prospective studies of HFOs. In 2017, an updated Cochrane review [59] investigated the clinical value of HFOs regarding decision making in epilepsy surgery. They identified only two prospective studies at the time and concluded that there is not enough evidence so far to allow for any reliable conclusions regarding the clinical value of HFOs as a marker for the SOZ. Today, five clinical trials are listed as using HFOs for surgical planning on clinicaltrials.gov as either recruiting, enrolling by invitation, or active and not enrolling and none have reported results. The fundamental limitation of the aforementioned studies lies in the fact that they approach the SOZ EEG marker discovery process as a signal processing and pattern recognition problem, concentrating on processing EEG observations to find events of interest (e.g., HFOs) as opposed to understanding how the observations were generated in the first place and how specific internal network properties can trigger seizures.

Why Neural Fragility Performs Well Rather than analyzing iEEG data at the channel and signal processing level, we seek to model the underlying network dynamics that might give rise to seizures in the form of neural fragility in a dynamical network model. A notion of fragility in networks is commonly seen in analysis of structural [60] networks. Although we are not directly analyzing the structural nature of neuronal network, there are studies that have characterized epilepsy in terms of structural fragility and network organization [61, 18]. Specifically, in cellular studies [34, 33], epilepsy is caused by changes in the structural network (i.e. chandelier cell loss, or abnormal axonal sprouting from layer V pyramidal cells), which causes loss of inhibition or excessive excitation respectively; these biological changes cause downstream aberrant electrical firing (i.e. seizures). By quantifying the fragility of each node in an iEEG network, we determine

how much of a change in that region’s dynamical connections is necessary to cause seizure-like phenomena (e.g. instability). As a result, high neural fragility is hypothesized to coincide with a region that is sensitive to minute perturbations, causing unstable phenomena in the entire network (i.e. a seizure).

Scientific and Technological Advances Emerging from Neural Fragility Neural fragility has the potential to re-define how epilepsy surgery is performed, departing from the classical “localization paradigms” and “en-bloc resections” to a personalized “network-based” user-friendly visualization and surgical strategy. By developing a 3D (brain region, time, fragility) network-based method for anatomical representation of the epileptiform activity, including the seizure onset areas and the early propagation zone, this study will have the potential to offer a safer, more efficient, and cost-effective treatment option for a highly challenging group of patients with disabling DRE. More precise SOZ localization using neural fragility would also guide of chronic implantation of neurostimulation devices aimed to suppress seizures with bursts of current when detected [62, 63].

Neural fragility may also be relevant in detecting epileptogenic regions in the brain when applied to interictal (between seizures) iEEG recordings. Ictal or seizure iEEG data are currently the gold standard in clinical practice for localizing the *SOZ* [21]. However, having patients with electrodes implanted for long periods of time, and requiring the monitoring of multiple seizure events over many weeks carries the risk of infection, sudden death, trauma and cognitive deficits from having repeated seizures. This contributes to the large cost of epilepsy monitoring [2]. If a candidate iEEG marker could be found that is able to provide strong localizing evidence using only interictal data, then it would significantly reduce invasive monitoring time [61].

Neural fragility is an EEG marker that can also further advance our knowledge of neural mechanisms of seizure generation that will then drive more effective interventions. For example, fragility can be used to identify pathological tissue that are then removed and tested in vitro for abnormal histopathological structure [34, 33]. Knowledge of structural abnormalities may inform new targeted drug treatments. In the future, specific fragility patterns can be correlated with specific pathological substrates. Likely, the specific pathological substrates will have different therapeutic approaches. As an example, epilepsy caused by focal cortical dysplasia is treated with focal surgical resection, but post-encephalic epilepsy may have a better therapeutic response with immunosuppressants and steroids.

Finally, neural fragility may have broader implications in understanding how underlying brain network dynamics change during intervention (e.g. drugs or electrical stimulation). Fragility analysis can be applied as a method of assessing the efficacy of specific drug trials to specific pathological groups, which include not only epilepsy but other neurological conditions such as Alzheimer’s disease or the spectrum of dementias. Commonly, the current optimal criteria to recognize therapeutic success in many neurological conditions is purely clinical, but clinical responses are not immediate. This delay in recognizing the appropriate drug and adequate therapeutic doses is highly detrimental. Computational methods as fragility could provide an additional criteria for drug responses, which can be immediate, guiding the treating physician to the correct treatment without delays and unnecessary drug trials.

References

- [1] A. T. Berg. "Identification of Pharmacoresistant Epilepsy." In: *Neurologic Clinics* 27.4 (2009), pp. 1003–1013.
- [2] WHO and W. H. Organization. *Epilepsy*. 2019.
- [3] B. C. Jobst. *Consensus over individualism: Validation of the ILAE definition for drug resistant epilepsy*. 2015.
- [4] C. E. Begley and T. L. Durgin. *The direct cost of epilepsy in the United States: A systematic review of estimates*. 2015.
- [5] L. Jehi. *The epileptogenic zone: Concept and definition*. 2018.
- [6] W. Penfield. "Epileptogenic lesions." In: *Acta neurologica et psychiatrica Belgica* 56.2 (1956), pp. 75–88.
- [7] B. C. Jobst, T. M. Darcey, V. M. Thadani, and D. W. Roberts. "Brain stimulation for the treatment of epilepsy: Brain Stimulation in Epilepsy." In: *Epilepsia* 51 (2010), pp. 88–92.
- [8] H. E. Scharfman. *The neurobiology of epilepsy*. 2007.
- [9] 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." In: *Clinical Neurophysiology* 127.2 (2016), pp. 1057–1066.
- [10] P. Shah, A. Ashourvan, F. Mikhail, A. Pines, L. Kini, R. T. Shinohara, D. S. Bassett, B. Litt, and K. A. Davis. "Local structural connectivity directs seizure spread in focal epilepsy." In: *bioRxiv* (2018), p. 406793.
- [11] A. Li, B. Chennuri, S. Subramanian, R. Yaffe, S. Gliske, W. Stacey, R. Norton, A. Jordan, K. A. Zaghloul, S. K. Inati, S. Agrawal, J. J. Haagensen, J. Hopp, C. Atallah, E. Johnson, N. Crone, W. S. Anderson, Z. Fitzgerald, J. Bulacio, J. T. Gale, S. V. Sarma, and J. Gonzalez-Martinez. "Using network analysis to localize the epileptogenic zone from invasive EEG recordings in intractable focal epilepsy." In: *Network Neuroscience* 02.02 (2018), pp. 218–240.
- [12] S. P. Burns, S. Santaniello, R. B. Yaffe, C. C. Jouny, N. E. Crone, G. K. Bergey, W. S. Anderson, and S. V. Sarma. "Network dynamics of the brain and influence of the epileptic seizure onset zone." In: *Proceedings of the National Academy of Sciences* 111.49 (2014), E5321–E5330.
- [13] R. B. Yaffe, P. Borger, P. Megevan, D. M. Groppe, M. A. Kramer, C. J. Chu, S. Santaniello, C. Meisel, A. D. Mehta, and S. V. Sarma. "Physiology of functional and effective networks in epilepsy." In: *Clinical Neurophysiology* 126.2 (2015), pp. 227–236.
- [14] A. N. Khambhati, K. A. Davis, T. H. Lucas, B. Litt, and D. S. Bassett. "Virtual Cortical Resection Reveals Push-Pull Network Control Preceding Seizure Evolution." In: *Neuron* 91.5 (2016), pp. 1170–1182.
- [15] D. S. Bassett and O. Sporns. "Network neuroscience." In: *Nature neuroscience* 20.3 (2017), pp. 353–364.
- [16] A. Li, S. Inati, K. Zaghloul, and S. Sarma. "Fragility in Epileptic Networks : the Epileptogenic Zone." In: *American Control Conference*. 2017, pp. 1–8.

