2024-06-21

Exploring Seizure Signal Processing and Methods to Characterize Seizure-like Activity in Mouse Brains

Benjamin Kearsley

Honors Thesis

EXPLORING SEIZURE SIGNAL PROCESSING AND METHODS TO CHARACTERIZE SEIZURE-LIKE ACTIVITY IN MOUSE BRAINS

By

Benjamin Kearsley

Submitted to Brigham Young University in partial fulfillment of graduation requirements for University Honors

Department of Statistics

Brigham Young University

August 2024

Advisor: Micah R. Shepherd

Honors Coordinator: Del Scott

ABSTRACT

EXPLORING SEIZURE SIGNAL PROCESSING AND METHODS TO CHARACTERIZE SEIZURE-LIKE ACTIVITY IN MOUSE BRAINS

Benjamin Kearsley

Department of Statistics

Bachelor of Statistics

Status Epilepticus (SE) is a dangerous type of seizure that is difficult to treat and can have permanent or fatal consequences. Developing methods to properly process and classify SE in Local Field Potential (LFP) signals are important steps towards being able to predict SE in clinic and save lives. This thesis explores methods for the seizure data processing LFP data with the goal of gaining a deeper understanding of the effect of brain region and tissue preparation paradigm on power output in specific frequency ranges. The brain regions compared are the neocortex and hippocampus, and the preparation paradigms are 4AP and 0 Mg^{2+}. This thesis also explores the use of statistical features in tree-based models for classifying SE-like behavior. A random-forest model was fit and tested on both intra-trace classification and inter-trace classification, with 99.58% and 64.97% accuracy, respectively, suggesting that other models may be better suited for this classification task.

ACKNOWLEDGMENTS

I am extremely grateful for all the assistance I received in conducting this research and support throughout this thesis. First and foremost, I would like to thank Dr. Micah R. Shepherd and Dr. R. Ryley Parrish for involving me in their research and taking a chance on a new research assistant with little prior experience in neuroscience. I am grateful for their encouraging me to ask new questions and experiment with new avenues of research on this project that were interesting to me. I would also like to thank my colleague Jacob Norby for his assistance in bringing me up to speed on prior research on this subject and his advice and counsel. Lastly, I would like to thank my parents, Chad and Pam Kearsley, and my wife, Mitzi Kearsley, for their support during my undergraduate tenure.

TABLE OF CONTENTS

CHAPTER 1 - INTRODUCTION

A seizure is an abrupt, unregulated surge of electrical activity in the brain. In human subjects, seizures pose serious risks, potentially causing brain impairment or fatality. Of particular concern is Status Epilepticus (SE), a seizure type characterized by prolonged seizure activity. Historically SE was defined clinically as a seizure persisting for over thirty minutes (Drislane). More recently, neuroscientists have refined this definition to encompass seizures exhibiting five or more minutes of continuous clinical or electrographic seizure activity, or recurrent seizure activity without recovery between seizures (Brophy, et. al.). SE may manifest as convulsive or non-convulsive events, with potential for generalization and propagation across brain cells, leading to more serious and permanent damage. Standard treatment of SE involves administration of benzodiazepines, a common medication used to control seizures. Nonetheless, during prolonged seizures, like SE events, many patients lose their sensitivity to these medications and the seizures become pharmacoresistant, rendering benzodiazepines an ineffective treatment (Goodkin, et. al.). Thus, it is critically important to be able to diagnose and even predict SE in patients in order to reduce the most serious neurological damage and mitigate mortality rates.

Various processes can precipitate SE, including both acute and chronic factors. Acute precipitants may include head trauma, hypoxia, drug withdrawal symptoms, and infections affecting the central nervous system. Additionally, chronic conditions such as pre-existing epilepsy, other seizure disorders, or prior brain injury and stroke can also be predisposing factors for future seizure occurrence (Wylie, et. al.). While these phenotypic precursors offer valuable insight into seizure risk, they do not readily translate into predictive features. Outside the realm of just SE, alternative methods have been investigated as indicators of seizure onset. Certain individuals undergoing seizure episodes report experiencing subjective symptoms referred to as auras, manifesting before and during the onset of the seizure. Auras commonly manifest as hallucinations, sensory disturbances such as numbness or paralysis, among other presentations (Leticia, et. al.). However, little research has examined the utility of auras as predictive features of SE. While auras may be relevant at an individual level, they lack the necessary standardization for widespread clinical application, and cannot be relied on for seizure prediction. Alternative avenues for seizure prediction have been explored, including training service animals to anticipate seizures in their owners and adjusting lifestyle factors such as sleep patterns, dietary habits, and physical activity (Dawit, et. al). Again, while these methods may be beneficial in select cases, they do not represent a universally applicable strategy for seizure management, and the extent of their efficacy in the context of SE remains insufficiently investigated.

Advancements in mathematics and data science have led to the development of sophisticated methods for managing vast datasets, including stochastic signals, thereby rendering machine learning and signal processing techniques more accessible and valuable for analyzing diverse forms of data, including biological data pertinent to seizures. One common form of seizure-related biological data are electroencephalogram (EEG) recordings, which capture electrical activity within brain tissue and serve as a primary modality used by neuroscientists to investigate SE and other seizures. A seminal study in 2011 demonstrated the efficacy of Support Vector Machines (SVMs) applied to EEG data, achieving a 97.5% sensitivity in seizure prediction (Park, et. al.). More recently other deep learning architectures, such as bidirectional Long Short-Term Memory (LSTM) networks, have been used to predict seizures with 94.6% sensitivity (Singh, et. al.). While a substantial body of literature exists describing the various machine learning and signal processing methods employed for seizure classification and prediction, comparatively limited research has been devoted specifically to classifying SE. A prevailing theme across this literature, irrespective of method, is the recognition that seizures follow non-random patterns that give rise to patient-specific probability distributions (Amengual, et. al.). However, SE poses unique challenges for classification owing to its distinct clinical and EEG manifestation compared to other seizures, which are self-terminating. Kramer, et al., observed that while self-terminating seizures converge toward a common dynamical mechanism upon termination, SE approaches this point but does not traverse it. Furthermore, the non-random nature of SE remains poorly understood, with little literature addressing the diurnal and nocturnal patterns of SE. Despite the inherent complexities arising from the stochastic nature of SE signals as portrayed by EEG recordings, data-driven approaches using signal processing and machine learning hold promise for developing robust classification and prediction models for SE-like events.

This thesis endeavors to explore the utilization of signal processing and machine learning methods for the visualization and comprehension of EEG-like data, termed Local Field Potential (LFP) data, related to SE-like events. The structure of this thesis is organized as follows. Chapter 2 delves into the data acquisition process and the data's structure. We will review prior research done on seizure data processing and feature engineering. We will discuss the significant spectral and statistical features that can be derived from raw LFP data. Comparative analysis of these features across baseline brain activity, seizure activity, and SE-like activity will be undertaken. Also, spectral analysis of seizure-like activity will be compared between the hippocampus and neocortex. The research on seizure data processing in Chapter 2 provides a strong foundation for the exploration of methodologies for SE classification in the next chapter.

