Adaptive Chirplet Transform-Based Neural Network for Seizure State Detection

Rocklen Jeong
Dept. of Elec. and Comp. Eng.
University of Toronto
Toronto, Canada
0009-0005-5596-9016

Nishant Kumar
Dept. of Elec. and Comp. Eng.
University of Toronto
Toronto, Canada
0009-0007-7169-9241

Steve Mann
Dept. of Elec. and Comp. Eng.
University of Toronto
Toronto, Canada
0000-0003-0363-3690

***Abstract*—Epilepsy is a neurological disease affecting millions worldwide. This paper proposes the use of deep learning, combined with advanced signal-processing techniques employing the Adaptive Chirplet Transform (ACT), to detect seizure states in real time. The ACT is capable of detecting and analyzing the fast rhythmic activity that occurs during seizure periods, which traditional time and frequency domain features often fail to effectively capture. Data from the CHB-MIT dataset was used, and each patient was analyzed to ensure sufficient and relevant data were available. A convolutional neural network using inception-style blocks was trained and evaluated using leave-one-out cross-validation (LOOCV) for each patient. Two different models using different ACT features and further postprocessing techniques were also evaluated for a total of four different sets of results. This study yielded promising results, with a recall of 84.54% and an F1 score of 81.58% for the best model. This study introduces an advanced signal-processing technique optimized for hardware implementation, demonstrating potential applications in future real-time seizure-state detection systems.**

***Index Terms*—Seizure, epilepsy, EEG, chirplet, neural network, detection**

I. Introduction

Epilepsy is a neurological disease that affects 50 million individuals worldwide [1]. It often entails recurring seizures, caused by excessive electrical discharges of brain cells. Epilepsy can be a devastating disease, as it can cause loss of control in movement and other cognitive activities. This can also lead to mental and financial burdens due to the medications necessary to treat the disease. Patients may also not be able to live safely even with medication, as it is reported that a third of adults living with epilepsy are considered drug-resistant [2]. Surgery is a potential solution, but it carries complications and may lead to further issues.

With recent advances in wearable technologies and implantable devices [3]–[7], there is potential to develop a system capable of detecting seizures and administering necessary treatments to support patients. Previous works have utilized various time and frequency domain features, such as power spectral density and wavelet features [8]–[14], to detect seizures. However, these models often require substantial computational resources or large amounts of data, as well as deep neural networks. Furthermore, window sizes are often large, which is detrimental for seizure state detection as the latency from the seizure onset would be too large to diagnose and react to the seizure in time.

This paper proposes the use of the Adaptive Chirplet Transform (ACT) to analyze and classify EEG data, leveraging both time and frequency domain features, which is particularly effective for processing the complex, intricate EEG data associated with seizures. With recent advances in chirplet processing, such as GPU Chirplet [15], we can now process EEG data in real time. Furthermore, our choice of a machine learning model that utilizes 1-D convolutional layers in an inception-style block aims to reduce computational complexity and runtime. We aim to efficiently and accurately classify EEG data to detect seizure states.

II. Background

A. Seizure Detection

Seizure detection has been a challenge for decades, with no definitive solution. Without a reliable measure to detect seizure onsets, millions of patients worldwide have suffered from seizures. There is a need for a system capable of identifying seizures and taking autonomous actions to reduce the harm.

EEG data is often classified into four states: interictal, ictal, preictal, and postictal. The interictal state is the period between seizures, whereas the ictal state is the period during the seizure, from onset to offset. The preictal state and postictal state are periods preceding and following a seizure, respectively. A significant challenge has been distinguishing between the preictal and interictal states for seizure prediction, as EEG activity in the preictal state often resembles that in the interictal state. The ictal state is the most notable, with activity typically occurring in the frequency ranges of 3-30 Hz, with rhythmic low frequency activity near 0.5 Hz, and fast rhythmic activity in the 40-50 Hz range [16]. For our work, we will examine only the interictal and ictal states, as our current goal is to evaluate the potential of the ACT for seizure detection rather than prediction.

B. Adaptive Chirplet Transform

The chirplet is a windowed chirp function in which the frequency varies linearly with an independent variable, such as time or space [17], [18], [20] as is demonstrated in Fig. 1.