- [17] A. Li, K. M. Gunnarsdottir, S. Inati, K. Zaghloul, J. Gale, J. Bulacio, J. Martinez-Gonzalez, and S. V. Sarma. “Linear time-varying model characterizes invasive EEG signals generated from complex epileptic networks.” In: *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*. 1. IEEE, 2017, pp. 2802–2805.
- [18] D. Sritharan and S. V. Sarma. “Fragility in Dynamic Networks: Application to Neural Networks in the Epileptic Cortex.” In: *Neural Computation* 26.10 (2014), pp. 2294–2327.
- [19] N. Jette, A. Y. Reid, and S. Wiebe. “What are the limitations of medical therapy for epilepsy? Surgical management of epilepsy Key points.” In: *CMAJ* 186.13 (2014), p. 186.
- [20] J. A. González-Martínez, T. Srikiyvilakul, D. Nair, and W. E. Bingaman. “Long-term seizure outcome in reoperation after failure of epilepsy surgery.” In: *Neurosurgery* 60.5 (2007), pp. 873–879.
- [21] S. Sheikh, D. Nair, R. Gross, and J. Gonzalez Martinez. “Tracking a changing paradigm and the modern face of epilepsy surgery: A comprehensive and critical review on the hunt for the optimal extent of resection in mesial temporal lobe epilepsy.” In: *Epilepsia* 60.9 (2019), pp. 1768–1793.
- [22] J. T. Greenamyre. “The Role of Glutamate in Neurotransmission and in Neurologic Disease.” In: *Archives of Neurology* 43.10 (1986), pp. 1058–1063.
- [23] W. Q. Zhang, P. M. Hudson, T. J. Sobotka, J. S. Hong, and H. A. Tilson. “Extracellular concentrations of amino acid transmitters in ventral hippocampus during and after the development of kindling.” In: *Brain Research* 540.1-2 (1991), pp. 315–318.
- [24] N. A. Janjua, H. Kabuto, and A. Mori. “Increased plasma glutamic acid in a genetic model of epilepsy.” In: *Neurochemical Research* 17.3 (1992), pp. 293–296.
- [25] A. Sherwin, Y. Robitaille, F. Quesney, A. Olivier, J. Villemure, R. Leblanc, W. Feindel, E. Andermann, J. Gotman, F. Andermann, R. Ethier, and S. Kish. “Excitatory amino acids are elevated in human epileptic cerebral cortex.” In: *Neurology* 38.6 (1988), pp. 920–923.
- [26] M. J. During and D. D. Spencer. “Extracellular hippocampal glutamate and spontaneous seizure in the conscious human brain.” In: *The Lancet* 341.8861 (1993), pp. 1607–1610.
- [27] R. Simantov, M. Crispino, W. Hoe, G. Broutman, G. Tocco, J. D. Rothstein, and M. Baudry. “Changes in expression of neuronal and glial glutamate transporters in rat hippocampus following kainate-induced seizure activity.” In: *Molecular Brain Research* 65.1 (1999), pp. 112–123.
- [28] P. B. Crino, H. Jin, M. D. Shumate, M. B. Robinson, D. A. Coulter, and A. R. Brooks-Kayal. “Increased expression of the neuronal glutamate transporter (EAAT3/EAAC1) in hippocampal and neocortical epilepsy.” In: *Epilepsia* 43.3 (2002), pp. 211–218.
- [29] J. A. González-Martínez, Z. Ying, R. Prayson, W. Bingaman, and I. Najm. “Glutamate clearance mechanisms in resected cortical dysplasia: Laboratory investigation.” In: *Journal of Neurosurgery* 114.4 (2011), pp. 1195–1202.
- [30] J. L. Noebels. *The biology of epilepsy genes*. 2003.
- [31] H. F. Bradford. *Glutamate, GABA and epilepsy*. 1995.
- [32] R. D. Traub, M. A. Whittington, E. H. Buhl, F. E. LeBeau, A. Bibbig, S. Boyd, H. Cross, and T. Baldeweg. “A possible role for gap junctions in generation of very fast EEG oscillations preceding the onset of, and perhaps initiating, seizures.” In: *Epilepsia* 42.2 (2001), pp. 153–170.

- [33] J. DeFelipe. “Chandelier cells and epilepsy.” In: *Brain* 122.10 (1999), pp. 1807–1822.
- [34] X. Jin, D. A. Prince, and J. R. Huguenard. “Enhanced Excitatory Synaptic Connectivity in Layer V Pyramidal Neurons of Chronically Injured Epileptogenic Neocortex in Rats.” In: *Journal of Neuroscience* 26.18 (2006), pp. 4891–4900.
- [35] J. C. Bulacio, P. Chauvel, and A. McGonigal. *Stereoelectroencephalography: Interpretation*. 2016.
- [36] R. Perry, T. M. Tomita, R. Mehta, J. Arroyo, J. Patsolic, B. Falk, and J. T. Vogelstein. “Manifold Forests: Closing the Gap on Neural Networks.” In: *Arxiv* (2019).
- [37] R. Couronné, P. Probst, and A. L. Boulesteix. “Random forest versus logistic regression: A large-scale benchmark experiment.” In: *BMC Bioinformatics* 19.1 (2018), p. 270.
- [38] S. Sheikh, N. Thompson, W. Bingaman, J. Gonzalez Martinez, I. Najm, and L. Jehi. “Redefining success in epilepsy surgery: The importance of relative seizure reduction in patient reported quality of life.” In: *Epilepsia* August (2019), pp. 2078–2085.
- [39] V. N. Vakharia, J. S. Duncan, J. A. Witt, C. E. Elger, R. Staba, and J. Engel. “Getting the best outcomes from epilepsy surgery.” In: *Annals of Neurology* (2018).
- [40] S. J. See, L. E. Jehi, S. Vadera, J. Bulacio, I. Najm, and W. Bingaman. “Surgical outcomes in patients with extratemporal epilepsy and subtle or normal magnetic resonance imaging findings.” In: *Neurosurgery* 73.1 (2013), pp. 68–76.
- [41] Y. Varatharajah, B. Berry, J. Cimbalknik, V. Kremen, J. Van Gompel, M. Stead, B. Brinkmann, R. Iyer, and G. Worrell. “Integrating artificial intelligence with real-time intracranial EEG monitoring to automate interictal identification of seizure onset zones in focal epilepsy.” In: *Journal of Neural Engineering* 15.4 (2018).
- [42] K. D. Laxer, E. Trinka, L. J. Hirsch, F. Cendes, J. Langfitt, N. Delanty, T. Resnick, and S. R. Benbadis. *The consequences of refractory epilepsy and its treatment*. 2014.
- [43] J. Engel, M. P. McDermott, S. Wiebe, J. T. Langfitt, J. M. Stern, S. Dewar, M. R. Sperling, I. Gardiner, G. Erba, I. Fried, M. Jacobs, H. V. Vinters, S. Mintzer, and K. Kieburz. “Early surgical therapy for drug-resistant temporal lobe epilepsy: A randomized trial.” In: *JAMA - Journal of the American Medical Association* 307.9 (2012), pp. 922–930.
- [44] A. M. McIntosh, C. A. Averill, R. M. Kalnins, L. A. Mitchell, G. C. Fabinyi, G. D. Jackson, and S. F. Berkovic. “Long-term seizure outcome and risk factors for recurrence after extratemporal epilepsy surgery.” In: *Epilepsia* 53.6 (2012), pp. 970–978.
- [45] R. Yardi, M. E. Morita-Sherman, Z. Fitzgerald, V. Punia, J. Bena, S. Morrison, I. Najm, W. Bingaman, and L. Jehi. “Long-term outcomes of reoperations in epilepsy surgery.” In: *Epilepsia* 61.3 (2020), pp. 465–478.
- [46] J. Jeppesen, S. Beniczky, P. Johansen, P. Sidenius, and A. Fuglsang-Frederiksen. “Detection of epileptic seizures with a modified heart rate variability algorithm based on Lorenz plot.” In: *Seizure* 24.C (2015), pp. 1–7.
- [47] G. Worrell and J. Gotman. *High-frequency oscillations and other electrophysiological biomarkers of epilepsy: Clinical studies*. 2011.
- [48] T. Fedele, S. Burnos, E. Boran, N. Krayenbühl, P. Hilfiker, T. Grunwald, and J. Sarnthein. “Resection of high frequency oscillations predicts seizure outcome in the individual patient.” In: *Scientific Reports* 7.1 (2017), p. 13836.
- [49] S. A. Weiss, Z. Waldman, F. Raimondo, D. Slezak, M. Donmez, G. Worrell, A. Bragin, J. Engel, R. Staba, and M. Sperling. “Localizing epileptogenic regions using high-frequency oscillations and machine learning.” In: *Biomarkers in Medicine* 13.5 (2019), pp. 409–418.

