# Detection and Characterization of Epileptic Seizures using Machine Learning

**Abhejit Rajagopal**
abhejit@ece.ucsb.edu
Scientific Computing Group

**Rohan Jain**
rohanjain@umail.ucsb.edu
Vision Research Lab

University of California, Santa Barbara 93106

***Abstract*—In this project, we explore various methods of constructing patient-specific classifiers and characterizations of epileptic seizure state using multichannel noninvasive scalp electroencephalogram (EEG). Specifically, we implement both a binary state classifier using a kernalized support-vector machine (SVM) and a generic characterization using kernel discriminant analysis (KDA). While the seizure-detection accuracy of our binary classifier is comparable with those reported in literature, we notice a considerable number of false positives before the onset of the seizure. By constructing a patient-specific manifold using the recently developed method of spectral regression for KDA (SR-KDA), we show how to track the state trajectory to observe the gradual (rather than abrupt) transition from normal to seizure-like dynamical states. Our experiments show that while one contiguous pre-seizure period exhibits a noticeably smooth transition, it is largely uncorrelated with subsequent seizures (exhibiting high non-stationarity). Finally, we discuss possible schemes of deducing a continuous measure of seizure state using kernel-space projections.**

***Keywords*—EEG, epilepsy detection, machine learning**

## I. Introduction

Epileptic seizures are brief aberrations in the working of the nervous system, specifically the brain. Colloquially called "fits", they are caused by abnormal excessive or synchronous neural activity in the brain [1]. Epilepsy patients are predisposed to suffer from frequent seizures that can result in lapse of concentration and whole-body convulsions. Most clinicians analyze EEG (electroencephalogram) data to study the characteristics, including the onset, of seizures. It becomes imperative to design a system which automatically detects the onset of seizures with high accuracy and low latency.

EEG measures the electrical activity of the brain using electrodes that are uniformly arrayed on the scalp. An EEG channel is formed by taking the difference between potentials measured at two electrodes, and captures the summed potential of millions of neurons. The 10-20 system of electrode placement (Fig.1) has been followed in the CHB-MIT dataset which is the subject of our study. It is common knowledge among neurologists that epileptic seizures result in a group of neurons firing in a synchronous manner as against their normal asynchronous firing. It is this feature that is critically important in our quest for characterizing seizure EEG data using machine learning techniques.

Typically, there many challenges associated with the formulation of seizure detection as a data science problem, some of which are enumerated as follows:

- Seizure onset exhibits cross-patient variability. Different patients have differing seizure EEG patterns. Hence any sezure prediction system has to be patient-specific.
- Seizure analysis is best studied as a multiclass problem as shown in Section 3. Therefore, conversion to a binary classification problem leads to complications.
- Adult Scalp EEG exhibits perceptually less difference between seizure and non-seizure activity. Therefore, for this study we are concentrating on pediatriv patients.
- The EEG of epilepsy patients constantly transitions between regimes both within the seizure and non-seizure states, and is therefore a non- stationary process.
- Algorithm designers must consider tradeoffs between seizure detection sensitivity and specificity as well as detector latency and specificity.
- Seizures are uncommon and hence their detection methods must be capable of working with unbalanced/sparse data.

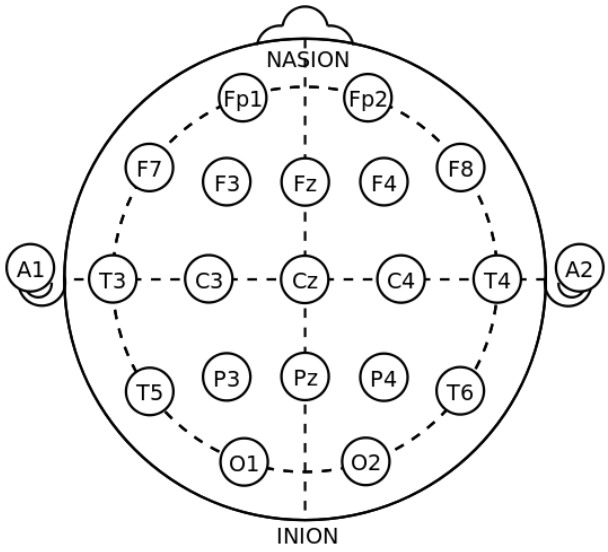



Figure 1. The $10 - -20$ system of electrode placement

## II. Binary Classification of Seizure State

We first implemented a binary classification of seizure state, inspired by the results reported by *Shoeb et al.*[2]. The goal for this section of the project was to learn a function $f(x)$

that maps a feature vector $x$, representative of a short-period of EEG, onto the binary labels $Y = \pm 1$ corresponding to seizure and non-seizure states respectively. We provide a short overview of our implementation, and subsequently discuss shortcomings in this approach.