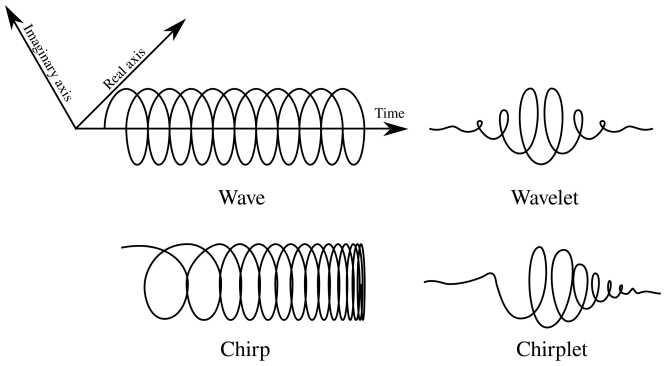

Fig. 1: Chirplet transform fundamentals, with a comparison between waves and wavelets to chirps and chirplets [17]–[19]

The ACT revolves primarily around the quadratic chirplet (q-chirplet), and the equation for the Gaussian chirplet is defined in Eq. 1. Our method for the ACT is based on the work of Bhargava et al. [21], in which we generated a dictionary of Gaussian chirplets parameterized by $I = (t_c, f_c, log(\Delta_t), c)$. With smaller step sizes, we can obtain smoother and more accurate approximations, albeit at the expense of higher computational cost. As such, it is crucial to strike the right balance, particularly in a classification as time-sensitive as seizure detection.

$$g_{t_c,f_c,\log(\Delta_t),c}(t) = \frac{1}{\sqrt{\sqrt{\pi}\Delta_t}} e^{\frac{1}{2}\left(\frac{t-t_c}{\Delta_t}\right)^2} e^{i2\pi f_c(t-t_c)} \quad (1)$$

Mann and Haykin initially introduced the ACT [17] to reduce the dimensionality of the output space of the non-adaptive transform. Previous studies have also demonstrated the effectiveness and potential of the ACT for EEG analysis purposes. Cui et al. [22] initially utilized the ACT to characterize visual evoked potentials (VEP), marking the first reported instance of the ACT successfully applied to EEG signals. Bhargava et al. improved on their ACT algorithm and showed success in classifying P300 signals. Recently, Mann et al. [19] further demonstrated the effectiveness of the ACT in multi-state classification through their work with sleep-state detection. With their prior success, there is potential in extending their work into seizure state detection. One caveat is that this work will utilize different parameters and frequency ranges, as seizures often extend into higher frequency bands.

III. Methodology

A. Dataset

This paper utilized the Children's Hospital Boston-Massachusetts Institute of Technology Scalp EEG dataset [23]–[25], which was collected from 22 patients: 5 males, aged 3 to 22, and 17 females, aged 1.5 to 19. Recordings were grouped into 23 cases, with chb21 taken 1.5 years after chb01 from the same patient. The EEG recordings were sampled at 256 Hz with a 16-bit resolution, utilizing electrodes placed according to the international 10-20 system of EEG electrode positions. The ictal data were obtained directly from the samples within the labeled seizure onset and offset periods from the dataset. In contrast, interictal data were obtained from samples in records at least 2 hours before seizure onset or after seizure offset. Fig. 2 shows an ictal and interictal period taken from the dataset. We also considered the number of lead seizures per patient, with a lead seizure defined as a seizure with at least 4 hours from the previous seizure. To ensure sufficient data from each patient, we included only patients with at least 4 lead seizures and 5 hours of interictal data, ensuring an adequate amount of preictal data relative to interictal data.

Furthermore, to reduce the dimensionality of the data, we selected only four channels for analysis across all patients: FP1-F3, FP2-F4, P7-O1, and P8-O2. These channels, as chosen by Chung et al. [11], represent four common locations for wearable electrodes to be easily attached, along with similar electrode locations for portable EEG systems, such as the Muse Athena, for potential real-time implementation. In Chung et al., neurologists reviewed the seizures of all patients and determined the seizure locations and electrodes. The patients with seizures without a precise location and patients without seizures close to the chosen electrodes were excluded in their work, which our work also followed. The final patients chosen for this work are chb01, chb05, chb08, and chb10.