- [50] A. Thomschewski, A. S. Hincapié, and B. Frauscher. *Localization of the epileptogenic zone using high frequency oscillations*. 2019.
- [51] V. Gnatkovsky, M. De Curtis, C. Pastori, F. Cardinale, G. Lo Russo, R. Mai, L. Nobili, I. Sartori, L. Tassi, and S. Francione. “Biomarkers of epileptogenic zone defined by quantified stereo-EEG analysis.” In: *Epilepsia* 55.2 (2014), pp. 296–305.
- [52] S. V. Gliske, Z. T. Irwin, C. Chestek, G. L. Hegeman, B. Brinkmann, O. Sagher, H. J. L. Garton, G. A. Worrell, and W. C. Stacey. “Variability in the location of high frequency oscillations during prolonged intracranial EEG recordings.” In: *Nature communications* 9.1 (2018), p. 2155.
- [53] C. Haegelen, P. Perucca, C. E. Châtillon, L. Andrade-Valença, R. Zelman, J. Jacobs, D. L. Collins, F. Dubeau, A. Olivier, and J. Gotman. “High-frequency oscillations, extent of surgical resection, and surgical outcome in drug-resistant focal epilepsy.” In: *Epilepsia* 54.5 (2013), pp. 848–857.
- [54] C. Cuello-Oderiz, N. von Ellenrieder, R. Sankhe, A. Olivier, J. Hall, F. Dubeau, and J. Gotman. “Value of ictal and interictal epileptiform discharges and high frequency oscillations for delineating the epileptogenic zone in patients with focal cortical dysplasia.” In: *Clinical Neurophysiology* 129.6 (2018), pp. 1311–1319.
- [55] Y. Höller, R. Kutil, L. Klaffenböck, A. Thomschewski, P. M. Höller, A. C. Bathke, J. Jacobs, A. C. Taylor, R. Nardone, and E. Trinka. “High-frequency oscillations in epilepsy and surgical outcome. A meta-analysis.” In: *Frontiers in Human Neuroscience* 9 (2015), p. 574.
- [56] A. M. Spring, D. J. Pittman, Y. Aghakhani, J. Jirsch, N. Pillay, L. E. Bello-Espinosa, C. Josephson, and P. Federico. “Interrater reliability of visually evaluated high frequency oscillations.” In: *Clinical Neurophysiology* 128.3 (2017), pp. 433–441.
- [57] M. A. van ’t Klooster, N. E. van Klink, W. J. Zwiethenning, F. S. Leijten, R. Zelman, C. H. Ferrier, P. C. van Rijen, W. M. Otte, K. P. Braun, G. J. Huiskamp, and M. Zijlmans. “Tailoring epilepsy surgery with fast ripples in the intraoperative electrocorticogram.” In: *Annals of Neurology* 81.5 (2017), pp. 664–676.
- [58] P. Höller, E. Trinka, and Y. Höller. *High-Frequency Oscillations in the Scalp Electroencephalogram: Mission Impossible without Computational Intelligence*. 2018.
- [59] D. Gloss, S. J. Nevitt, and R. Staba. *The role of high-frequency oscillations in epilepsy surgery planning*. 2017.
- [60] W. C. Chang, J. Kudlacek, J. Hlinka, J. Chvojka, M. Hadrava, V. Kumpost, A. D. Powell, R. Janca, M. I. Maturana, P. J. Karoly, D. R. Freestone, M. J. Cook, M. Palus, J. Otahal, J. G. Jefferys, and P. Jiruska. “Loss of neuronal network resilience precedes seizures and determines the ictogenic nature of interictal synaptic perturbations.” In: *Nature Neuroscience* 21.12 (2018), pp. 1742–1752.
- [61] M. A. Kramer and S. S. Cash. *Epilepsy as a disorder of cortical network organization*. 2012.
- [62] K. J. Meador, R. Kapur, D. W. Loring, A. M. Kanner, and M. J. Morrell. “Quality of life and mood in patients with medically intractable epilepsy treated with targeted responsive neurostimulation.” In: *Epilepsy and Behavior* 45 (2015), pp. 242–247.
- [63] D. W. Loring, R. Kapur, K. J. Meador, and M. J. Morrell. “Differential neuropsychological outcomes following targeted responsive neurostimulation for partial-onset epilepsy.” In: *Epilepsia* 56.11 (2015), pp. 1836–1844.
- [64] P. Perucca, F. Dubeau, and J. Gotman. “Intracranial electroencephalographic seizure-onset patterns: Effect of underlying pathology.” In: *Brain* 137.1 (2014), pp. 183–196.

- [65] A. H. Hawasli, S. K. Bandt, R. E. Hogan, N. Werner, and E. C. Leuthardt. “Laser ablation as treatment strategy for medically refractory dominant insular epilepsy: Therapeutic and functional considerations.” In: *Stereotactic and Functional Neurosurgery* 92.6 (2014), pp. 397–404.
- [66] J. Gonzalez-Martinez and D. Lachhwani. “Stereo-electroencephalography in children with cortical dysplasia: technique and results.” In: *Child’s Nervous System* 30.11 (2014), pp. 1853–1857.
- [67] B. Kemp, A. Värri, A. C. Rosa, K. D. Nielsen, and J. Gade. “A simple format for exchange of digitized polygraphic recordings.” In: *Electroencephalography and Clinical Neurophysiology* 82.5 (1992), pp. 391–393.
- [68] C. Holdgraf, S. Appelhoff, S. Bickel, K. Bouchard, S. D’Ambrosio, O. David, O. Devinsky, B. Dichter, A. Flinker, B. L. Foster, K. J. Gorgolewski, I. Groen, D. Groppe, A. Gunduz, L. Hamilton, C. J. Honey, M. Jas, R. Knight, J. P. Lachaux, J. C. Lau, C. Lee-Messer, B. N. Lundstrom, K. J. Miller, J. G. Ojemann, R. Oostenveld, N. Petridou, G. Piantoni, A. Pigorini, N. Pouratian, N. F. Ramsey, A. Stolk, N. C. Swann, F. Tadel, B. Voytek, B. A. Wandell, J. Winawer, K. Whitaker, L. Zehl, and D. Hermes. “iEEG-BIDS, extending the Brain Imaging Data Structure specification to human intracranial electrophysiology.” In: *Scientific data* 6.1 (2019), p. 102.
- [69] K. J. Gorgolewski, T. Auer, V. D. Calhoun, R. C. Craddock, S. Das, E. P. Duff, G. Flandin, S. S. Ghosh, T. Glatard, Y. O. Halchenko, D. A. Handwerker, M. Hanke, D. Keator, X. Li, Z. Michael, C. Maumet, B. N. Nichols, T. E. Nichols, J. Pellman, J. B. Poline, A. Rokem, G. Schaefer, V. Sochat, W. Triplett, J. A. Turner, G. Varoquaux, and R. A. Poldrack. “The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments.” In: *Scientific Data* 3.1 (2016), p. 160044.
- [70] A. Gramfort, M. Luessi, E. Larson, D. A. Engemann, D. Strohmeier, C. Brodbeck, R. Goj, M. Jas, T. Brooks, L. Parkkonen, and M. Hämäläinen. “MEG and EEG data analysis with MNE-Python.” In: *Frontiers in Neuroscience* 7 (2013), p. 267.
- [71] A. Gramfort, M. Luessi, E. Larson, D. A. Engemann, D. Strohmeier, C. Brodbeck, L. Parkkonen, and M. S. Hämäläinen. “MNE software for processing MEG and EEG data.” In: *NeuroImage* 86 (2014), pp. 446–460.
- [72] S. Appelhoff, M. Sanderson, T. L. Brooks, M. van Vliet, R. Quentin, C. Holdgraf, M. Chaumon, E. Mikulan, K. Tavabi, R. Höchenberger, D. Welke, C. Brunner, A. P. Rockhill, E. Larson, S. K. Herbst, R. Luke, A. Li, A. Gramfort, and M. Jas. “MNE-BIDS: MNE-Python + BIDS = easy dataset interaction.” In: *Organization for Human Brain Mapping*. 2020.
- [73] S. Van Der Walt, S. C. Colbert, and G. Varoquaux. “The NumPy array: A structure for efficient numerical computation.” In: *Computing in Science and Engineering* 13.2 (2011), pp. 22–30.
- [74] P. Virtanen et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python.” In: *Nature Methods* 17.3 (2020), pp. 261–272.
- [75] J. D. Hunter. “Matplotlib: A 2D graphics environment.” In: *Computing in Science and Engineering* 9.3 (2007), pp. 99–104.
- [76] M. W. O. B. D. O. P. H. S. L. D. C. G. T. A. Y. H. J. B. C. J. W. J. d. R. C. P. S. H. J. V. S. V. G. K. Qalieh. *mwaskom/seaborn: v0.8.1 (September 2017)*. 2017.