Chapter 3 presents findings from supervised machine learning paradigms, followed by discussions on unsupervised clustering techniques and changepoint detection methods. A comprehensive evaluation of the merits and limitations of each approach is provided, along with suggestions for potential avenues of advancement in subsequent research endeavors. Chapter 4 synthesizes the outcome of our investigations and concludes this thesis.

CHAPTER 2 – SEIZURE DATA PROCESSING [1]

Integral to this chapter are two main questions:

1. How do different preparation paradigms affect the spectral character of seizure-like activity in LFP recordings from mice brain slices in the hippocampus and neocortex?
2. How do statistical features differ between baseline, Non-SE seizure-like activity, and SE-like seizure activity?

To answer these primary questions, we will begin by discussing prior research conducted on seizure data analysis and signal processing. We will also review the use of some statistical feature engineering. Next, we will review the data acquisition process, and briefly describe the different preparation paradigms and seizure activity classes. We will discuss the methods used to answer the two main questions of this chapter. We will present specific examples of differences found between preparation paradigms and differences between baseline and seizure-like activity. We will also examine differences found in aggregate across a larger sample. Lastly, we will discuss our findings and their implications, in particular how they pertain to the research presented in Chapter 3.

Prior Research on Seizure Data Processing

Signal processing stands as a pivotal domain in contemporary society, integral to the functionality of ubiquitous technologies such as cellphones, computers, and radios. Rooted in electrical engineering, signal processing encompasses the modeling and analysis of data

[1] The content in this chapter is a selected review of research done as part of the following paper which is currently in the process of being published. Stubbs, I. W., Blotter, M. L., Jacob H. Norby, Holmes, M, Kearsley, B, et. al. (n.d.). High Quality Seizure-Like Activity from Acute Brain Slices Using a Complementary Metal-Oxide-Semiconductor High-Density Microelectrode Array Systems.

representations of physical phenomena (IEEE Signal Processing Society). An illustrative example of using signal processing techniques would be the work done by Prasad and Gaitonde in 2022, in which they presented a model for jet noise, leveraging sophisticated signal processing methodologies. Despite disparate manifestations, the chaotic nature of both jet noise and LFP seizure data underscores the utility of decomposing these signals into constituent components for comprehensive analysis.

The Fourier transform, a cornerstone of signal processing, serves as a mathematical tool to dissect signals in the time domain into their constituent frequencies, revealing the power of the signal in frequencies up to one half of the signal's sampling rate (also known as the Nyquist Frequency.) Its widespread utilization spans various domains of science and engineering, facilitating the analysis and manipulation of signals in the frequency domain. Core implementations of the Fourier transform are pervasive across programming languages like Python and MATLAB, with the Fast Fourier transform emerging as a computationally efficient alternative for computing Discrete Fourier Transforms across successive windows in a time trace. In finance, the Fourier transform plays a crucial role in the Black-Scholes-Merton formula for pricing volatile options (Schmelzle). Again, similar to jet noise and seizures, option pricing represents another stochastic signal that can be modeled using signal processing techniques.

LFP signals represent ideal candidates for the application of Fourier analysis as a means for generating descriptive statistics for signals, as well as creating features for predictive models. In conventional neuroscience, there are typically seven defined frequency bands as shown in Table 1:

Table 1 – Conventional Frequency Bands for Seizure Analysis

Band	**Frequency**
Delta	1-4 Hz
Theta	4-8 Hz
Alpha	8-12 Hz
Beta	12-30 Hz
Low Gamma	30-60 Hz
High Gamma	60-150 Hz
High Frequency Oscillations	150Hz

Observing the power distribution in each of these frequency bands during different events during an LFP recording or across LFP recordings is a common and well-defined method for describing significant differences between signals (Wang and Mengoni). In our research, we will use them as the primary features to differentiate between paradigms, regions, and events.

Other research has been done on the capacity to use these descriptors as features in a predictive model. Mehla and Singal, et. al., used the Fourier transform to decompose and engineer features for a Support Vector Machine (SVM) classification model, that correctly identified baseline and seizure activities with an accuracy of over 99%. Using a multi-step algorithm with a discrete wavelet transform to reduce the noise and a Fourier transform to identify features, Gao, et. al., fit a Pattern Recognition Network with 92% accuracy at classifying seizure and baseline activity. In both of these studies, the Fourier transform was effectively applied to EEG recordings, similar to our LFP recordings, and used to differentiate between seizure and baseline activity. It is therefore logical to think that it might similarly be used to further differentiate between seizure and SE-like activity. In summary, Fourier analysis has shown great promise for differentiating seizure and baseline activity, but little work has been done to apply Fourier analysis techniques to the differentiation of SE-like vs. non-SE-like activity.

Besides spectral analysis, statistical analysis is also a useful tool to analyze and find differences in stochastic signals. In particular, the four statistical moments (mean, variance, skewness, kurtosis) are simple metrics that can be calculated over sliding windows across a time series to identify discrepancies between sections. Alam, et. al., used empirical mode decomposition and the higher order statistical moments to inform a model for detecting seizures in EEG recordings. This model proved faster than spectral models while retaining the same accuracy, suggesting that many of the differentiating factors of seizures can be captured simply with the statistics. In another paper, the extreme events theory, which has roots in statistics and probability, was used to study the underlying mechanisms responsible for seizure in mice (Frolov). Results from this analysis "evidenced a possibility for early (up to 7 s) prediction of epileptic seizures" (Frolov, et. al.). Other research suggests that basic statistical moments can be used to create a feature set capable of characterizing basic EEG patterns, enabling seizure classification (Haderlein, et. al., Roy, et. al.). These statistical methods may also be promising for classification, but have not been thoroughly applied to the seizure vs. SE case or baseline vs. SE case.

Data Collection

Data for this exploration of seizure data processing was acquired in a neuroscience lab at Brigham Young University. For the first question introduced in this chapter, data was acquired from the Parrish Lab. In the lab researchers prepared acute brain slices from several male and female C57BL-6 mice using two different preparation paradigms, namely 0 Mg^{2+} and 4AP.[2] These brain slices were then placed on a 3Brain C-MOS Accura chip with 4096 electrodes,

[2] Exact details regarding mice breeding and brain slice preparation can be found in the paper cited in footnote 1. In simple terms, preparation paradigm, refers to the type of convulsive media used on the sample.

spaced 60 micrometers apart, and perfused with pro-convulsant media such as will simulate seizures and SE-like events. Seventeen three-hour LFP recordings at a sampling rate of 10,000Hz have been recorded with the 3Brain BioCam DupleX across all 4096 array channels of the Multiple Electrode (MEA) device. These raw recordings were incredibly large, with each raw recording containing more than $4x10^{11}$ samples across all channels and the entire time series. Two different sampling techniques were necessary to reduce the data load to a more feasible size. First, promising candidate channels were selected by research assistants in the Parrish Lab using the Xenon LFP Analysis Platform, reducing the 4096 array channels down to only a few handfuls per recording (Mahadevan, et. al). Second, each trace was downsampled to a sampling rate of 1,000Hz. This rate was deemed low enough to grant computational advantages, but high enough to still allow for detailed analysis in the spectral domain at frequencies up to its Nyquist frequency (1/2 the sampling rate, in this case 500Hz). Each one-second sample in the selected traces was then labeled as either baseline activity or seizure-like activity.

