

Cross-site Epileptic Seizure Detection Using Convolutional Neural Networks

Danielle Currey*, David Hsu[†], Raheel Ahmed[‡], Archana Venkataraman[‡], Jeff Craley[‡]

*Department of Computer Science, Johns Hopkins University, Baltimore, MD

[†] School of Medicine and Public Health, University of Wisconsin, Madison, WI

[‡]Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD

Abstract—Automated epileptic seizure detection has been an active area of research for the last two decades. Yet few, if any, of these methods are used in clinical practice because they fail to generalize across different patient populations. We present three simple Convolutional Neural Network (CNN) architectures for seizure detection that are capable of generalizing across sites. The convolutional layers automatically learn robust and discriminative correlations directly from both the raw multi-channel scalp electroencephalography (EEG) signal and its short-time spectral representation. The models are trained on the publicly available Children's Hospital of Boston (CHB) data set in a leave-one-patient-out cross validation strategy. The trained model is then tested on a data set recorded at the University of Wisconsin (UW). We demonstrate that our CNNs achieve higher sensitivity than competing baselines, with only a minor increase in false positive rate. To our knowledge, this is the first work to achieve inter-hospital seizure detection without a significant drop in performance, thus providing an important benchmark for the seizure detection field.

I. INTRODUCTION

Epilepsy is a neurological disorder, characterized by bursts of abnormal electrical activity in the brain that manifest as seizures [1]. In 2015, more than 3.4 million people suffered from epilepsy in the United States, which underscores its public health relevance [2]. Scalp electroencephalography (EEG) is the first modality used for epilepsy diagnosis and characterization. Scalp EEG is typically recorded in the hospital over several days in order to capture only a handful of seizures. These recordings are visually inspected for seizure activity which is a time consuming and error-prone process.

While many automated seizure techniques have been investigated, the heterogeneity of epilepsy presentations adds complexity to the problem. Furthermore, to deploy pre-trained models at a new site, it is important to make sure that they generalize and are robust to different patient populations. Studies have shown that results suffer when models are used on new clinical data sets [3], [4]. In this work, we use deep learning to address the problem of cross-site generalization. Namely, we propose a group of simple network architectures based on Convolutional Neural Networks (CNNs).

Traditional machine learning approaches have applied many classifiers to the problem of seizure detection. Support vector machines are a popular classifier and have been used in [5]–[8]. This method identifies representative data points and uses

them to construct a boundary between seizure and non-seizure data. Similarly, Random Forests (RFs) have been used in seizure detection in [9], [10]. RFs use an ensemble of simple decision trees using random subsampling of the training data to train a classifier that is robust to overfitting.

As deep learning methods continue to mature, they have been applied to the problem of seizure detection [11]. The simplest deep learning architecture is the Multi-Layer Perceptron (MLP) which uses successive fully connected neural network layers to learn complex decision boundaries. Though simple compared to other deep learning architectures, the MLP has sufficient complexity to outperform traditional machine learning classifiers. This is shown in [10], where MLP classifiers outperform RF and support vector machines in detecting seizures using time-frequency feature representations. While typical MLP approaches rely on predetermined feature sets, [12] uses the MLP architecture to learn features for classification directly from the EEG signal.

Alternatively, CNNs have become popular due to their success in various computer vision applications [13]. Applying ideas directly from visual perception, time-frequency spectrogram images are often used as input for 2D CNNs for seizure detection. Yuan et al. constructed spectrograms to use in a seizure detection CNN [14]. The model used by [14] integrates inter- and intra-channel information in an auto-encoder structure. This is combined with a supervised seizure classifier to learn decision boundaries. Another approach by Khan et al. used wavelets to capture time-frequency structure, using them as input to a CNN architecture [15].

The use of one-dimensional convolutions has also been proposed, where convolutions are applied directly to short windows of signals. These filters are capable of partitioning the signals into discriminative frequency ranges, similar to a filterbank. In this way, 1D CNNs automatically learn the bands of interest from the data. This approach is used by Wei et al. where a 5-layer convolutional network is used [16]. Zou et al. applies a 1D CNN approach, sharing filters across channels to learn more robust representations of the data [17]. Similarly, [18] couples a 1D CNN with a recurrent neural network to analyze seizure activity at short and long timescales.

In this paper, we introduce a set of simple CNN architectures to address the challenges of generalization. The convolution operations automatically learn important features about the frequency and phase information from the data,

This work was supported by the NSF CAREER award 1845430 and a Johns Hopkins University Discovery Award.