- [77] M. Allen, D. Poggiali, K. Whitaker, T. R. Marshall, and R. A. Kievit. “Raincloud plots: A multi-platform tool for robust data visualization [version 1; peer review: 2 approved].” In: *Wellcome Open Research* 4 (2019), p. 63.
- [78] J. Ho, T. Tunkaya, S. Aryal, H. Choi, and A. Claridge-Chang. *Moving beyond P values: data analysis with estimation graphics*. 2019.
- [79] R. Vallat. “Pingouin: statistics in Python.” In: *Journal of Open Source Software* 3.31 (2018), p. 1026.
- [80] S. Raschka. “MLxtend: Providing machine learning and data science utilities and extensions to Python’s scientific computing stack.” In: *Journal of Open Source Software* 3.24 (2018), p. 638.
- [81] F. Pedregosa, V. Michel, M. Blondel, P. Prettenhofer, R. Weiss, J. Vanderplas, D. Cournapeau, G. Varoquaux, A. Gramfort, B. Thirion, O. Grisel, V. Dubourg, A. Passos, M. Brucher, M. Perrot, and É. Duchesnay. “Scikit-learn: Machine Learning in Python.” In: *Journal of Machine Learning Research* 12.85 (2011), pp. 2825–2830.
- [82] B. Frauscher. “Localizing the epileptogenic zone.” In: *Current Opinion in Neurology* (2020), pp. 1–9.
- [83] J. Gonzalez-Martinez and I. M. Najm. “Indications and selection criteria for invasive monitoring in children with cortical dysplasia.” In: *Child’s Nervous System* 30.11 (2014), pp. 1823–1829.
- [84] L. Ross, A. M. Naduvil, J. C. Bulacio, I. M. Najm, and J. A. Gonzalez-Martinez. “Stereoelectroencephalography-Guided Laser Ablations in Patients With Neocortical Pharmacoresistant Focal Epilepsy: Concept and Operative Technique.” In: *Operative Neurosurgery* (2018).
- [85] J. Chen, Z. Huang, L. Li, L. Ren, and Y. Wang. “Histological type of focal cortical dysplasia is associated with the risk of postsurgical seizure in children and adolescents.” In: *Therapeutics and Clinical Risk Management* 15 (2019), pp. 877–884.
- [86] E. Bertram. *The relevance of kindling for human epilepsy*. 2007.
- [87] L. Breiman. “Random Forests.” In: *Machine Learning* 45 (2001), pp. 5–32.
- [88] T. M. Tomita, J. Browne, C. Shen, J. Chung, J. L. Patsolic, B. Falk, C. E. Priebe, J. Yim, R. Burns, M. Maggioni, and J. T. Vogelstein. “Sparse projection oblique randomer forests.” In: *Journal of Machine Learning Research* 21 (2020).
- [89] M. Skocik, J. Collins, C. Callahan-Flintoft, H. Bowman, and B. Wyble. “I Tried a Bunch of Things: The Dangers of Unexpected Overfitting in Classification.” In: *bioRxiv* (2016), p. 078816.
- [90] A. Alyakin, Y. Qin, and C. E. Priebe. “LqRT: Robust Hypothesis Testing of Location Parameters using Lq-Likelihood-Ratio-Type Test in Python.” In: *Arxiv* (2019).
- [91] Y. Qin and C. E. Priebe. “Robust hypothesis testing via Lq-likelihood.” In: *Statistica Sinica* 27.4 (2017), pp. 1793–1813.

4 Acknowledgements

AL is supported by NIH T32 EB003383, the NSF GRFP (DGE-1746891), the Arcs Scholarship, Whitaker Fellowship and the Chateaubriand Fellowship. SVS is supported by NIH R21 NS103113, the Coulter Foundation, Maryland Innovation Initiative, US NSF Career Award 1055560 and the Burroughs Wellcome Fund CASI Award 1007274. Computational resources were also provided by

the Maryland Advanced Research Computing Center (MARCC). The authors would like to thank Carey Priebe for useful discussions on statistical analyses, Joshua Vogelstein for introducing AL to the concepts of Structured Random Forests, Sarah Kim and Rachel June Smith for very helpful reviews of the manuscript and Macauley Breault for help in creating brain figures. Finally, the authors would like to thank the editor and anonymous reviewers for their insightful comments and reviews.

5 Author Contributions

AL, JM and SVS conceived the project. AL, IC, DB, JJ, AC, AK, JHopp, SC, JHaagensen, EJ, WA, NC, ZF, JB and JM collected and supervised human epilepsy recordings and processes related to IRB at their respective centers. IC, DB, JJ, AC, AK oversaw data from UMH. JHopp, SC, and JHaagensen oversaw data from UMMC. EJ, WA and NC oversaw data from JHH. SI and KZ oversaw data from NIH. ZF, JB and JM oversaw data from CClinic. AL organized the data from all centers, converted data to BIDS, conducted the analyses and generated the figures. AL and SVS wrote the paper with input from the other authors. CH helped develop the RF model and statistical analyses.

6 Competing Interests

The authors declare the following competing interests: AL, SVS and JM have equity in a stealth-mode startup related to epilepsy data analysis.

7 Supplementary Material

Raw data are present in the link in Section 8.2. Any supplementary Information and Source Data files are available in the online version of the paper.

7.1 Dataset and definition elaborations

Clinical Center Patients						
Clinical Covariate	NIH	JHH	UMMC	UMH	CClinic	Total
Number Patients	14	4	7	5	61	91
Number Success	9	1	7	3	24	44
Number Failure	5	3	0	2	37	47
Number Male	6	n/a	7	1	30	44
Number Female	8	n/a	0	4	31	43

Table S1: A table of number of patients per clinical center (NIH = National Institute of Health; JHH = Johns Hopkins Hospital; UMMC = University of Maryland Medical Center; UMH = University of Miami Jackson Memorial Hospital; CClinic = Cleveland Clinic). Includes number of success, failures, and gender per group. Note that JHH did not retain gender information for these group of patients.

Figure S1: After non-invasive analysis is inconclusive, a complex and relatively lengthy iEEG workflow for localization takes place. During invasive monitoring, clinicians identify the clinically hypothesized *SOZ* as the electrodes with the earliest visually identifiable EEG signatures at seizure onset (see [21] for a review). Note it is impossible to observe the EZ directly, so clinicians are required to estimate where to place iEEG electrodes and then successfully interpret the network patterns and raw iEEG patterns (e.g. HFOs, or spiking activity) to formulate the clinical *SOZ* hypothesis. An estimate of the true EZ is then created using a combination of the clinical *SOZ* hypothesis, as well non-invasive findings. When possible, subsequent surgical resection or laser ablation, generally including the *SOZ* along with a variable extent of additional tissue, is performed. Post-operatively, patients are followed for 12+ months and categorized as either success, or failure as defined in Supplementary Materials and Methods - Dataset collection, resulting in a Engel score, or ILAE score. As a result of seizure freedom, one can assume that the clinical *SOZ* included the EZ in some manner. In contrast, if seizures continue, then the clinical *SOZ* was not sufficient to cover the EZ.

Figure S2: Computational experiment setup for all candidate *SOZ* features and statistical analysis. **(a)** Any candidate feature that can produce a spatiotemporal heatmap was computed from EEG data and then partitioned by the clinically annotated *SOZ* set and the complement, *SOZ^C* (i.e. non-*SOZ* electrodes) to compute a confidence statistic measuring the feature's belief of the clinician's hypothesis. Here F_{SOZ} and F_{SOZ^c} were the feature values within their respective sets. $\hat{f}_\theta$ is the function depending on the Random Forest model parameters, θ that maps the statistics of the F_{SOZ} and F_{SOZ^c} to a confidence statistic. An ideal feature would have high and low confidence for success and failed outcomes respectively. Each point on the final CS distribution comparisons represent one patient. **(b)** A more detailed schematic of how our proposed fragility and baseline features were computed from EEG data for a single snapshot of EEG data. See fragility methods section for description of x , A and Δ . For a similar schematic of how the baseline features were computed, see Supplemental figure S5.

Figure S3: Distributions of the 91 patient dataset based on a variety of clinical factors, such as gender (a), handedness (b), clinical complexity (c), and Engel class (d). The plots show distributions over the 91 patients used in analysis.