### A. Dataset

EEG data and expert annotations were acquired from the open-source MIT-CHB database, which includes a total of 23 patients, most of which were pediatric subjects [3]. The EEG was sampled uniformly in 23 channels at 256Hz for roughly 46 non-contiguous hours for each patient. The provided annotations included the start an end times of seizure-like activity. On average, each patient experienced 3-5 seizures over 46-hour recording period. It is important to note here that measuring a patient's non-seizure brain activity over the course of a day is important because many normal brain activities exhibit synchronicity and temporal coherence that must not be mistaken for seizure episodes (e.g.*sleep spindles*) . As such, our training period includes nearly 24 hours of non-seizure data, which is possibly sufficient to characterize the dynamic range of normal brain activity.

### B. Feature Vector Design

Following the suggestion of [2], we extract spectral, temporal, and spatial features of a patient's EEG signal, and assemble these into a single 1-D featurevector. This is accomplished by first filtering each hour-long patient record using an 8-channel uniform filterbank ranging from $0 - 25$Hz (Figure 2) using a 2-second sliding window amd computing the power in each band. The 2-second window captures short-time spectral information, providing a maximum frequency resolution of 0.5Hz. Passing the each channel of the EEG through the same filterbank captures the spatial information related to electrode placement, producing a matrix $R^{8x23}$ which we subsequently vectorize to produce a $R^{184x1}$ feature vector corresponding to each 2-second sliding window.

In order to capture long-time spectral information, such as the temporal coherence or connectivity of state, we concatenate each 2-second window with two preceding non-overlapping 2-second windows to produce a final featurevector $x \in R^{552x1}$ representing a 6-second period (termed an epoch). The reasoning here is that seizure-events will be highly temporally correlated and, as such, should be identified by a significant change in long-term spectral characteristics rather than by instantaneous spikes in synchronicity. In summary, we incorporate short-term spectral, short-term spatial, and long-term temporal information into a single feature vector by concatenating multichannel spectral power densities (Figure 3).

### C. Kernalized SVM Classification

In order to learn a classification rule $f(x)$, we use LibSVM $v3.20$ [4] to construct a kernalized support-vector machine (SVM). The SVM provides a mapping from each 6-second epoch to binary labels corresponding to seizure and non-seizure states. The SVMs are trained using data of a single patient, and thus are patient-specific (i.e. the learning procedure must be applied for each patient independently).

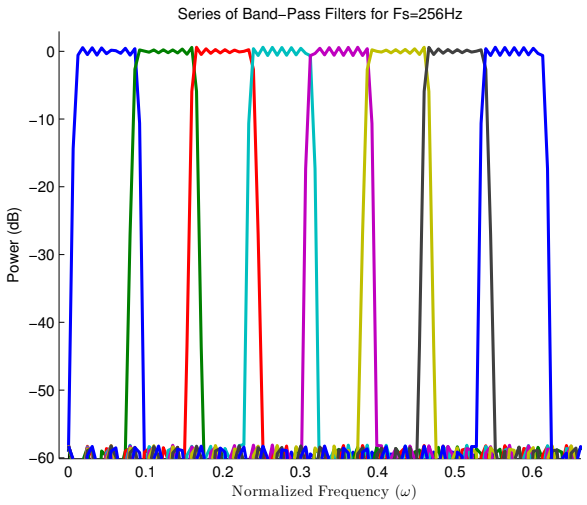



Figure 2. Magnitude response of M=8 uniformly-spaced FIR filterbank ($0 - 25$Hz), with passband gain of 1 and stopband attenuation $-60$dB. The corresponding phase transitions are nearly linear.