TABLE I
CLINICAL ATTRIBUTES OF THE CHB AND UW DATA SETS.

	CHB Training data set	UW Generalization data set
Total Recording Time	3527743 s (980 hours)	35164s (9.8 hours)
Number of Recordings	683	53
Average Seizure Duration	60 s	68 s
Minimum/Maximum Seizure Duration	6 / 752 s	13 / 212 s
Average Seizures per Patient	5.75	4.4
Minimum/Maximum Number of Seizures per Patient	3 / 14	1 / 18
Average Recording Time per Patient	40.7 hrs	0.8 hrs
Minimum/Maximum Recording Time	19 hrs / 156 hrs	0.1 hr / 3.1 hrs
Average Seizure Time per Patient	7.7 min	6.9 min
Minimum/Maximum Seizure Time per Patient	1.4 min / 33.2 min	0.83 min / 42 min

and the low complexity of the proposed models reduces the likelihood of overfitting to a single site. The CNNs capture time-frequency information in different ways. For the 2D CNN this information is obtained directly from the spectrogram input, and for the 1D CNN it is learned from the signal itself. We address the problem of cross-site generalization by training the models on a data set acquired from Children’s Hospital Boston (CHB), [5], [19] and evaluating it on a data set from the University of Wisconsin, Madison (UW). We compare our models to baselines that cover a representative range of machine learning and neural network approaches. In particular, we show that the CNN architectures outperform the Multi-Layer Perceptron (MLP) and Random Forest (RF) based approaches. We demonstrate that our proposed CNN architectures retain higher sensitivity on the UW data set. To our knowledge, this is the first demonstration of inter-patient cross-site generalizations in the seizure detection literature.

II. METHODS

A. EEG Data

To evaluate cross-site generalization, we use two disjoint data sets for training and testing. Patient characteristics are summarized in Table I. For training we use the publicly available Children’s Hospital of Boston (CHB) data set [5], [19]. This data set contains EEG recordings of 24 pediatric patients from the entire multi-day recording period. Importantly, it contains a large amount of both baseline and seizure EEG activity from a diverse group of patients. EEG files are provided in the longitudinal bipolar montage and were sampled at 256 Hz. This data set consists of 683 EEG files and approximately 980 hours of recordings. Each file is approximately 1 hour long and there are on average 5.75 seizures per patient.

The generalization data set was acquired at the University of Wisconsin (UW) and contains seizure recordings extracted from longer continuous monitoring recordings. This data set contains a diverse set of patients but is smaller and contains far less baseline EEG than the CHB data set. While its smaller size and prominence of seizure complicates training, the data set is still appropriate for validating generalization performance. This data set includes 12 pediatric patients with an average of 4.4 seizures per patient. The data was recorded at 256 Hz using the 10-20 common reference and was converted to the longitudinal bipolar montage for this work. In total, we have 53




Fig. 1. Rhythmic seizure activity. Corresponding peaks and troughs indicating a phase reversal localize this activity to the T8 electrode.

EEG recordings, each containing at least one seizure, totalling approximately 9.8 hours worth of recordings.

For each data set, the EEG signals were low-pass filtered at 30 Hz, as higher frequencies carry significant artifact activity and little relevant information for seizure detection. In addition, the signals were high-pass filtered at 1.6 Hz to remove DC trends and physiological artifacts. Since the data can be variable across patients and hospitals, it was normalized to mean zero and standard deviation one for each channel in each EEG file individually. One second, non-overlapping intervals were extracted and each second was labeled as seizure or baseline using annotations made by clinicians.

B. CNN Architectures

In this work we evaluate three different CNN architectures designed to encode time-frequency information. Motivated by the abnormal rhythmic, highly correlated EEG activity indicative of seizure, our first architecture is a 1D CNN. Figure 1 illustrates this activity, as rhythmic signal is observed in the difference channels F8-T8 and T8-P8. The change in polarity between the two difference channels, termed a phase reversal in the clinical literature, indicates that the rhythmic activity is maximal in the T8 channel [20]. 1D CNN architectures offer a natural way to capture both the rhythmicity of seizure activity and the correlation of neighboring channels. Each 1D convolution layer applies a set of filters to the input signal with access to cross-channel signals to capture correlated activity. Through hierarchical application of these layers, a filterbank-like representation that can be learned directly from the multi-




Fig. 2. The models used for seizure detection. (a) 1D CNN, (b) 2D CNN-Spectrogram, (c) CNN-Combined, (d) MLP baseline.

channel EEG signal is constructed. Taken together, the 1D CNN architecture is capable of learning to extract both relevant frequency and correlation information for seizure detection from the raw multi-channel EEG input signal.