Figure S4: An overview of the clinical pathways that lead to epilepsy surgery. After a neurosurgical consultation, in few cases, concordant non-invasive data allows a patient to proceed to surgery immediately. In most other cases, invasive monitoring takes place, where localization is attempted based on stereotypical electrographic signatures, described in Supplementary Materials and Methods - Dataset collection and in [64]. If the SOZ was localized confidently, and eloquent regions are not affected, then a resection, or laser ablation is performed. The dashed lines indicate where a decision split was made depending on the clinical analyses of the data.

7.2 Benchmark analysis

Figure S5: An accompanying schematic for Figure S2, describing how we processed baseline features, such as spectral power, and graph metrics. The feature heatmap processing is exactly the same as fragility, allowing us to compare the feature representations of neural fragility, spectral power and graph metrics of correlation and coherence derived graphs for the purposes of SOZ localization.

Figure S6: Pooled fragility distribution analysis for all 91 patients: failed (a), no surgery (b) and successful surgery (c) datasets. Each *SOZ* (soz in blue bars) and *SOZ^C* ('nsOZ' in orange bars) distribution per patient was bootstrap sampled (see Supplementary Materials and Methods for more information on sampling) and then compared using the one-sided Mann-Whitney U test. The corresponding test yielded a statistic of 2776334 (PValue = 0.355) for the failed patient outcomes and a statistic of 36836739 (PValue = 3.326e-70) for the successful patient outcomes. The patients without resection were not included in the analysis comparing to outcome, but these patients can present as interesting case studies where the SOZ was hypothetically localizable, but perhaps was too close to eloquent areas.

7.3 Pooled subject and *SOZ* vs *SOZ^C* analysis

Figure S7: Red SOZ vs black SOZ^C signals for patients presented in Figure 4: Patient_01 (a), Patient_26 (b), Patient_40 (c). For each patient, the ictal snapshots available are visualized around seizure onset with 5 seconds before onset until the first 20% of the seizure. Not necessarily all electrodes in the clinically annotated SOZ are part of the EZ when the patient had a successful outcome. Therefore, if neural fragility had value in contrasting true EZ electrodes from non-EZ electrodes, then any extra electrodes clinically annotated in the SOZ should have relative lower fragility.

Figure S8: Red SOZ vs black SOZ^C fragility signals for pooled patients within each of the five centers with successful (a) and failed outcomes (b) for NIH (n=14), JHH (n=4), CC (n=61), UMH (n=5), and UMMC (n=7) (top to bottom respectively). Note UMMC only had successful outcomes, so there was no curve for the failures.³¹ Seizure periods were resampled and normalized to 100 samples for averaging and viewing purposes. In JHH and UMH, there were only one and two patients in successful outcomes respectively.

7.4 Results and interpretations

Figure S9: Frigility heatmaps with electrodes on y-axis and time on x-axis with the dashed white-lines denoting seizure onset and offset. Shows a period of 30 seconds before seizure onset and 30 seconds after seizure offset. (a) Shows clinically annotated maps of the implanted ECoG/SEEG electrodes with red denoting SOZ contacts. (b) shows spatiotemporal fragility heatmaps for example of successful outcome (Patient_01), and failed outcome (Patient_26 and Patient_40). The color scale represents the amplitude of the normalized fragility metric, with closer to 1 denoting fragile regions and closer to 0 denoting relatively stable regions. The contacts in red and orange are part of the SOZ and RZ, respectively as defined in Methods section. Note that the red contacts are also part of the RZ. Within the seizures, estimating the linear systems are not as stable, which can be seen by fragility "everywhere" in the map. Visualized with Turbo continuous colormap. Best seen if viewed in color.

Figure S10: Calibration curve showing the fraction of actual successful surgical outcomes on the y-axis vs the average CS output on the x-axis. The curve measures how calibrated the predicted success probability values are to the true risk stratification of the patient population. The closer a curve is to the $y = x$ line, then the more calibrated a model is. It is quantified by the Brier-loss (closer to 0 is better), which is shown in the legend, and is significantly lower than the next best feature (an improvement of 15%).

Figure S11: (a) The ROC curve over 10 folds of cross-validation of the held-out test set obtained by applying a Random Forest model onto the spatiotemporal heatmaps to predict surgical outcome (see Supplementary Materials and Methods). Fragility and the top-3 baseline features in terms of AUC are visualized. The shaded area represents the standard deviation of the curve obtained by linear interpolation for visualization purposes. The AUC of fragility obtained a 0.88 ± 0.064 over the 10 standard deviation with a relative improvement of 7.2% improvement in AUC compared to the next best feature representation (i.e. the beta frequency band). At the Youden point (stars), neural fragility obtains a balanced accuracy score of 0.76 ± 0.06 , and an improvement of 0.32 in TPR and 0.32 in FPR compared to the clinical operating point (red star). (c) A paired estimation plot showing how the same test set of patients differed in AUC depending on whether it was using the fragility, or beta feature heatmap representation. The paired Cohen's D effect size was computed at -0.975 (-1.97 to -0.29 ; 95% CI). The p-values associated with the difference between Neural Fragility and the Beta frequency band were 0.0204, 0.0273, and 0.0225 using the Wilcoxon rank-sum test, permutation test, and the paired student t-test respectively. Comparing Fragility to all other feature representations were more different in terms of larger effect sizes and smaller p-values.

Figure S12: Estimated feature importance (mean and stdev) of the associated fragility heatmap used estimated using permutation. The metric of interest was the concordance statistic (i.e. AUC) of the ROC curve. The original feature map is transformed into a 20-dimensional set of time-varying statistics of its *SOZ* and *SOZ^C* electrodes describing the quantiles of the spatiotemporal heatmap (10% - 100% quantiles). This time-varying summary allows these heatmaps to be pooled together across subjects when training a Random Forest classifier as described in Supplementary Materials and Methods.

Figure S13: Fragility success probabilities split by clinical factors, such as handedness (a), gender (b), ethnicity (c) and age at surgery (d). Not all patients had data for each of these categories, so the subset of available data was used. Note the sample sizes vary across different groups shown. For ethnicity, we also had 1 Asian subject, but left it out because the permutation effect size estimation procedure does not work for 1 sample. Effect sizes were estimated using the permutation test and Mann Whitney U test described in Supplementary Materials and Methods. The corresponding effect sizes and p-values were (0.1/0.99) for handedness, and (0.12/0.7) for gender. The slope was negligibly close to 0 for surgery age linear fit.

Figure S14: Fragility versus frequency power in the delta, theta, alpha, beta, gamma and highgamma band for Patient_01, Patient_26, and Patient_40. For band definitions, refer to Supplementary Materials and Methods - Baseline features - spectral features. Every point represents the spectral power and neural fragility value from a randomly chosen window and electrode from one of the patients. No significant correlation is seen or computed from the data. Each spectral feature and fragility are normalized as described in Supplementary Materials and Methods.

8 Supplementary Materials and Methods

8.1 Code and data availability

All code related to generate the figures are at https://github.com/adam2392/fragility_in_ieeg (will be made public once published). We include a jupyter notebook written in Python to help reproduce figures. We also released the raw iEEG data for patients from NIH, UMH, UMMC, and JHH in the OpenNeuro repository in the form of BIDS-iEEG (<https://openneuro.org/datasets/ds003029>). Due to restrictions on data sharing from CClinic, we were unable to release the raw iEEG data that we received from this site.

8.2 Dataset collection

iEEG data from 91 DRE patients who underwent intracranial EEG monitoring, which included either electrocorticography (ECoG), or depth electrodes with stereo-EEG (SEEG) were selected from University of Maryland Medical Center (UMMC), University of Miami Jackson Memorial Hospital (UMH), National Institute of Health (NIH), Johns Hopkins Hospital (JHH), and the Cleveland Clinic (CClinic). Patients exhibiting the following criteria were excluded: no seizures recorded, pregnant, sampling rate less than 250 Hz, previous surgeries more than 6 months before, and no surgery performed (possibly from *SOZ* localizations in eloquent areas). We define successful outcomes as seizure free (Engel class I and ILAE scores of 1 and 2) at 12+ months post-op and failure outcomes with seizure recurrence (Engel classes 2-4) [65, 66, 35]. Of these 91 patients, 44 experienced successful outcomes and 47 had failed outcomes (age at surgery = 31.52 $\pm$ 12.32 years) with a total of 462 seizures (seizure length = 97.82 $\pm$ 91.32 seconds) and 14703 total number of recording electrodes (159.82 $\pm$ 45.42 per subject).