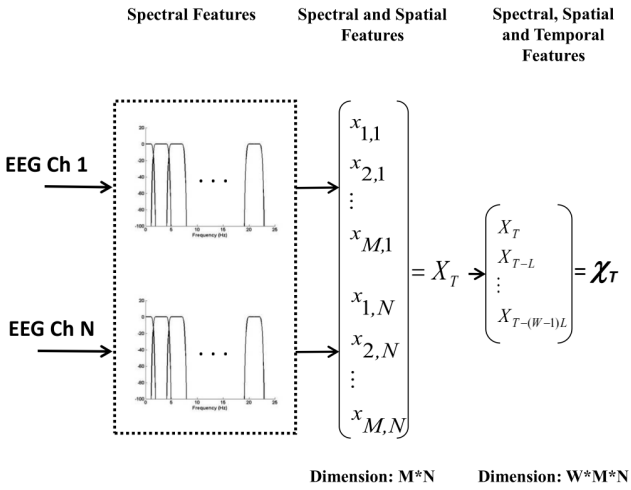



Figure 3. Incorporation of short-term spectral, spatial, and long-term temporal information into a single feature vector by concatenation[2].

We trained the C-SVM model on the first few seizure records and 24 hours of non seizure data belonging to each patient. We used a radial basis function(RBF) kernel, as did [2], as the two classes are expected to not be linearly separable. Moreover, we note that the physical location (value) of particular dimensions (frequencies) is not particularly useful since EEG is highly non-stationary; a kernalized approach uses the relative ampltiude and frequency information to discriminate between training samples. Hyperparameters (e.g. slack and gamma of RBF) were cross validated initially, but in later runs set as constants ($C = 1$, $\gamma = 0.1$).

### D. Classification Results and Discussion

Our trained model was reasonably effective at predicting seizures, giving a excellent sensitivity of $90.575\%$ on an average over five patients as compared to the base paper[2] which detects $96\%$ of all seizures in the full dataset (Figure 4).

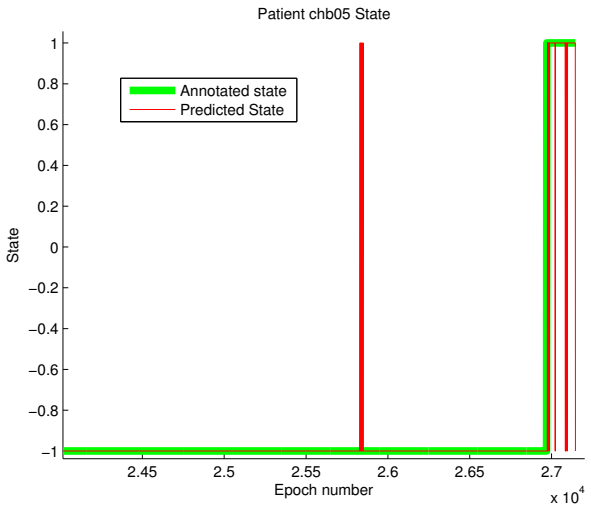



Figure 4. Predicted seizure labels vs groundtruth for patient CHB05 piror to seizure no. 3. There is an uncanny number of false positives detected prior to the onset of the seizure. Even still, the overall label prediction accuracy is high.

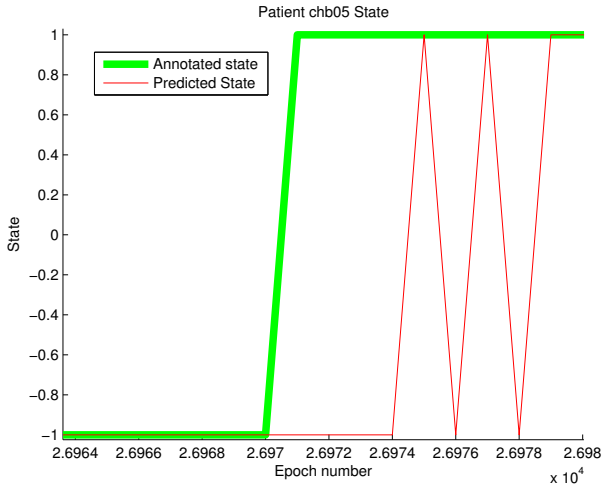



Figure 5. Latency of detection for seizure no. 3 for patient CHB05. is 10 seconds.

However we were not able to replicate the specificity of the base paper which detects only 2 false positives per 24 hours (Figure 5).

In our view, this can largely be attributed to the highly unbalanced dataset that we were working with. At most a single patient had 6 seizures in a 24 hour period, each lasting between 1-2 minutes. This scarcity of one class of data heavily biases the C-SVM and essentially calls for other methods of analysis. These can include One-Class SVM's, Ensemble of One-class SVM's, and Boosting methods among others. Our efforts to utilize One class SVM's did not give satisfying results and hence we broadened our approach to characterize seizures better. Specifically, we aim to understand the noticeable spike in the number of false positives preceding each annotated seizure period.