The 1D CNN architecture is shown in Figure 2 (a). The model consists of 6 1D convolutional layers with kernels of size 5, stride of 1 and padding of 2, each followed by a 1D max pooling layer with kernels of size 2 and a stride of 2. 1D batch normalization was used after the max pooling layers. This was followed by 2 fully connected layers with 200 nodes each and the binary output layer.

The second approach, shown in Figure 2 (b), uses a 2D CNN operating on spectrogram images. Spectrograms use the fast Fourier transform to analyze frequency content of the EEG signal. Spectrograms were extracted from each channel over 1 second windows and images were formed with frequency bands from 0 to 30 Hz on one axis and channel on the other for an image size of 30×18 . The spectrogram images were first passed through a single 2D convolutional layer with a 5×5 kernel. Following the convolutional layer are two fully connected layers, each with 100 nodes, and a final fully connected layer which outputs a binary seizure detection prediction for the input image.

The third CNN model, shown in Figure 2 (c), concatenates the outputs from both the 1D and 2D CNNs to form an aggregated encoded representation. This concatenated hidden representation was classified using two fully connected layers with 200 nodes each and an output layer for classification. All networks used LeakyReLU for all activations and dropout with probability 0.2 between each fully connected layer.

The networks were trained for 20 epochs using ADAM with a learning rate of 0.001 and batch size of 512. The loss function was Cross Entropy Loss. Due to the prominence of baseline activity, weighted sampling was used to over-sample the seizure instances during training. In addition, the models

were trained only on files that contain at least one seizure.

Model outputs of seizure versus baseline can be noisy, so predictions are smoothed over 20 second windows to encourage contiguous predictions during evaluation. As this is a binary classification problem, outputs are between 0 and 1. The threshold to determine what is seizure versus baseline represents a trade off between sensitivity and false positives. This threshold was calibrated on the training set by maximizing sensitivity while keeping the total false positive detections below 120 seconds per hour. The threshold obtained on the training data was used on test data. Calibration was performed separately on the UW data to obtain new thresholds, but the parameters and architecture were not tuned in any way.

C. Baselines

In addition, we compare our proposed CNN networks to two MLP classifiers. Where the use of the convolution operation in CNN networks make implicit assumptions about the underlying structure of the input data, MLP classifiers treat each dimension of the input data as a separate feature. Here we use two time-frequency features, spectrogram and bandpass, to compare the time-frequency type extraction used by the CNN models. The former concatenates the spectrograms used in the 2D CNN-Spectrogram model into a single vector for input into the MLP. The latter was used in [21] and compresses the information in the spectrograms according to clinically observed brain wave frequencies in the theta (1–4 Hz), delta (4–8 Hz), alpha (8–13 Hz), and beta (13–30 Hz) bands. Both the MLP-Spectrogram and MLP-Bandpass baselines use two fully connected layers with 100 hidden nodes before a fully connected binary classification layer. Finally, RFs are used on the same bandpass and spectrogram features in the RF-Spectrogram and RF-Bandpass baselines.




Fig. 3. Model outputs from the generalization experiment. The 1D CNN identifies two seizures while alternative methods identify the fourth seizure only.

D. Cross-Site Generalization Experimental Design

The models were trained on the CHB data set using a leave one patient out cross validation strategy. Specifically, all files from a single patient were used for validation while the network was trained on the data from the remaining 23 patients. This procedure was performed for each patient to evaluate the ability of the model to generalize to unseen patients recorded at the same site. The best performing training model was then scored on the UW generalization data set to evaluate cross-site generalization.

E. Evaluation

We evaluate metrics at the level of independent EEG windows and seizure sequences. Window level metrics treat each 1 second window as an independent sample and as such do not consider the seizure as a contiguous time interval. Here we evaluate sensitivity, false positive rate, and area under the receiver operating characteristic (AUC-ROC) curve. Seizure level metrics evaluate the performance of each model by treating the entire seizure as one unit. At the seizure level, seizure detection sensitivity, detection latency, and false positive count per hour were computed. Predictions are considered to be seizure if the prediction is above the calibrated detection threshold and baseline if it is below. False positives per hour calculates the number of times each model incorrectly predicts a seizure per hour. Seizure level sensitivity considers each seizure prediction within an annotated seizure interval as a correct classification. Latency time is computed by taking the time difference between the first seizure prediction in the seizure interval and the clinically annotated onset.