Our clinician team categorized patients by surgical outcome, Engel class and ILAE score. In addition, we categorized patients by their clinical complexity (CC) as follows: (1) lesional, (2) focal temporal, (3) focal extratemporal, and (4) multi-focal (Figure 3) [35, 21]. Each of these were categorized based on previous outcome studies that support this increasing level of localization difficulty. Lesional patients have success rates of $\sim$ 70%, experiencing the highest rate of surgical success because the lesions identified through MRI are likely to be part of the *SOZ* [40]. Localization and surgical success are more challenging patients with non-lesional MRI, with average surgical success rates in temporal, extratemporal and multi-focal epilepsy of $\sim$ 60%, 45% and 30%, respectively [38, 39]. Patients that fit into multiple categories were placed into the more complex category. Next, clinicians identified electrodes that they hypothesized as *SOZ*. In general, this was a subset of the resected region for all patients, unless otherwise noted. The epileptologists define the clinically annotated *SOZ* as the electrodes that participated the earliest in seizures. Every patient's *SOZ* was labeled by 1-3 trained epileptologists (depending on the center). The corresponding *SOZ* complement that we define, or *SOZ*^C, are the electrodes that are not part of the *SOZ*. For more detailed information regarding each patient, see Supplemental Figure S3 and Supplemental clinical data summary Excel file.

At all centers, data were recorded using either a Nihon Kohden (Tokyo, Japan) or Natus (Pleasanton, USA) acquisition system with a typical sampling rate of 1000 Hz (for details regarding sampling rate per patient, see Supplementary file table). Signals were referenced to a common electrode placed subcutaneously on the scalp, on the mastoid process, or on the subdural grid. At all centers, as part of routine clinical care, up to three board-certified epileptologists marked the electrographic onset and the termination of each seizure by consensus. The time of seizure onset was indicated by a variety of stereotypical electrographic features which included, but were not limited to, the onset of fast rhythmic activity, an isolated spike or spike-and-wave complex followed by rhythmic activity, or an electrodecremental response. Each snapshot available for a

patient was clipped at least 30 seconds before and after the ictal event. We discarded electrodes from further analysis if they were deemed excessively noisy by clinicians, recording from white matter, or were not EEG related (e.g. reference, EKG, or not attached to the brain) which resulted in 97.23 ± 34.87 (mean $\pm$ std) electrodes used per patient in our analysis. Data was initially stored in the form of the European Data Format (EDF) files [67]. We preprocessed data into the BIDS-iEEG format and performed processing using Python3.6, Numpy, Scipy, MNE-Python and MNE-BIDS [68, 69, 70, 71, 72, 73, 74]. Figures were generated using Matplotlib and Seaborn [75, 76, 77]. Statistical analyses were performed with mlxtend, pingouin, dabest, scipy and scikit-learn [78, 79, 80, 74, 81].

Decisions regarding the need for invasive monitoring and the placement of electrode arrays were made independently of this work and part of routine clinical care. All data were acquired with approval of local Institutional Review Board (IRB) at each clinical institution. The acquisition of data for research purposes was completed with no impact on the clinical objectives of the patient stay. Digitized data were stored in an IRB-approved database compliant with Health Insurance Portability and Accountability Act (HIPAA) regulations.

8.3 Surgical Workup

Epilepsy surgery is a complicated workflow that depends on many clinical factors, such as the history of the patient, and non-invasive testing [82]. We give a brief overview here to provide the reader with a sense of the complexities inherent in evaluating efficacy of localization algorithms, and even the surgical outcomes (see a schematic in Supplementary Figure S4). Every DRE patient generally begins with what is known as a Phase 1 evaluation, where a plethora of non-invasive procedures are done, such as: MRI to look for structural abnormalities, SPECT and PET to look at blood flow and metabolism, fMRI/MEG to map eloquent brain regions and/or hyperactive regions, and neuropsychological testing to determine if cognitive deficits are related to seizure onsets and provide a picture of likely post-operative cognitive deficits. These are just samples of tests that can be done, and active research is being done to explore their utilities in the context of SOZ localization.

When a patient has undergone all the necessary tests, then Phase 2 evaluation begins, where invasive monitoring is considered. iEEG electrodes are implanted into the patient’s brain, and are monitored for several weeks to observe many seizures. While they are being monitored, patients are taken off anti-epileptic drugs (AEDs), which present serious dangers to the patient because they may have recurring seizures that could lead to sudden death. Once clinicians have sufficiently formed an SOZ hypothesis based on the iEEG data they have collected, and deem that surgery will not affect eloquent regions, then they proceed with surgery. The types of procedures can be broadly classified as a resection, or laser ablation [83, 84].

Assuming that surgery includes the true epileptic tissue, then the patient will be seizure free. On the other hand, there can be many reasons for a failed surgery, such as: neoepileptogenesis [85], kindling [86], a mismatch between the *SOZ* and resected region, or insufficient resection. Due to the complex nature of epilepsy and the brain, it is difficult to determine the cause of failure. We simply require that a useful feature representation of the electrodes separate the *SOZ* from the *SOZ^C* in successful surgeries, and vice-versa in failed surgeries.

8.4 Preprocessing of data

Every dataset was notch filtered at 60 Hz (with a cutoff window of 2 Hz) and bandpass filtered between 0.5 and the Nyquist frequency with a fourth order Butterworth filter. A common average reference was applied to remove any correlated noise. EEG sequences were broken down into

sequential windows and the features were computed within each window (see Supplementary Materials and Methods - Neural fragility of a network, Baseline features - spectral features and Baseline features - graph analysis of networks for details). Each proposed feature representation produces a value for each electrode for each separate window, and results in a full spatiotemporal heatmap when computed over sequential windows with electrodes on the y-axis, time on the x-axis and feature value on the color axis. In total, we computed 20 different baseline feature representations from the iEEG data: 6 frequency power bands, 7 eigenvector centralities (one for each frequency band coherence connectivity matrix and one for a correlation connectivity matrix), 7 degrees (one for each frequency band coherence connectivity matrix and one for a correlation connectivity matrix). Values at each window of time were normalized across electrodes to values that could range from 0 up to at most 1, to allow for comparison of relative feature value differences across electrodes over time; the higher a normalized feature, the more we hypothesized that electrode was part of the SOZ [16].

8.5 Neural fragility of a network

When one observes iEEG data during interictal, or preictal periods (e.g. Supplementary Figure 1, top left), activity recorded from each channel is noisy and hovers around a baseline value. In contrast, when one observes iEEG data during a seizure event, activity (i) grows in amplitude, (ii) oscillates, and (iii) spreads in the brain. From a dynamical systems perspective, the iEEG network has switched from a stable (non-seizure) to an unstable (seizure) network. The only difference between the iEEG networks in Figure 2 is the connection strengths representing the dynamical interactions between a few channels, i.e., the *SOZ*. Our conjecture is that small changes in connection strengths at *SOZ* nodes cause an imbalance in connectivity between inhibitory (I) and excitatory (E) populations (nodes) in the region. Either inhibition is decreased and/or excitation is increased; thus, if the *SOZ* is perturbed then over excitation can occur manifesting in a seizure.

To compute fragility heatmaps from iEEG recordings, we constructed simple linear models as described above and in equation 1, but one for each 250 ms iEEG window. We used least-squares algorithm with a $10e-5$ l2-norm regularization to ensure that the model identified was stable as in [17, 16]. This can be solved using the pseudoinverse. Then, we slid the window 125 ms and repeated the process, generating a sequence of linear network models in time as in Supplemental Figure S2b). We systematically computed the minimum perturbation required for each electrode's connections (Supplemental Figure S2b) to produce instability for the entire network as described in [16]. This is represented in equation 2, where the Δ_i is the desired column perturbation matrix for channel i , and the $\lambda = r$ is the desired radii to perturb one single eigenvalue to. The electrodes that were the most fragile were hypothesized to be related to the SOZ in these epilepsy networks (seen as the dark-red color in the turbo colormap in Figure 4).

$$x(t+1) = Ax(t) \quad (1)$$

$$\Delta_i = \min ||\Delta_i|| \text{ s.t. } \exists \lambda = r \in \sigma(A + \Delta) \quad (2)$$

8.6 Baseline features - spectral features

We constructed spectral-based features from frequency bands of interest by applying a multi-taper Fourier transform over sliding windows of data with a window/step size of 2.5/0.5 seconds [12]. We required relatively longer time windows to accurately estimate some of the lower frequency

bands. Each EEG time series was first transformed into a 3-dimensional array (electrodes X frequency X time), and then averaged within each frequency band to form six different spectral feature representations of the data. We break down frequency bands as follows:

1. Delta Frequency Band [0.5 - 4 Hz]
2. Theta Frequency Band [4 - 8 Hz]
3. Alpha Frequency Band [8 - 13 Hz]
4. Beta Frequency Band [13 - 30 Hz]
5. Gamma Frequency Band [30 - 90 Hz]
6. High-Gamma Frequency Band [90 - 300 Hz]

The High-Gamma frequency band includes frequencies of "ripples" - a category of HFOs.