## III. Discriminant Analysis of EEG

Following up on our discussion on the false positive rate, we noticed that a binary classifier misses certain important information in the transient EEG waveform. Specifically, the popular consensus amongst neurologists is that the onset of seizures occurs gradually [5], [6], [7], [8]. As such, labeling epochs of EEG data as "seizure-like" or "non-seizure like" is considerably misleading. Although [2] addresses this concern, it is passed off as a difficult multi-class identification problem. While their approach is justified by means of their result, a highly accurate low-latency detector, the result provides little inside into the dynamics of seizures. The support vectors, for example, can be completely arbitrary and do not highlight biomarkers clinicians can look for manually in patient waveforms.

In this light, we set out to characterize EEG epochs using a discriminant analysis. For our work, we chose to employ the recently developed method of spectral regression for kernel discriminant analysis (SR-KDA) [9]. The goal of SR-KDA in this case is to find a regression model which leads to similar subspace projections $y_i \in Y$ for input data samples (i.e. pulses) $x_i \in X$ that are morphologically similar. Similar to the RBF used in the kernelized SVM, we employ a similarity measurement using a function of the Euclidean distance between feature vectors. However, unlike the previously presented approach, we do not provide hard labels to classify the result.

### A. SR-KDA of EEG epochs

Here, we take two approaches to characterize the data. In all cases, we apply KDA to the input data and visualize the results in 3 dimensions by selecting the top three eigenvectors of the affinity matrix $W$ corresponding to the top three eigenvalues. Next, input samples are projected to a reproducing kernel Hilbert space (RKHS) by constructing a kernel $K$ with respect to the training data and subsequently multiplying by the training eigenvectors.

We first applied an unsupervised KDA to differentiate EEG epochs, and visualized the resulting state transition (Figure 6). We notice a clear demarcation between seizure states based on the provided labels. Although this provides some measure of seizure state, it does not provide any insight into the transition in state.

To highlight the hypothesized temporal correlation of EEG epochs, we constrained the projections into RKHS by training on soft labels provided by annotations. The soft labels are normalized units of time, corresponding to the distance of an epoch from the onset and offset of annotated seizure periods. Specifically, each epoch's featurevector was appended with a weighted variable $w_i * t_i$ defined as follows:

$$t_{mid} = \frac{t_{onset} + t_{offset}}{2} \quad (1)$$

$$t_i = t(t) - t_{mid} \quad (2)$$

Finally, the vector $t$, corresponding to the relative time instances of each epoch, was normalized between 0 and 1 and multiplied by the empirical weight factor 200. $t$ is then used to

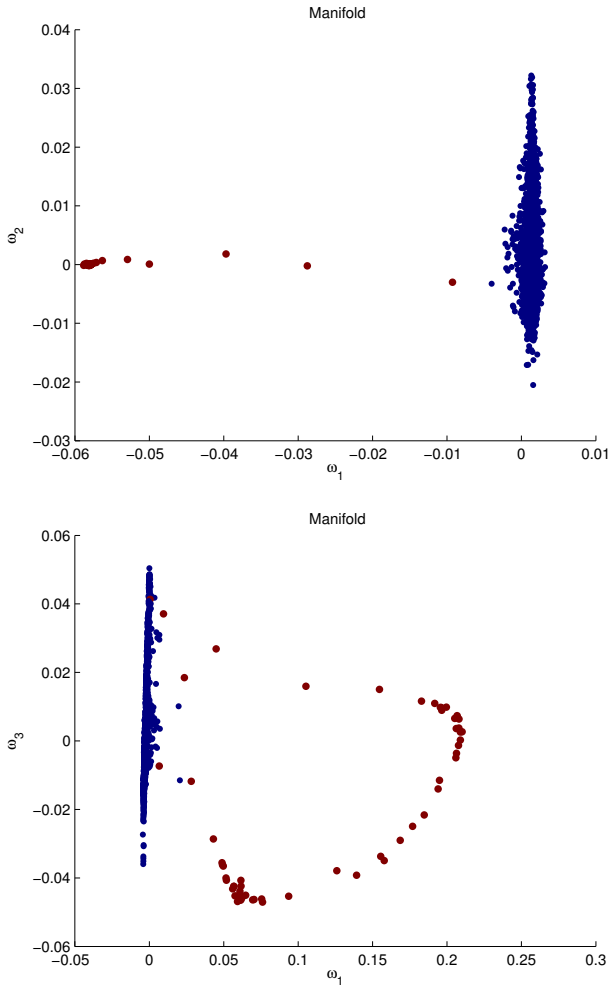



Figure 6. Both a supervised (a) and unsupervised (b) results are shown. Although there is a demarcation between seizure and nonseizure data, even for the unsupervised case, the threshold must be learned. As such, there is no clear improvement over an SVM.