In addition to our cross-site generalization experiment, we report these same metrics on models trained and tested on




Fig. 4. An EEG recording from the UW data set. The 1D CNN classification is shaded in blue. The model finds the seizure onset shown by the dashed line at 151 seconds, in accordance with the clinical annotation.

a single site. By comparing cross-site generalization performance to models trained in the original data sets, we can determine the performance loss due to generalization.

III. EXPERIMENTAL RESULTS

A. Intra-Site Results

Table II reports the intra-site and cross-site seizure detection performance. When trained on the UW data set, the 1D CNN achieves the highest window level AUC-ROC with the lowest false positive rate. Although the RF classifiers have the highest seizure level sensitivity, this performance comes at the cost of greatly increased false positive rates. While exhibiting worse overall performance, in the UW intra-site experiment the MLP

TABLE II
WINDOW AND SEIZURE LEVEL RESULTS FOR BOTH THE CHB AND UW DATA SETS.

	Sensitivity (Seizure)	Latency (s)	False Positives (Count / hr)	Sensitivity (Window)	False Positive Rate (Seconds / hr)	AUC-ROC
Cross-site Generalization, CHB → UW						
1D CNN	0.893	<u>-3.25</u>	8.12	0.584	<u>91.00</u>	<u>0.850</u>
2D CNN-Spectrogram	0.796	-11.42	5.70	0.524	127.37	0.836
CNN-Combined	0.783	-10.98	5.98	0.530	161.83	0.841
MLP-Spectrogram	0.564	7.03	6.94	0.255	134.68	0.753
MLP-Bandpass	0.469	9.81	7.20	0.226	229.22	0.685
RF-Spectrogram	0.871	18.70	4.52	0.476	88.87	0.873
RF-Bandpass	0.707	1.13	6.29	0.421	354.10	0.782
Single Site Training, UW						
1D CNN	0.967	-11.60	<u>16.80</u>	0.796	215.05	0.884
2D CNN-Spectrogram	0.981	-21.77	45.14	0.823	595.41	0.836
CNN-Combined	0.972	-17.15	27.84	0.822	361.41	0.867
MLP-Spectrogram	0.938	-22.69	41.24	0.707	632.74	0.746
MLP-Bandpass	0.992	-65.68	24.66	0.752	825.38	0.644
RF-Spectrogram	1.0	-44.70	26.10	0.909	565.97	0.870
RF-Bandpass	0.997	-108.54	13.55	0.916	293.93	0.847
Single Site Training, CHB						
1D CNN	0.957	0.06	9.54	0.881	131.47	0.989
2D CNN-Spectrogram	1	-6.34	8.21	0.998	127.26	0.998
CNN-Combined	1	-5.03	7.54	0.998	117.46	0.998
MLP-Spectrogram	0.477	16.21	4.54	0.227	53.79	0.864
MLP-Bandpass	0.768	1.27	7.71	0.533	120.18	0.863
RF-Spectrogram	0.803	2.87	6.17	0.694	113.17	0.902
RF-Bandpass	0.850	-139.74	8.79	0.72	312.71	0.893

baselines show seizure detection efficacy, achieving AUC-ROCs of 0.814 and 0.721 for the MLP-Spectrogram and MLP-Bandpass, respectively.

In the intra-site CHB experiment, all CNN models perform comparably with the combined network exhibiting the best performance across most metrics. Notably, the 1D CNN shows lower window level sensitivity (0.881) than the 2D CNN-Spectrogram and CNN-Combined models (0.998 and 0.998), while all three networks have similar seizure level sensitivity (0.957, 1, and 1) and AUC-ROC (0.989, 0.998, and 0.998). When compared to the CNN methods, the MLP and RF methods show a similar decrease in performance as in the intra-site UW experiment. Most methods show false positive rates close to 120 seconds per hour, indicating that the models misclassify false positives at the original calibration rate.

