Epileptic Seizure Detection using EEG data

Hannah Li[1], Debashri Mukherjee[2]

CS 230, Winter 2018, Stanford University

Motivation:

- Epilepsy afflicts 200,00 people in the US every year, making it the fourth most common neurological disorder.
- The primary method for Epilepsy diagnosis is by reading EEG data- a process that is both expensive and fraught with inter-doctor and inter-patient differences.
- Automating this process can help make an accurate diagnosis of epilepsy in a resource-limited setting.
- This project aims to build such a system using Deep Learning methods on electroencephalogram (EEG) data.

Data:

- Data obtained in Adrzejak et.al. Indications of nonlinear deterministic and finite dimensional structures in time series of brain electrical activity
- 1-dimensional time-series data of 11500 observations.
- Each sample is a 1 second block of EEG recordings, split into 178 data points. The category label classifies each observation as a recording of seizure activity ('positive') or not ('negative').
- **Distribution of labels:**

	Positive	Negative
Occurrences	2300	9200
Mean	-4.729	-8.396
Avg. Variance (across time)	115702	4735

- For certain samples, the difference between a positive and negative EEG is not obvious

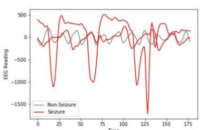

Model:

Model Architecture:

- 1D Convolutional Neural Net with L convolutional, ReLu, and MaxPool layers, followed by a fully connected layer
- Adam Optimizer with a learning rate of 0.009
- Trained using mini-batches of size 64
- Loss function: softmax cross entropy with logits

$$J(\theta) = -\frac{1}{m}\left[\sum_{i=1}^{m}(1-y^{(i)})\log(1-h_\theta(x^{(i)})) + y^{(i)}\log h_\theta(x^{(i)})\right]$$
$$= -\frac{1}{m}\left[\sum_{i=1}^{m}\sum_{j=0}^{1} 1\{y^{(i)}=j\}\log p(y^{(i)}=j|x^{(i)};\theta)\right]$$

Hyperparameter Search:

- Random coarse search over number of layers and number of channels for each layer
- Number of layers - 2 to 5
- Number of channels - 4 to 16

Optimal values found:
Number of layers: 4
Number of channels in each layer: 7, 5, 9, 7

Model Architecture:

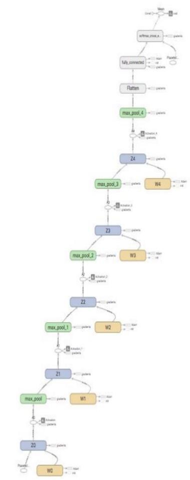

Results:

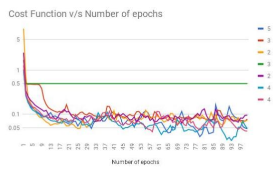

Cost functions for various models. Label indicates the number of convolutional layers. The number of filters at each layer is random between 4 and 16.

	Training Set	Test Set
Observations	**10,350**	**1,150**
Accuracy	**0.98**	**0.97**
Precision	**0.99**	**0.97**
Recall	**0.99**	**0.99**

F1 Score on test set- 0.98

Discussion:

- Model based on 1-Dimensional ConvNets performs better than other traditional machine learning models. eg. SVMs achieve an accuracy of 0.93 on the same task (Nicolaou et al., 2012)
- Despite using raw, minimally-processed input EEG data, the model performance is close to that of a GRNN that classifies hand-crafted features like Shannon Entropy, Energy and Standard Deviation to achieve 100% accuracy. (Swami et al., 2016)
- The system can be useful for other similar classification problems based on EEG brain signals
- Our algorithm correctly classified many seizure examples that resemble a non-seizure example, although it misclassified samples that appear to correspond to the start of a seizure

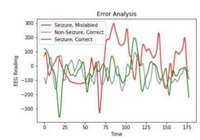

Future Work:

- Perform further error analysis specifically to lower False Negative rate.
- Conduct a finer hyperparameter search in the number of filters to use
- Implement an online forecasting algorithm by training NN to recognize signals specific to the start of a seizure

[1] hannahli@stanford.edu, [2] debm@stanford.edu