8.7 Baseline features - graph analysis of networks

We computed a time domain model using Pearson correlation (equation 3) and a frequency domain model using coherence (equation 4). We computed the connectivity matrix using MNE-Python and used the default values [71]. In the equations, (i, j) are the electrode indices, Cov is the covariance, σ is the standard deviation, f is the frequency band, and G is cross-spectral density. These connectivity models attempt to capture linear correlations either in time, or in a specific frequency band, but are not dynamical system representations of the data (i.e. $x(t+1) = Ax(t)$). For each network-based feature, a sliding window/step size of 2.5/0.5 seconds were used, resulting in a sequence of network matrices over time resulting in 3-dimensional arrays (electrodes X electrodes X time) [11, 12, 13].

$$Corr_{ij} = \frac{Cov(x_i, x_j)}{\sigma_{x_i} \sigma_{x_j}} \quad (3)$$

$$Coh_{ij}(f) = \frac{|G_{ij}(f)|^2}{G_{ii}(f)G_{jj}(f)} \quad (4)$$

From each network matrix, we computed the eigenvector centrality [12, 11], and the degree features of the network for each electrode across time. Since coherence and pearson correlation are symmetric matrices, the in-degree is equivalent to the out-degree. Centrality describes how influential a node is within a graph network. The degree is the weighted sum of the connections that connect to a specific node. Both features are potential measures that attempt to capture the influence of a specific electrode within an iEEG network, through the lens of graph theory.

8.8 Experimental design

We tested if the neural fragility representation of iEEG data localized the *SOZ* better compared to other proposed features, compared to clinicians and compared to chance. Electrodes with extreme feature values deviating from the average were hypothesized as part of the *SOZ*.

To compute a probability of successful surgical outcome for each patient, we trained a non-parametric machine learning classifier, a Random Forest classifier to output a probability value that the surgery was a success. We input the distribution of feature values in the clinically *SOZ* and the *SOZ*^C. We considered a set of hyperparameters that we then performed evaluation over a 10-fold nested cross-validation scheme. We then performed statistical analysis on the final classification performance to determine the most robust feature representation.

Pooled Patient Analysis First, we analyzed the difference in the distributions of neural fragility between *SOZ* and the *SOZ^C*. We pooled all patients together, stratified by surgical outcome and compared the neural fragility distributions using a one-sided Mann-Whitney U test (Success pvalue = 3.326e-70, Fail pvalue = 0.355; Supplementary Figure S6). This suggested that there is some effect on average where fragility is higher in the *SOZ* for successful outcomes, so we next looked at the distributions per patient’s seizure snapshot around seizure onset. In Supplementary Figure S7, success outcome patients have a higher neural fragility in the *SOZ*. This effect is seen when pooling patients across all centers as well, where neural fragility is either i) higher before the seizure onset, or ii) has a marked difference starting at seizure onset (Supplementary Figure S8). Next, we performed a classification experiment (Figure 3) that would determine the robustness of the neural fragility representation at the patient level benchmarked against 20 other features.

Non-parametric Decision-Tree Classifier To determine the value of a feature representation of the iEEG data, we posed a binary classification problem where the goal would be to determine the surgical outcome (success or failure) for a particular patient’s spatiotemporal heatmap. Each spatiotemporal heatmap was split into its *SOZ* and *SOZ^C* set (F_{SOZ} and F_{SOZ^C} in Figure S2). Then each set of electrodes’ feature distributions were summarized with its quantile statistics over time, resulting in twenty signals: ten quantiles from 10-100 of the *SOZ* and ten quantiles of the *SOZ^C* over time. We used a Random Forest (RF) classifier [87]. Specifically, it was a variant known as the Structured Random Forest, or Manifold Random Forests [36, 88]. The manifold RF approach allows one to encode structural assumptions on the dataset, such as correlations in time of the fed in data matrix. The input data matrix for each RF is a multivariate time series that summarizes the quantile statistics of the *SOZ* and *SOZ^C* over time. As a result, we obtained better results using the manifold RF because it took advantage of local correlations in time, rather than treating all values in the data matrix input as independent as done in a traditional RF model. This approach allowed our classifiers to learn faster with less data, compared to treating all inputs as independent in the traditional RF model. For more information on how manifold RF improves on traditional RF, we refer readers to [36, 88]. For every model, we used the default parameters from scikit-learn and the rerf package [81, 88]: 500 estimators, max depth is None, minimum samples split is 1, max features is auto, feature combinations is 1.5, image height of 20 (20 quantiles total), patch height max of 4, patch height minimum of 1, patch width max of 8, patch width minimum of 1. The output of the decision tree classifier is:

$$\hat{f}_{\theta}(F_{SOZ}, F_{SOZ^C}) = P(success)$$

where θ are the trained RF parameters and F_{SOZ} are the heatmap values for the *SOZ* and F_{SOZ^C} are the heatmap values for the *SOZ^C* for a specific feature representation F (shown in Supplemental Figure S2). $\hat{f}$ is the function we are trying to estimate, which predicts the probability of successful surgical outcome for a given feature heatmap.

Hyperparameters When looking at iEEG data, clinicians inherently select windows of interest due to external and prior information about when the seizure clinically manifests. Similarly, we consider a window of 10 seconds before seizure onset to the first 5% of the seizure event. This window was chosen apriori to analysis, and we repeated the analyses with slightly varying windows, but the results were consistent. To provide further contrast to the spatiotemporal heatmaps, we considered also a threshold between 0.3 and 0.7 (spaced by 0.1) that would be applied to the heatmap such that values below were set to 0. This is analogous to clinicians being able to look at a iEEG data and hone in only on areas that are extreme relative to the rest. We

selected a fixed threshold through nested cross-validation, where thresholds are selected on the train/validation set, and then performance is measured on the held out test set.

Structured Heatmap Input To train a RF classifier on patients with varying number of channels, we structured the spatiotemporal heatmap inputs to summarize the distributions of the feature values in the SOZ and SOZ^C sets. We converted these into a vector of 20 quantile values (10 quantiles for SOZ and 10 for SOZ^C), taken evenly from the 10th to the 100th quantiles. This forms a data matrix per heatmap, which summarizes the SOZ and SOZ^C distributions over time, fixed around the seizure onset time.

Nested Cross-Validation Feature Evaluation It is common practice when building a supervised machine learning model to incorporate cross-validation (CV) where the dataset is split into a training and testing set. However, when one has hyperparameters (in addition to the machine learning model parameters) it is crucial to control against over-fitting to the test dataset [89]. Due to our hyperparameter selection of the optimal heatmap thresholds, we used a nested CV scheme, where we split the dataset into a training, validation and a "lock-box" dataset, where the hyperparameters were tuned on an inner CV (70% of the dataset) with the training and validation data, and then performance was evaluated on the lock-box dataset. All features were optimized separately, so that their hyperparameters were optimal for that feature. We computed the discrimination statistic known as the area under the curve (AUC). We repeated the nested-cross validation 10 times, resulting a 10-fold CV. In addition, we performed patient-level CV, ensuring that no patient was in multiple splits of the dataset (i.e. any one patient is only in the train, test, or lock-box dataset). We performed sampling with respect to clinical complexity to ensure that there was an approximately some lesional, temporal, extratemporal and multi-focal patients in the training dataset over each CV fold. We found that this performed slightly better for all features compared to random sampling.

Statistical analysis - Hypothesis Testing Success probability values were computed from the RF classifiers trained on the spatiotemporal feature representations of the iEEG data (fragility, spectral features, and graph metrics from correlation and coherence graphs) resulting in a distribution of probabilities for each feature.