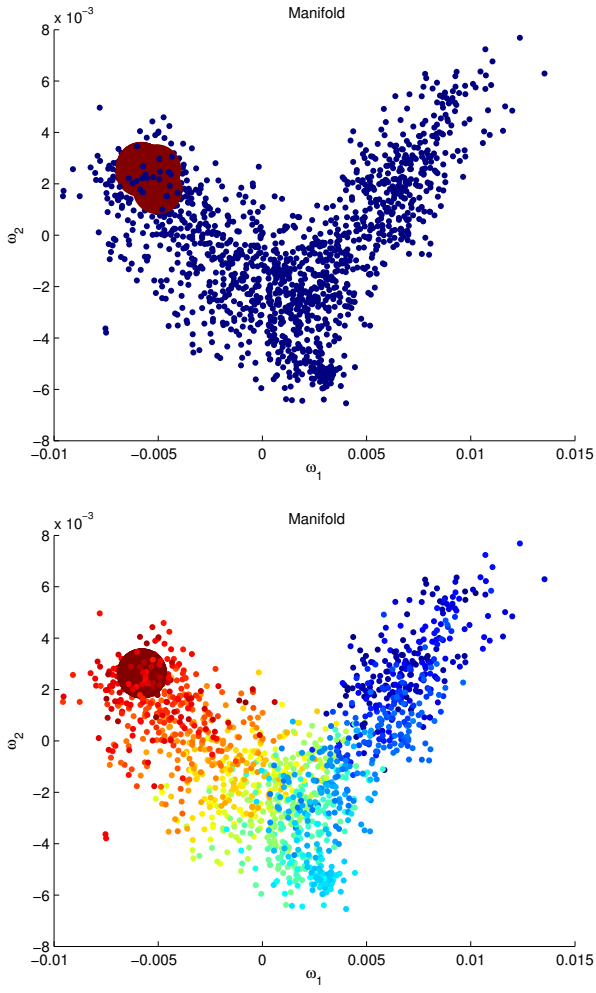



Figure 7. Seizure manifold trained using 1 seizure. The resulting RKHS plot shows a smooth transition between non-seizure state (dark blue) and seizure state (dark red). There is a remarkable correspondence between the state (a) and proximity to the onset (b).

warp distances in the RKHS by modifying the affinity matrix $W$. Because we do not redefine the kernel $K$, new epochs need not contain temporal information to be projected along the warped distance axes[10]. Our result, depicted in (Figure 7), demonstrates a clear transition in seizure state.

### B. Preliminary Results

Using the temporally constrained eigenvectors, we projected testing data onto the same Hilbert space. Employing a simple k-nearest neighbor (KNN) decision rule, we quantized each projected featurevector to a binary label. Although in some instances, we were able to predict the seizures before they occurred (Figure 8), we found that this characterization was not sufficient to build a robust classifier that consistently reports seizure and non-seizure states. We postulate that this is primarily because subsequent seizures are high uncorrelated to previous seizures, and as such do not project to meaningful locations on the constructed manifold. Visually, we can see that subsequent projections of testing data (or even training on a greater number of seizures) results in cluttered subspace projections (Figure 9). In a sense, our training period is

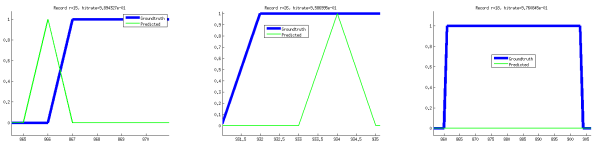


Figure 8. On patient CHB05, two seizures were used in training. Testing using the remaining data, the third seizure is predicted with negative latency (a), while the fourth seizure is detected with positive latency (b), and the fifth seizure is non detected at all.

produces a patient-specific, seizure-specific detector that is only valid for a short amount of time before and after a seizure.