B. Cross-Site Generalization Results

In the cross-site generalization experiment, the 1D CNN achieves the highest seizure level and window level sensitivity. At the seizure level, the model correctly identifies 0.893 of seizures with a window level sensitivity of 0.548. In addition, the 1D CNN maintains a false positive rate below the calibration point. The 2D CNN-Spectrogram and CNN-Combined models achieve seizure level sensitivities of only 0.796 and 0.783, respectively. However, we note that the sensitivity performance in the 1D CNN comes at the cost of increased false positive count per hour, with 8.124. Comparatively the 2D CNN-Spectrogram and CNN-Combined models show false positives per hour of 5.703 and 5.982. With the exception of false positive rate, the MLP baselines show drastically worse performance across all metrics. This indicates that

these baselines fail to identify seizure activity when applied cross-site. Overall, the CNN models have significantly higher window level sensitivities than the MLP or RF baselines.

Figure 4 shows a seizure prediction from the 1D CNN model in a seizure from the UW data set. In the figure, windows labeled seizure are shaded blue. The 1D CNN model correctly predicts seizure activity corresponding to the clinically annotated seizure onset occurring at 151 seconds. Figure 3 shows model outputs from a patient with seizures that are difficult to identify. Although these seizures were not as well identified by any of the models, the 1D CNN identified the second and last seizure, and the other two CNN models identified parts of the final seizure, whereas the MLP and RF baselines only partially identify the last seizure.

IV. DISCUSSION

In this work we propose 3 simple CNN architectures for seizure detection and successfully show performance in a cross-site generalization experiment. To our knowledge, this experiment is the first in the seizure detection literature to show generalization between data sets recorded at different sites. Our first model uses 6 1D convolutional layers to encode the multi-channel EEG directly from the signal itself. Effectively, this method extracts representations from hierarchical filtering operations. As phase information between channels is preserved through the 1D convolution layers, we hypothesize that the 1D CNN encoder structure is uniquely able to learn filterbank type representations that encode cross-channel correlations as well. This structure encodes robust hidden representations capable of detecting seizures across different recording sites and patient populations.

Comparatively, the 2D CNN approach relies on spectrogram extraction as a preprocessing step. While spectrograms are a natural way to encode time-frequency information, they fail to encode cross-channel correlations which may be preserved by the 1D CNN. In the 2D CNN-Spectrogram and CNN-Combined models we see similar performance on intra-site experiments, however these models exhibit a slight drop in performance when applied across sites. As shown in Figure 3, the CNN models are able to identify seizures that the RF and MLP baselines miss. This can be explained by the higher seizure level sensitivity of the CNN models.

The CNN models have negative latency in the generalization task. As false positive rate remains near the calibration point in this task, this phenomena is likely due to a combination of low detection threshold and temporal smoothing, resulting in early detections. This phenomena is likely responsible for the low negative latencies in the intra-site CHB experiment as well. However, the high false positive rates in the UW intra-site experiment indicate the calibration point set during training does not generalize to the left out test patient. This highlights the difficulty of training and calibrating in to the same data in a small data set where overfitting is likely.

While showing efficacy in the intra-site seizure detection task, the MLP-Spectrogram and MLP-Bandpass baselines fail in cross-site generalization. Notably, the MLP-Spectrogram and 2D CNN-Spectrogram share the same inputs. However, the 2D CNN-Spectrogram significantly outperforms the MLP-Spectrogram approach. This is likely due to the 2D CNN encoding local similarities between neighboring frequencies and channels in the spectrogram image. Thus despite containing only one convolutional layer, the 2D CNN approach more easily captures relevant cross-channel and cross-frequency information in the spectrogram.

V. CONCLUSION

We have evaluated 3 CNN approaches for seizure detection which generalize both to new patients within the data set on which it is trained on as well as patients on an unseen data set. To our knowledge, this is the first work to observe inter-hospital generalization performance, where other studies have noted significant drops in cross-site performance. We compare the CNN approaches to MLP and RF approaches using time-frequency decompositions and observe superior generalization in the CNNs. Extensions to this work include training on larger and more diverse data sets in order to boost performance and further improve generalization. This work explores how models can be applied to different data sets and different patient cohorts, but further work is needed to further develop robust models that can be transferred across hospitals, epilepsy types, and patient populations.