To compare the RF model performance across the 20 proposed features of fragility, spectral power and graph metrics, we computed Cohen's D effect size differences between the groups of interest, and then Mann-Whitney U tests for unpaired data, and paired Wilcoxon rank-sum tests for paired data [78]. We corrected for any multiple comparisons using the Holm-Bonferroni step-down method. In some cases, we also present the Likelihood Q-ratio test (LQRT) results to show that the results are robust to statistical test chosen. The LQRT has been shown to be more robust compared to both the Mann-Whitney U-test and t-test in the presence of noise [90, 91]. LQRT utilizes a bootstrapping procedure, so it's resolution is 0.001 (i.e. if it produces a p-value of 0, then that means it is < 0.001).

We compared the success probabilities stratified by different clinical covariates: surgical outcome, Engel class, clinical complexity, handedness, gender, onset age, and surgery age. We then estimated the effect size differences between distributions in the form of a Cohen's d statistic. Cohen's d was estimated using a non-parametric bootstrapping on the observed data with 5000 re-samples used to construct a 95% confidence interval [78]. The null hypothesis of our experimental setup was that the success probabilities came from the same population. The alternative hypothesis was that the populations were different (i.e. a feature could distinguish success from failed outcomes, or between different Engel classes).

8.9 Feature evaluation using predicted probability of successful surgery

The fragility and all baseline features proposed generated a spatiotemporal heatmap using EEG snapshots of ictal data (outlined in Figure S2). To compare spatiotemporal heatmaps across features, we computed a probability of success (i.e. a confidence score) that is hypothesized to be high for success outcomes and low for failures for "good" features. We expected that fragility would follow a trend of decreasing confidence as CC, Engel score and ILAE score increase. For each clinical covariate group, we measured the effect size difference via bootstrapped sampling, and the statistical p-value between the distributions. We **hypothesized that**: i) fragility would have an effect size difference significantly different from zero when comparing success vs failed outcomes, ii) in addition, this effect size would correlate with meaningful clinical covariates, such as CC and Engel class and iii) both the effect size and p-value would be better than the proposed baseline features.

The higher the probability of success (closer to 1), the more likely the feature indicated a successful surgery, and the lower it was (closer to 0), the more likely the feature indicated a failed surgery. To compute this value, we first partition the heatmap into a *SOZ* and *SOZ^C* as seen in Supplemental Figure S2. This forms the two sets of signals that represent the spatiotemporal feature values of the *SOZ* set vs the *SOZ^C* set of electrodes. Then we take windows of interest, where clinicians find most valuable for *SOZ* localization: the period before seizure and right after seizure onset, and performing a nested CV of RF models. The efficacy of each proposed feature is evaluated based on how well the trained RF model is able to predict the surgical outcomes. We tested our hypotheses stated above by computing a probability of success from each feature heatmap for each patient, and estimated the distribution differences of the CS between various clinical covariates.

8.10 Spatiotemporal feature heatmap interpretability

To determine how valid the output probability values are, we first computed a calibration curve, which told us how well-calibrated the probability values were. We furthermore compared the calibration curves across neural fragility, spectral features and graph metrics. With relatively good discrimination measured by the AUC and good calibration, we were able to then compare the CS across clinical complexity (CC) scores, Engel scores and ILAE scores. Since epilepsy and reasons for failed surgeries are so complex, these are clinical methods for stratifying patient groups based on observed etiology. We analyzed how the success probability differs across each of these categories. In addition, we looked at how the success probabilities might differ across other clinical variables, such as sex (M vs F), handedness (R vs L), epilepsy onset age (years), and age during surgery (years).

Qualitatively reading off a spatiotemporal heatmap is highly interpretable as one can match raw iEEG segments to certain time periods in the heatmap. For every patient, the time-varying quantile signals (from 10th to the 100th) of the *SOZ* and *SOZ^C* were computed for every iEEG heatmap, as specified in Supplementary Materials and Methods - Experimental design. We used permutations feature importance sampling to obtain a relative weighting of each signal over time and how important it was in allowing the RF to correctly determine surgical outcome. This was visualized as a heatmap showing the mean and std importance scores of the *SOZ* and *SOZ^C* statistics, as shown in Figure S12.

In addition, we claim that human interpretability of the spatiotemporal heatmap relies on contrast between the *SOZ* and *SOZ^C* regions in successful outcomes. Using the results of the feature importance permutation test, we computed for every heatmap, an interpretability ratio, which was defined as:

$$I = \frac{F_{SOZ}(90th)}{F_{SOZ^C}(90th)}$$

where I is the interpretability ratio for a specific subject's feature heatmap F . $F_{SOZ}(90th)$ is the 90th quantile and up feature values of the SOZ , and $F_{SOZ^C}(90th)$ is for the SOZ^C . This ratio is then stratified according to surgical outcome, which we expect higher ratios for successful outcomes and lower ratios for failed outcomes. We quantified this difference using Cohen's D effect size and Mann Whitney U-test Pvalues (see Supplementary Materials and Methods - Statistical analysis - Hypothesis Testing).

8.11 Fragility heatmaps are more interpretable than all other EEG feature maps

We can parse the RF model to determine important aspects of the spatiotemporal heatmap that are important in predicting surgical outcome. To do so, we perform permutations on the heatmap inputs to the RF models to measure their relative importances over space and time. We observed that the AUC metric was affected primarily by a combination of the highest (90th quantiles and up) neural fragility values of both the SOZ and SOZ^C contacts (Figure S12). The SOZ^C neural fragility as much as 10 seconds before seizure onset impacted the AUC as did neural fragility of the SOZ right at seizure onset. For both the SOZ and the SOZ^C fragility distributions, the 80th quantile and below did not contribute to the predictive power. In different patients across different clinical centers, the practice of marking seizure onset times can vary due to varying methodologies [21]. As a result, in the SOZ^C , there was some variation in terms of which time point in the heatmap mattered most.

Although strong predictive capabilities of an EEG marker of the SOZ are promising, it is also important that the marker be presented in an interpretable manner to clinicians. We next show how fragility heatmaps are the most interpretable over all baseline features. In Figure S15a, there are two heatmaps computed for the same seizure event in Patient_01: one is a beta band map (left) and one is a neural fragility map (right). Both maps are normalized across channels and both are computed with similar sliding windows. However, less contacts "stand-out" as pathological in the beta-band heatmap before seizure onset, with the majority of the map being different shades of blue. In contrast, in the fragility heatmap, one contact (**ATT1** from Figure 4) is fragile the entire duration before seizure onset (solid white line), and then a few more contacts become fragile after the electrographic onset of the seizure. These fragile areas that "pop-out" in the heat-map as red-hot areas occur in the clinically annotated SOZ and this patient had a successful surgery.

To quantify interpretability, we compute an interpretability ratio: the ratio of the feature values in the 90th quantile between the SOZ and the SOZ^C over the section of data used by the feature's RF model. This measures the contrast that one sees between the extreme values in the SOZ versus the extreme values they see in the SOZ^C . The larger the ratio, then the more contrast the map will have. In Figure S15b, we show that the effect size difference between successful and failed outcome of this interpretability ratio is largest in neural fragility when compared to all baseline features. It is well-known that RF models are scale-invariant [87], so it is plausible that there are portions of heatmaps that distinguish one channel from another that can be parsed out via a decision tree and not the naked eye, which leads to high AUC in the beta band representation in Figure 5. For example a decision tree can discriminate between 0.3000 and 0.2995, which on a normalized color-scale is difficult to parse by visual inspection. Neural fragility on the other hand, shows marked differences between the clinically annotated SOZ (red electrodes on the y-axis) and the actual fragility values even before seizure onset.

Figure S15: Interpretability ratio of feature heatmaps. (a) Two heatmap examples of a seizure snapshot of Patient_01 (NIH treated, ECoG, CC1, Engel I, ILAE 1) with the beta frequency band (left) and the neural fragility heatmap (right). Both colormaps show the relative feature value normalized across channels over time. The black line denotes electrographic seizure onset. (b) A bar plot of the interpretability ratio that is defined in Results Section 8.11 computed for every feature. The y-axis shows an effect size difference between the interpretability ratios of success and failed outcomes. The interpretability ratio for each patient's heatmap is defined as the ratio between the feature values in the two electrode sets ($\frac{SOZ}{SOZ^c}$). Neural fragility is significantly greater than the beta band (alpha level=0.05). The black bars denote 1 std.