For the second question, we use a body of single channel LFP recordings from MEA data obtained by Dr. Ryley Parrish. With this data, the channels had already been selected from the array and downsampled to 100Hz. While this data was less fine than that used to answer the first question, it was deemed appropriate for the exploratory analysis conducted. Each candidate trace is labeled as previously described.

Figure 1 is included as an example of a full-length recording from a single electrode from the MEA device with the previously described labels. Notably, the three distinct activity patterns in this example appear visually different from one another. Baseline activity can be characterized in an LFP recording by a low amplitude, low voltage signal. It is most easily identified at the beginning of a trace. A non-SE-like seizure event is higher in amplitude than baseline activity, but is self-terminating.[3] SE-like seizure events are similar to non-SE-like seizure events in their high amplitude, but as described in Chapter 1, they last longer than five minutes, and are not self-terminating. Because of the chaotic nature of seizure signals, these patterns do not always hold to a specific voltage threshold across electrodes or recordings, and ofttimes, the patterns that indicate an SE-like event in one electrode are not helpful comparisons for another.

Figure 1 - Example Single Electrode Trace Recording From 3Brain C-MOS Accura Chip with Annotations Illustrating Non-SE and SE-like Seizure Events.

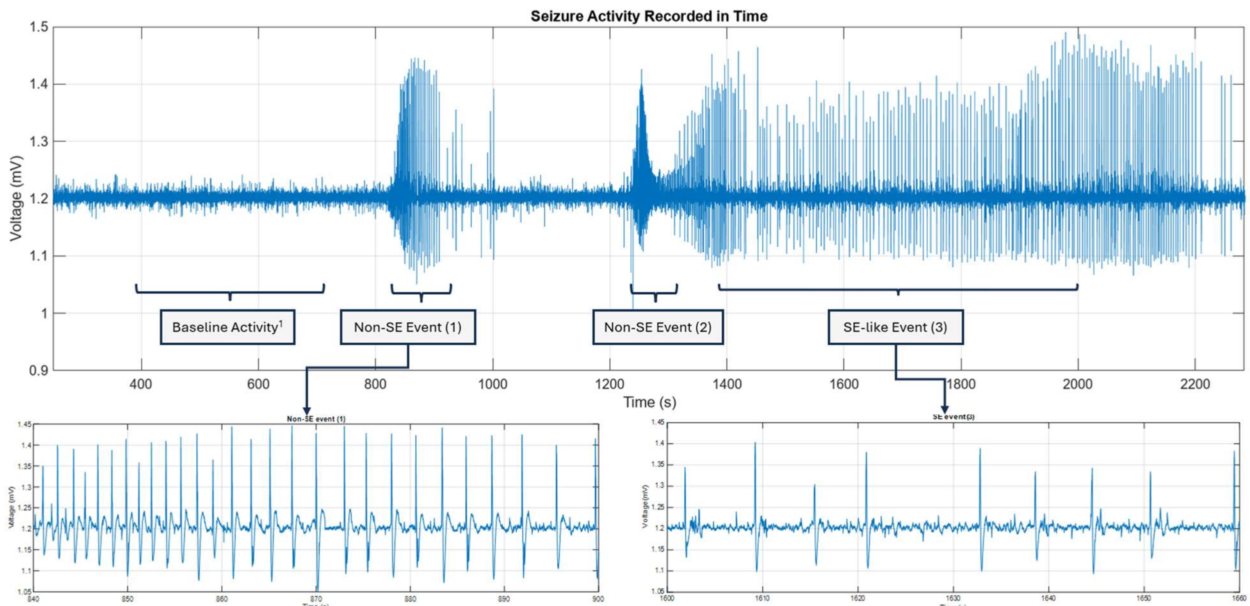

Methods

To answer the first primary question of this chapter, each downsampled single channel trace in the sample was processed as follows. First, each trace was transformed using the Fourier transformation. Because the Fourier transformation produces output with an imaginary component, the absolute magnitude of the transformed value was then computed. This gave the

[3] Most researchers will differentiate between types of self-terminating, non-SE-like seizure events, however, for the sake of this research, they have been lumped non-discriminatorily as non-SE-like seizure behavior.

magnitude for frequencies up to the Nyquist frequency. We created filter masks at multiples of 60Hz up to 480Hz to filter out frequencies that were artifacts of the machine used to record the LFP traces.

Next, for each of the bands previously mentioned, a bandpass filter was engineered and applied to a baseline, and seizure-like segment within each trace. A Hanning filter was also applied to each respective segment to taper the transition period between phases of the event. It was necessary to apply a correction to more completely capture the entire filtered signal based on the discrete nature of the trace. This correction was applied on both sides of the band, and the corrections were summed with the main calculation for the single sided power spectral density to achieve the final power estimate within that band.

The final power estimates within each band were then standardized by length of the segment, and proportion of ambient power. Ambient power was measured during a baseline period before a brain slice underwent any seizure like activity. This method of standardization made it possible to compare estimates between recordings and traces. The final, standardized power measures for each band in each segment class were then recorded along with the preparation paradigm and brain region of the class. Spectrograms were created across the entire signal and as seen in the results section of this chapter, visually demonstrate the differences between segments.

Lastly, this data was tested in aggregate using a Tukey's Honestly Significant Difference (HSD) Test to determine the statistical significance of brain region and preparation paradigm on average standardized power during seizure-like activity within a specified frequency band. This test included main factors for brain region and preparation paradigm, as well as an interaction term.

The second primary question in this chapter required different preprocessing. Sliding averages of the derivative, mean, variance, skewness, and kurtosis of each standardized trace were recorded. Anecdotally, these metrics were compared within traces as attempted features to differentiate baseline, non-SE-like seizure activity, and SE-like seizure activity. This analysis was not intended to be systematic in its structure, but instead serve as a testing ground for hypotheses on the usefulness of these statistical features for further differentiation between seizure types. These results will be discussed and reviewed in the following section.