REFERENCES

- [1] R. S. Fisher, C. Acevedo, A. Arzimanoglou, A. Bogacz, J. H. Cross, C. E. Elger, J. Engel Jr, L. Forsgren, J. A. French, M. Glynn *et al.*, "Ilae official report: a practical clinical definition of epilepsy," *Epilepsia*, vol. 55, no. 4, pp. 475–482, 2014.
- [2] M. M. Zack and R. Kobau, "National and state estimates of the numbers of adults and children with active epilepsy—united states, 2015," *MMWR. Morbidity and mortality weekly report*, vol. 66, no. 31, p. 821, 2017.
- [3] K. Saab, J. Dunnmon, C. Ré, D. Rubin, and C. Lee-Messer, "Weak supervision as an efficient approach for automated seizure detection in electroencephalography," *npj Digital Medicine*, vol. 3, no. 1, pp. 1–12, 2020.
- [4] J. R. Zech, M. A. Badgeley, M. Liu, A. B. Costa, J. J. Titano, and E. K. Oermann, "Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study," *PLoS medicine*, vol. 15, no. 11, p. e1002683, 2018.
- [5] A. H. Shoeb, "Application of machine learning to epileptic seizure onset detection and treatment," Ph.D. dissertation, Massachusetts Institute of Technology, 2009.
- [6] A. H. Shoeb and J. V. Gutttag, "Application of machine learning to epileptic seizure detection," in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 2010, pp. 975–982.
- [7] V. Sridevi, M. R. Reddy, K. Srinivasan, K. Radhakrishnan, C. Rathore, and D. S. Nayak, "Improved patient-independent system for detection of electrical onset of seizures," *Journal of Clinical Neurophysiology*, vol. 36, no. 1, p. 14, 2019.
- [8] M. Bandarabadi, C. A. Teixeira, J. Rasekhi, and A. Dourado, "Epileptic seizure prediction using relative spectral power features," *Clinical Neurophysiology*, vol. 126, no. 2, pp. 237–248, 2015.
- [9] D. Wu, Z. Wang, L. Jiang, F. Dong, X. Wu, S. Wang, and Y. Ding, "Automatic epileptic seizures joint detection algorithm based on improved multi-domain feature of ceeg and spike feature of aeeg," *IEEE Access*, vol. 7, pp. 41 551–41 564, 2019.
- [10] E. Alickovic, J. Kevric, and A. Subasi, "Performance evaluation of empirical mode decomposition, discrete wavelet transform, and wavelet packed decomposition for automated epileptic seizure detection and prediction," *Biomedical signal processing and control*, vol. 39, pp. 94–102, 2018.
- [11] A. Craik, Y. He, and J. L. Contreras-Vidal, "Deep learning for electroencephalogram (eeg) classification tasks: a review," *Journal of neural engineering*, vol. 16, no. 3, p. 031001, 2019.
- [12] A.-M. Tăuțan, M. Dogariu, and B. Ionescu, "Detection of epileptic seizures using unsupervised learning techniques for feature extraction," in *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2019, pp. 2377–2381.
- [13] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [14] Y. Yuan, G. Xun, K. Jia, and A. Zhang, "A multi-view deep learning framework for eeg seizure detection," *IEEE journal of biomedical and health informatics*, vol. 23, no. 1, pp. 83–94, 2018.
- [15] H. Khan, L. Marcuse, M. Fields, K. Swann, and B. Yener, "Focal onset seizure prediction using convolutional networks," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 9, pp. 2109–2118, 2017.
- [16] Z. Wei, J. Zou, J. Zhang, and J. Xu, "Automatic epileptic eeg detection using convolutional neural network with improvements in time-domain," *Biomedical Signal Processing and Control*, vol. 53, p. 101551, 2019.
- [17] L. Zou, X. Liu, A. Jiang, and X. Zhou, "Epileptic seizure detection using deep convolutional network," in *2018 IEEE 23rd International Conference on Digital Signal Processing (DSP)*. IEEE, 2018, pp. 1–4.
- [18] J. Craley, E. Johnson, C. Jouny, and A. Venkataraman, "Automated inter-patient seizure detection using multichannel convolutional and recurrent neural networks," *Biomedical Signal Processing and Control*, vol. 64, p. 102360.
- [19] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals," *circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
- [20] L. V. Marcuse, M. C. Fields, and J. J. Yoo, *Rowan's Primer of EEG E-Book*. Elsevier Health Sciences, 2015.
- [21] J. Craley, E. Johnson, and A. Venkataraman, "A spatio-temporal model of seizure propagation in focal epilepsy," *IEEE Transactions on Medical Imaging*, pp. 1–1, 2019.