## IV. Conclusion

In conclusion, we have investigated a robust way of detecting seizures using a binary kernelized SVM classifier, although this method ignores the transient characteristics of EEG and provides little insight into the mechanics of epilepsy. Furthermore, we have have demonstrated a method of discriminating between subtle changes in state by learning a patient-

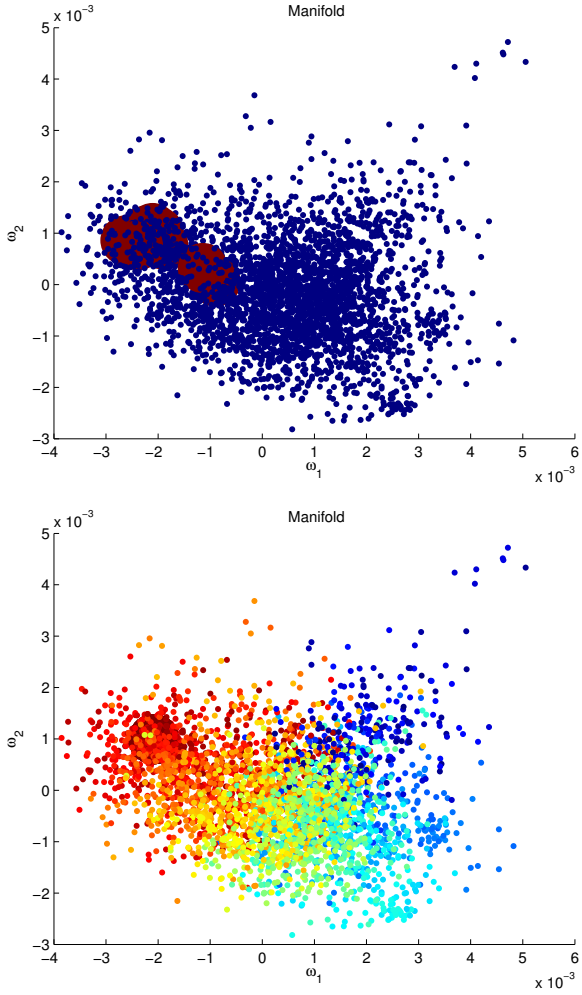



Figure 9. Seizure manifold trained using 2 seizures. Although the resulting RKHS plot shows a transition between non-seizure state (dark blue) and seizure state (dark red), the plot looks cluttered with no clear path. This is indicative of uncorrelated EEG signals (i.e. considerable non-stationarity in seizure signals).

specific manifold. Further research into this field could study EEG as a dynamical system. Specifically, we could model the transition from nonseizure to seizure states in a number of scenarios, where there seizure state corresponds to the nullspace (vanishing point) of the nonseizure manifold.

## REFERENCES


[1] Wikipedia. Wikipedia article on eplileptic seizure. [Online]. Available: http://en.wikipedia.org/wiki/Epileptic_seizure

[2] A. H. Shoeb and J. V. Guttag, "Application of machine learning to epileptic seizure detection," *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*.

[3] Goldberger, A. AL, G. LAN, H. L, I. JM, M. PCh, M. RG, M. JE, P. GB, S. C-K, and HE, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals," *Circulation Electronic Pages; http://circ.ahajournals.org/cgi/content/full/101/23/e215*.

[4] C.-C. Chang and C.-J. Lin, "Libsvm: a library for support vector machines," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 2, no. 3, p. 27, 2011.

[5] T. Zhang, J. Zhou, R. Jiang, H. Yang, P. R. Carney, and H. Jiang, "Pre-seizure state identified by diffuse optical tomography," *Scientific reports*, vol. 4, 2014.

[6] M. Goodfellow, C. Rummel, D. Garry, G. Baier, K. Schindler, and P. Glendinning, "State transitions in a model of intermittent seizure dynamics," 2012.

[7] M. Le Van Quyen, J. Martinerie, V. Navarro, P. Boon, M. D'Havé, C. Adam, B. Renault, F. Varela, and M. Baulac, "Anticipation of epileptic seizures from standard eeg recordings," *The Lancet*, vol. 357, no. 9251, pp. 183–188, 2001.

[8] J. Gotman, V. Levtova, and B. Farine, "Graphic representation of the eeg during epileptic seizures," *Electroencephalography and clinical neurophysiology*, vol. 87, no. 4, pp. 206–214, 1993.

[9] D. Cai, X. He, and J. Han, "Speed up kernel discriminant analysis," *The VLDB Journal*, vol. 20, no. 1, pp. 21–33, 2011.

[10] A. Rajagopal, R. Hamilton, and F. Scalzo, "Denoising of intracranial pressure signals using causal shape manifolds," *Biomedical Signal Processing and Control–under review*.