Results

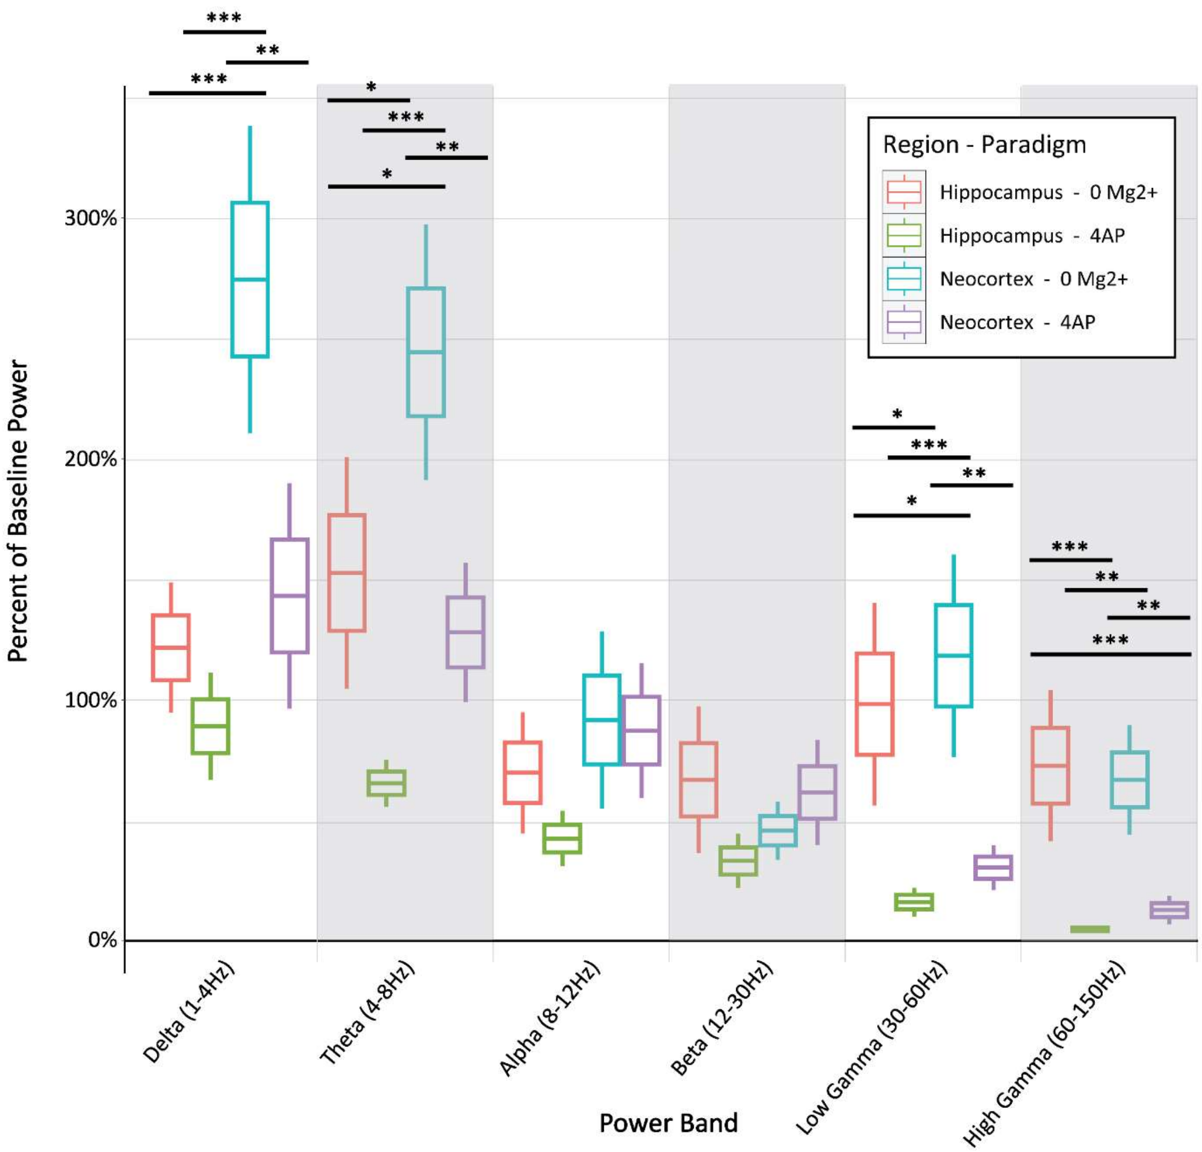

Figure 2 - Boxplot Comparing Banded Power Estimates in Aggregate Across Sample of Recordings

Figure 2 illustrates the results of Tukey's HSD test and encapsulates our findings on question one of this chapter. Specifically, significant differences between preparation paradigm and brain region were found in each frequency band with the exception of the alpha and beta bands. The striking similarities in percentage of baseline power among alpha and beta rhythms across traces is interesting and suggests it may be important to place more emphasis on other frequency bands for future analysis of this type. The 0 Mg^{2+} paradigm shows a higher percentage of baseline power in the low gamma and high gamma bands when compared to the 4AP paradigm. This finding is true across the neocortex and hippocampus. Traces from recordings in the hippocampus using the 0 Mg^{2+} paradigm also show the highest power levels in the delta and theta bands. Across all frequency bands, recordings from the hippocampal region prepared using the 4AP paradigm show the lowest percentage of baseline power. For researchers interested in that specific brain region, the 4AP prep style may not lead to satisfactory results with high power outputs.

Figure 2 also illustrates some similarities between the paradigms and brain regions. Namely, in both the Delta and Theta bands, more power is found in the 0 Mg^{2+} samples when compared with the 4AP samples. These observations are interesting because they may aid in the experimental design of seizure projects by helping researchers choose a brain region and preparation paradigm based on what specific frequency band they are most interested in studying.

Rolling Mean

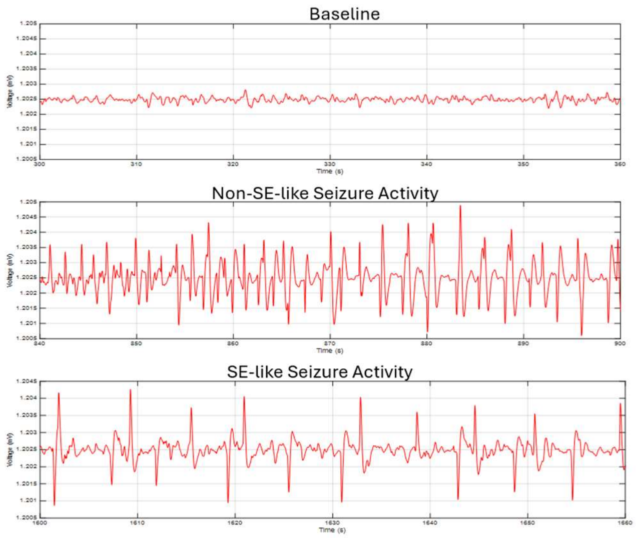

	Baseline[2]	Non-SE-like Seizure Activity	SE-like Seizure Activity
Min	1.2022	1.2006	1.2009
Max	1.2028	1.2049	1.2043
Mean	1.2025	1.2025	1.2025

Figure 3 - Compared Traces and Metrics of a Rolling Mean Calculated Over 10 Second Windows

Rolling Variance

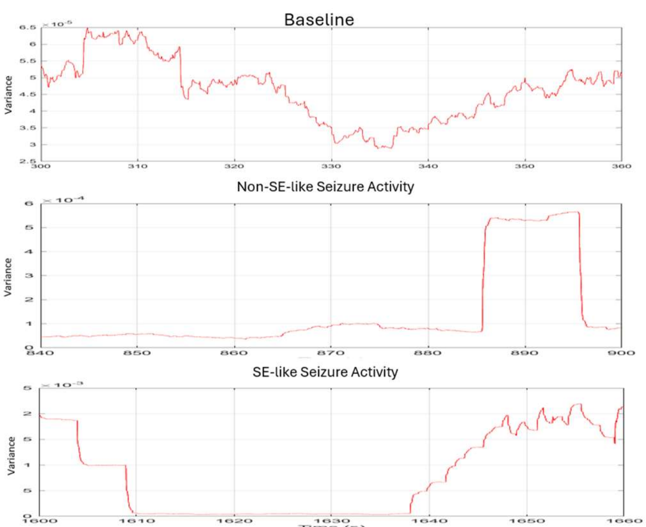

	Baseline[2]	Non-SE-like Seizure Activity	SE-like Seizure Activity
Min	2.8781e-05	3.5640e-05	3.2932e-05
Max	6.4980e-05	5.6580e-04	0.0022
Mean	4.6036e-05	1.4439e-04	7.7938e-04

Figure 4 - Compared Traces and Metrics of a Rolling Variance Calculated Over 10 Second Windows

Rolling Skewness

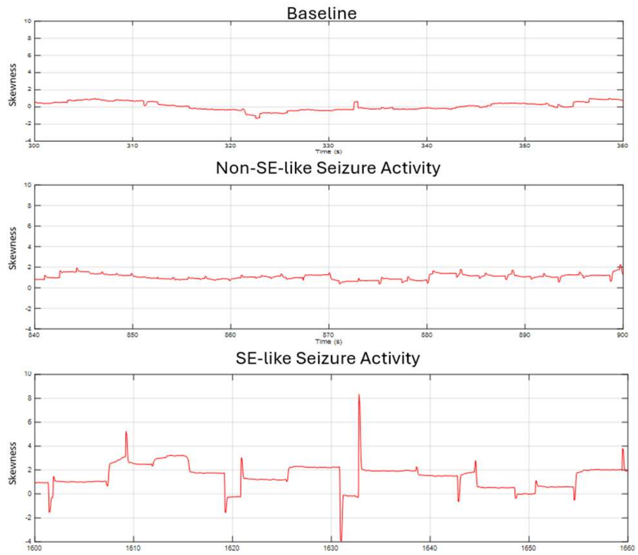

	Baseline[2]	Non-SE-like Seizure Activity	SE-like Seizure Activity
Min	-1.3400	-0.7267	-3.9566
Max	1.5433	2.2579	8.3464
Mean	0.0409	0.8269	1.5570

Figure 5 - Compared Traces and Metrics of a Rolling Skewness Calculated Over 10 Second Windows

Rolling Kurtosis

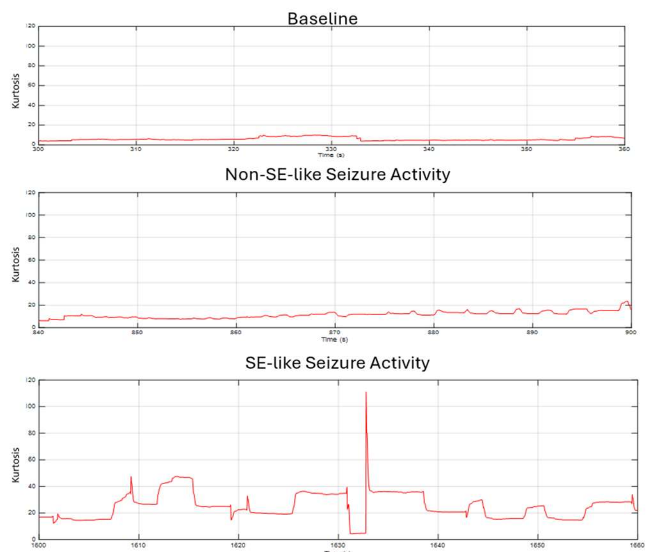

	Baseline	Non-SE-like Seizure Activity	SE-like Seizure Activity
Min	3.6688	3.4099	4.4219
Max	11.2080	23.4853	111.0814
Mean	6.1711	10.6160	25.0416

Figure 6 - Compared Traces and Metrics of a Rolling Kurtosis Calculated Over 10 Second Windows

Figures 3-6 tackle question two of this chapter and explore the anecdotal differences between statistical features of baseline, non-SE-like seizure activity, and SE-like seizure activity in a single channel recording. In figure 3, we see only small differences between the rolling mean of the three different behaviors. The greatest difference can be seen between the baseline and non-SE like seizure activity behaviors. In figure 4, we see interesting, unpredictable patterns appear during all three events. Specifically, during the non-SE-like seizure activity and the SE-like seizure activity, the rolling variance approaches very small values. Interestingly, across the larger body of traces anecdotally observed, other traces showed high variance during non-SE-like and SE-like activity. This seems to indicate that this statistical features may be difficult to use in future predictive models.

Larger differences in metric values can be found when using skewness and kurtosis, as seen in figure 5 and figure 6. Interestingly, in this recording, the non-SE-like seizure activity appears to be more similar to the baseline activity than to SE-like seizure activity. This was unexpected behavior and perhaps is indicative of some unique characteristic of the signal that clearly differentiates seizure types. While more exhaustive research must be done to systematize this statistical feature engineering process, these initial, exploratory calculations are informative as we proceed to the next chapter of the thesis.

Conclusion

Chapter 2 of this thesis paints a picture that will set up Chapter 3's dive into machine learning applications on this topic. In Chapter 2, we reviewed prior uses of signal processing in neuroscience and explained how a statistical feature set may be useful for differentiating between seizure types. We introduced the data and how it was collected, as well as explained the specific methodology used to analyze it. We explained how the results of our research in this chapter can be useful to researchers down the road who look to do research on seizure behavior in specific frequency bands. Lastly, we explained how some statistical features may assist in the classification of baseline, non-SE-like seizure activity, and SE-like seizure activity.

CHAPTER 3 – EXPLORING MACHINE LEARNING ON LFP RECORDINGS AND METHODS FOR SEIZURE CLASSIFICATION

In this chapter, we will examine the efficacy of tree-based supervised machine learning algorithms employing statistical features for the classification of SE-like events in LFP recordings. Additionally, we will review relevant literature on unsupervised changepoint detection and discuss its potential uses for SE classification.

We will begin by reviewing previous attempts to utilize tree-based supervised machine learning algorithms for classifying seizures in LFP and EEG recordings. This will include an investigation into feature engineering and methods for assessing the reliability of these approaches, as well as an exploration of the challenges they present and potential solutions. Following this, we will examine several changepoint detection models, discussing their respective advantages and disadvantages. We will then describe the methods and tests employed to evaluate a Random Forest model. Finally, this chapter will conclude with a discussion of the results from our Random Forest model tests and outline future steps for employing machine learning methods.

Prior Research on Machine Learning with LFP and EEG Recordings

Supervised machine learning tasks require both input features and output labels of a dataset, and can be divided into two primary tasks, prediction and classification. One common method for classification tasks is to use tree based algorithms, such as the random forest algorithm. The random forest algorithm is an ensemble method that relies on multiple simple decision trees (the number is specified by the researcher) to vote and reach a conclusion. Each simple decision tree is straightforward to interpret and utilize, however, they can be inaccurate on their own. Combining the decision trees into a ‘forest’ increases its accuracy. The forest is more accurate than the individual trees because it introduces randomness through bootstrap sampling (sampling with replacement) and the consideration of a random subset of features at each branch in the decision tree. Based on the sample and features associated with a decision tree in the random forest, the tree will cast a vote as to what class an observation belongs to. The random forest algorithm has many hyperparameters that researchers can tune to influence how many features are included in the random subset at tree branches, how many trees are trained, what level of impurity is required to warrant another branch, and more. This gives researchers more control and influence on the algorithm and makes it flexible.

Previously the random forest algorithm has been used for time-series classification tasks and has even been used on EEG data akin to what was cited in Chapter 2 of this thesis. In 2021, researchers used a random forest classifier to classify generalized and local seizures in EEG scalp recording data with an average accuracy of 96% and an F-measure of over 90% (Basri and Arif). In this study, power in specific frequency bands were used as features for the model, using the Fourier Transform over mutually exclusive 10 second windows to calculate these features. This study also noted that to reach high accuracy levels, it was important to over-sample the seizure classes to reduce the class imbalance present in the dataset.

Earlier in 2019, Wang, et. al., implemented a “random forest model combined with grid search optimization” to classify seizures. In this research, Wang fed the “mean energy, standard deviation, and high amplitude gamma frequency” as processed by a short time Fourier transform into a principal component analysis to form the features for the random forest. Compared to other machine learning models, Wang points to the following advantages of the random forest algorithm; namely that the random forest model is parallelizable and efficient, it can handle large features sets, and it can perform multi-class classification. This research produced a model with a 96.7% accuracy and an area under the Receiver Operating Curve of 99.0%, indicating excellent performance (Wang, et. al.).

The random forest algorithm also has some disadvantages that may impact its suitability for this classification task. First, because it has so many tunable parameters, it is necessary to do intensive tuning, which can be difficult and time-consuming. Also, the random forest algorithm does not inherently incorporate time-series information as part of its feature selection process. For an LFP recording, it is logical to assume that data points that are temporally correlated will likely belong to similar classes, and this would be important to encode in the model. It is important to keep these constraints in mind during any use of the random forest classifier.

In summary, previous research into seizure classification with tree-based algorithms has shown how the random forest algorithm can produce very accurate results when the proper feature selection and sampling methods are used. This research sets up a strong foundation for the research we conduct introduced later in the chapter. Our research differs in our feature set selection, and in the seizure type we are classifying.

While supervised machine learning requires labeled data, unsupervised machine learning is applicable when data lacks labels. This is particularly advantageous in scenarios like ours, where manually labeling large LFP datasets is challenging and time-consuming. Unlike supervised learning, unsupervised learning is not used for classification or prediction but is instead employed for tasks such as clustering, anomaly detection, and changepoint detection in time-series data. Next, we will briefly review literature on changepoint detection methods and explore their potential use case in detecting transitions from baseline activity to SE-like activity.

Offline changepoint detection refers to analyzing a time-series to find changepoints after all the data has been collected. Offline changepoint detection methods rely on a cost function and a search function. Cost functions can be parametric or non-parametric, and search functions can be exact or approximate. Truong, et. al., published a review of changepoint detection methods and various cost functions in 2020 that explains the implementations of these methods. Some of these methods are implemented in the Python package “ruptures.”

Unsupervised changepoint detection can be used for real time or online monitoring of EEG signals. Using a publicly available EEG database, researchers found favorable results using an auto-regressive (AR) linear model to model observed data and detect changes between baseline and seizure activity (Gao, et. al.). These researchers used statistical moments calculated across a sliding window as the features of the AR model, which fed into a randomized power martingale to make the decision as to whether there was a regime shift. These unsupervised methods resulted in very rapid identification of changepoints and achieved precision of 96.97% and recall of 97.66% (Gao, et. al.).

Both offline and online changepoint detection have a particularly interesting use case in the specific task of detecting transition from baseline to SE-like activity. Online changepoint detection clearly has valuable clinical applications that could save lives. With pharmacoresistant seizures, such as SE, response time is critical to preventing death or permanent injury, and high-performing online methods can significantly reduce response times. Offline changepoint detection methods have a different, albeit also important use case in SE classification. While standard clinical definitions exist for SE, much research is still being done on the diagnosis of SE and its definition in an EEG signal. Zafar and Aljaafari conducted a review on “EEG criteria for diagnosing nonconvulsive status epilepticus” and concluded that, “despite advances in EEG technology, the diagnostic dilemma of [SE] remains." A robust offline changepoint detection algorithm could be a powerful tool in helping to resolve this dilemma. The combination of this type of algorithm and the trained eye of a professional to help with tuning it could yield new insights into the statistical or time-frequency components of SE-like activity in EEG recordings, and result in more standardized methods for diagnosing it. This application would then enable easier labeling of EEG recordings and make the application of supervised machine learning methods a simpler task.

Methods

To answer the primary questions posed in this chapter, we will use the same collection of 100Hz downsampled EEG recordings as used to answer the second question in Chapter 2.[4] Each single channel recording used was labeled by research assistants for testing the supervised tree based methods.

First, we will discuss the methods used to engineer features for the data and prepare it for tree based modeling. A random forest model was fit with 24 time-series of training data selected randomly from the body of labeled time-series. Each single channel time-series was scaled to have a mean of 0. The model initially included features for the sliding mean, variance, skewness, and kurtosis of the scaled time-series over ten second windows. Initial iterations of the model tested the use of polynomial features, and found that the squared mean, skewness, and kurtosis were also helpful for model performance. The model was then used to predict the probability of an SE-like event at each sample point in the remaining six test time-series. With those probabilities, the area under the Receiver Operating Curve (ROC) was calculated and the

[4] See page 12 for more details on these traces.

optimal threshold was chosen to maximize the distance between the true positive rate and the false positive rate. This threshold was used to generate an intermediary label for each sample point. A Hanning window was applied across a 10 second window of these intermediary labels and the mean was calculated at each sample point. This final mean was then compared to a tuned threshold value to yield the final classification as SE-like behavior or non-SE-like behavior. This final step allowed information about sample points that were temporally correlated to increase the probability that they were of the same class. Due to performance constraints, grid search hyperparameter tuning was not feasible. Accuracy, precision, and recall were deemed appropriate measures for assessing the performance of the random forest model.

Results

Initial testing with the Random Forest model seemed promising. When the model was trained on a randomly selected sample from one time series and used to classify test data from the same time series, the model performed with 99.58% accuracy, 98.76% precision, and 99.52% recall.

As the training set for the model was expanded to include more time-series, its performance worsened significantly, performing with only a 64.97% accuracy, a 71.83% precision, and less than a 5% recall. This poor performance and serious digression from early success may be the result of a few factors. First, seizure-like activity, in particular, SE-like seizure activity, in EEG data can vary significantly in character between single channel time series. In chapter 2, we noted how preparation paradigm and brain region result in significantly different power outputs in certain frequency bands, which can have large impacts on the manifestation of SE-like activity. Besides this, EEG recordings are very noisy, and all this variability can seriously impact a supervised machine learning model's ability to learn the true behavior. Also, due to computational constraints, this model took longer than 12+ hours to train on the larger body of traces. This made tuning the model particularly difficult. Lastly, further measures could have been taken with this model to oversample the minority SE-like behavior class, which likely would have improved its performance.

Future Directions

Moving forward there is much exciting research to be done in the task of classifying SE-like activity. While tree based modeling for this task underperformed, there are other supervised machine learning techniques that have worked well for multiclass seizure classification or other similar tasks. For example, a hidden Markov model (HMM) is a statistical model used to describe observable events depending on internal factors that are not observable. When optimized with human learning, researchers have used HMMs to detect seizure like activity with greater than 92% accuracy, sensitivity, and specificity across multiple datasets (Chavan, et. al.). HMMs do require a signal to be stationary (i.e., the probability of transition from one state to another is independent of time), which may be a difficult assumption to prove true with this type of data. The largest advantage of an HMM is its ability to handle hidden states. In the context of EEG recordings, the true phenotypical state of the brain (pre-seizure, seizure, post-seizure) is not directly observable but can be inferred from the EEG signal.

Other researchers have found great success using long short-term memory (LSTM) neural networks as the primary engines for seizure-like activity classification. LSTM networks are especially powerful because they were built for time-series data and can capture temporal dependencies and patterns over a prolonged period. Khan, et. al, used a LSTM to perform binary and multiclass seizure-like activity classification on the EEG data, and achieved a 100% accuracy for the binary classification, and greater than 90% accuracy for the multiclass task.

This is particularly promising because the task of classifying SE-like activity must either be undertaken as a multiclass problem, or a series of binary classification steps.

Lastly, more research should be conducted on offline changepoint detection methods and its potential for introducing a more standard definition for SE-like activity in EEG recordings.

Conclusion

In Chapter 3, we reviewed the use of tree based models with time-frequency features to perform seizure classification tasks, and found literature suggesting excellent performance. We discussed other types of machine learning and how they could prove useful for future research. We found that a tree based model with a statistical feature set may work well for classification tasks constrained to one time-series, but does not generalize to larger datasets very well. Potential weaknesses with the model were investigated, and potential remedies to improve the model and fix said weaknesses were proposed.

CHAPTER 4 - CONCLUSION

In this thesis, we thoroughly examined how LFP data can be processed to extract time-frequency components and inform decisions on different types of seizure-like activity, such as SE. We reviewed common signal processing methods, and their applications on traces from MEA data collected in the Parrish Lab at BYU. Specifically, we compared the power found in different frequency bands during seizure-like events between different preparation paradigms and brain regions. We found significant differences between the 4AP and 0 Mg^{2+} paradigms in both the neocortex and hippocampus. Our findings suggest that researchers interested in studying the hippocampus should avoid the 4AP paradigm. We also concluded that the 0 Mg^{2+} paradigm was strongly suited for obtaining high power in the low and high gamma frequency bands. While further research may be necessary to confirm these findings, they are informative for future researchers interested in studying power levels using these preparation methods. In chapter 2 we also explored statistical differences between select LFP signals displaying SE-like seizure activity, non-SE-like seizure activity, and baseline activity. We concluded that the largest differences between SE-like activity and non-SE-like activity could be found using moving metrics of skewness and kurtosis. This analysis also requires further validation to confirm its findings but was informative in the exploration of machine learning methods for seizure classification in chapter 3 of this thesis.

In chapter 3 we explored how tree-based machine learning methods with statistical features classified SE-like activity. Prior research suggested favorable performance, and a random forest model was fit to single trace EEG recordings and tested in two different test scenarios. While it performed well when used to predict seizure activity in the same trace that some of its sample data came from, it struggled with entirely new traces. This finding confirms our hypothesis that the highly variable nature of LFP data between traces is a serious challenge. Further testing using more advanced methods of noise reduction, standardization, and oversampling of the minority class is necessary to try and mitigate said challenge. New methods of using unsupervised machine learning methods to learn more about SE-like activity and classify it were explored. This thesis recommends a more detailed investigation of these unsupervised machine learning methods, including changepoint and anomaly detection methods, as next steps to consider in future projects.

REFERENCES

Alam, S. M. S., & Bhuiyan, M. I. H. (2013). Detection of Seizure and Epilepsy Using Higher Order Statistics in the EMD Domain. *IEEE Journal of Biomedical and Health Informatics*, *17*(2), 312–318. https://doi.org/10.1109/JBHI.2012.2237409

Amengual-Gual, M., Sánchez Fernández, I., & Loddenkemper, T. (2019). Patterns of epileptic seizure occurrence. *Brain Research*, *1703*, 3–12. https://doi.org/10.1016/j.brainres.2018.02.032

Basri, A., & Arif, M. (2021). Classification of Seizure Types Using Random Forest Classifier. *Advances in Science and Technology Research Journal*, *15*(3), 167–178. https://doi.org/10.12913/22998624/140542

Brophy, G. M., Bell, R., Claassen, J., Alldredge, B., Bleck, T. P., Glauser, T., LaRoche, S. M., Riviello, J. J., Shutter, L., Sperling, M. R., Treiman, D. M., & Vespa, P. M. (2012). Guidelines for the Evaluation and Management of Status Epilepticus. *Neurocritical Care*, *17*(1), 3–23. https://doi.org/10.1007/s12028-012-9695-z

Chavan, P. A., & Desai, S. (2023). Effective Epileptic Seizure Detection by Classifying Focal and Non-focal EEG Signals using Human Learning Optimization-based Hidden Markov Model. *Biomedical Signal Processing and Control*, *83*, 104682. https://doi.org/10.1016/j.bspc.2023.104682

Dawit, S., & Crepeau, A. Z. (2020). When Drugs Do Not Work: Alternatives to Antiseizure Medications. *Current Neurology and Neuroscience Reports*, *20*(9). https://doi.org/10.1007/s11910-020-01061-3

Different Ways to Predict Seizures. (2016, January 13). Epilepsy Foundation. https://www.epilepsy.com/stories/seizure-prediction-elusive-yet-important-goal

Drislane, F. W. (2024, February 1). *Convulsive status epilepticus in adults: Classification, clinical features, and diagnosis*. Www.uptodate.com. https://www.uptodate.com/contents/convulsive-status-epilepticus-in-adults-classification-clinical-features-and-diagnosis#:~:text=Historically%2C%20the%20International%20League%20Against

Frolov, N. S., Grubov, V. V., Maksimenko, V. A., Lüttjohann, A., Igor Khmelinskii, Pavlov, A. N., Evgenia Sitnikova, Pisarchik, A. N., Jürgen Kurths, & Hramov, A. E. (2019). Statistical Properties and Predictability of Extreme Epileptic Events. *Scientific Reports*, *9*(1). https://doi.org/10.1038/s41598-019-43619-3

Gao, Q., Alaa Hamza Omran, Yasamin Baghersad, Mohammadi, O., Mohammed Ayad Alkhafaji, Kareem, A., Sara Hakem Al-Khafaji, Emami, N., D. Toghraie, & Mohammad Javad Golkar. (2023). Electroencephalogram signal classification based on Fourier transform and Pattern Recognition Network for epilepsy diagnosis. *Engineering Applications of Artificial Intelligence*, *123*, 106479–106479. https://doi.org/10.1016/j.engappai.2023.106479

Gao, Z., Lu, G., Yan, P., Lyu, C., Li, X., Shang, W., Xie, Z., & Zhang, W.-M. (2018). Automatic Change Detection for Real-Time Monitoring of EEG Signals. *Frontiers in Physiology*, *9*. https://doi.org/10.3389/fphys.2018.00325

Goodkin, H. P., & Kapur, J. (2003). Responsiveness of Status Epilepticus to Treatment with Diazepan Decreases Rapidly as Seizure Duration Increases. *Epilepsy Currents*, *3*(1), 11–12. https://doi.org/10.1046/j.1535-7597.2003.03104.x

Haderlein, J. F., Peterson, A. D. H., Zarei Eskikand, P., Cook, M. J., Burkitt, A. N., Mareels, I. M. Y., & Grayden, D. B. (2023). *Path Signatures for Seizure Forecasting*. https://arxiv.org/pdf/2308.09312

Khan, S. U., Jan, S. U., & Koo, I. (2023). Robust Epileptic Seizure Detection Using Long Short-Term Memory and Feature Fusion of Compressed Time–Frequency EEG Images. *Sensors*, *23*(23), 9572. https://doi.org/10.3390/s23239572

Kramer, M. A., Truccolo, W., Eden, U. T., Lepage, K. Q., Hochberg, L. R., Eskandar, E. N., Madsen, J. R., Lee, J.-H., Maheshwari, A., Halgren, E., Chu, C. J., & Halgren, E. (2012). *Human seizures self-terminate across spatial scales via a critical transition*. *109*(51), 21116–21121. https://doi.org/10.1073/pnas.1210047110

Letícia, A., Hossam Tharwat Ali, Elrefaey, A., Elejla, S. A., & Jamir Pitton Rissardo. (2023). Somatosensory Auras in Epilepsy: A Narrative Review of the Literature. *Medicines*, *10*(8), 49–49. https://doi.org/10.3390/medicines10080049

Mahadevan, A., Codadu, N. K., & Parrish, R. R. (2022). Xenon LFP Analysis Platform Is a Novel Graphical User Interface for Analysis of Local Field Potential From Large-Scale MEA Recordings. Frontiers in Neuroscience, 16. https://doi.org/10.3389/fnins.2022.904931

Mehla, V. K., Singhal, A., Singh, P., & Pachori, R. B. (2021). An efficient method for identification of epileptic seizures from EEG signals using Fourier analysis. *Physical and Engineering Sciences in Medicine*. https://doi.org/10.1007/s13246-021-00995-3

Park, Y., Luo, L., Parhi, K. K., & Netoff, T. (2011). Seizure prediction with spectral power of EEG using cost-sensitive support vector machines. *Epilepsia*, *52*(10), 1761–1770. https://doi.org/10.1111/j.1528-1167.2011.03138.x

Prasad, C., & Gaitonde, D. V. (2022). A time-domain linear method for jet noise prediction and control trend analysis. *Aerospace Science and Technology*, *121*, 107377. https://doi.org/10.1016/j.ast.2022.107377

Roy, S., Asif, U., Tang, J., & Harrer, S. (2020). *Seizure Type Classification using EEG signals and Machine Learning: Setting a benchmark*. https://arxiv.org/pdf/1902.01012

Schmelzle, M. (2010). *Option Pricing Formulae using Fourier Transform: Theory and Application* . https://pfadintegral.com/docs/Schmelzle2010%20Fourier%20Pricing.pdf

Signal Processing 101. (2015, December 17). IEEE Signal Processing Society. https://signalprocessingsociety.org/our-story/signal-processing-101

Singh, K., & Malhotra, J. (2022). Two-layer LSTM network-based prediction of epileptic seizures using EEG spectral features. *Complex & Intelligent Systems*. https://doi.org/10.1007/s40747-021-00627-z

Stubbs, I. W., Blotter, M. L., Jacob H. Norby, & Holmes, M. (n.d.). *High Quality Seizure-Like Activity from Acute Brain Slices Using a Complementary Metal-Oxide-Semiconductor High-Density Microelectrode Array Systems*.

Truong, C., Oudre, L., & Vayatis, N. (2020). Selective review of offline change point detection methods. *Signal Processing*, *167*, 107299. https://doi.org/10.1016/j.sigpro.2019.107299

Wang, X., Gong, G., Li, N., & Qiu, S. (2019). Detection Analysis of Epileptic EEG Using a Novel Random Forest Model Combined With Grid Search Optimization. *Frontiers in Human Neuroscience*, *13*. https://doi.org/10.3389/fnhum.2019.00052

Wang, Z., & Paolo Mengoni. (2022). Seizure classification with selected frequency bands and EEG montages: a Natural Language Processing approach. *Brain Informatics*, *9*(1). https://doi.org/10.1186/s40708-022-00159-3

Wylie, T., Sandhu, D. S., Goyal, A., & Murr, N. (2020). *Status Epilepticus*. PubMed; StatPearls Publishing. https://www.ncbi.nlm.nih.gov/books/NBK430686/

Zafar, A., & Danah Aljaafari. (2023). EEG criteria for diagnosing nonconvulsive status epilepticus in comatose - An unsolved puzzle: A narrative review. *Heliyon*, *9*(11), e22393–e22393. https://doi.org/10.1016/j.heliyon.2023.e22393