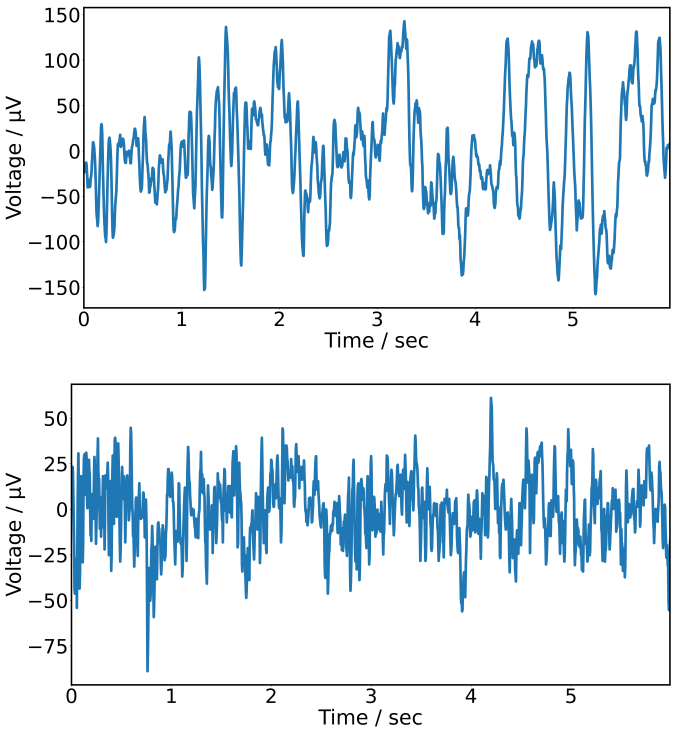

Fig. 2: EEG data of an ictal and interictal period taken from patient chb01, measured in quantigraphic units.

B. Preprocessing

The EDF files provided from the dataset were preprocessed using the MNE-Python library. A bandpass filter and a notch filter were utilized in the process.

The bandpass filter utilized a finite impulse response filter set between 1 Hz to 50 Hz, as most seizure activity can be observed within this range [16].

The MNE library utilizes a windowed time-domain design (firwin) method, employing a Hamming window with a 0.0194 passband ripple and 53 dB of stopband attenuation. The lower transition bandwidth was 0.50 Hz with a -6 dB cutoff at 0.75 Hz, and the upper transition bandwidth was 12.50 Hz with a -6 dB cutoff at 56.25 Hz.

A notch filter at 60 Hz was also applied to remove the powerline noise. A firwin method was employed, utilizing a Hamming window with 0.0194 passband ripple and 53 dB stopband attenuation. The lower passband edge was 59.35 Hz, with a lower transition bandwidth of 0.50 Hz and a -6 dB cutoff frequency at 59.10 Hz. The upper passband edge was 60.65 Hz, with an upper transition bandwidth of 0.50 Hz and a -6 dB cutoff frequency at 60.90 Hz.

C. Adaptive Chirplet Transform Features

The features used as inputs to the neural network were the parameters and coefficients, similar to Bhargava et al. and Mann et al. Furthermore, we hypothesized that the RMSE of our ACT approximation relative to the actual EEG could also serve as a valuable feature for the neural network, given the EEG's unpredictable behaviour during seizure onset. During ictal periods, seizures were more volatile and the peaks were more challenging to capture through the approximation with the range of chirp rates, compared to interictal periods. The volatility could potentially be captured by a higher RMSE, which could improve the model's accuracy. As such, we trained two different models: one with RMSE as a feature and one without, to compare the results. Each epoch was flattened to a 4x3 vector in the RMSE model and a 4x2 vector in the model without RMSE to account for the four channels.

D. Convolutional Neural Network

We utilized a 1-D Convolutional Neural Network (1-D CNN) with multi-branch inception-style blocks as the classifier, as seen in Fig. 3. The 1-D CNN layers are utilized due to their ability to identify temporal patterns within time sequences. The inputs to the CNN, the 1-second epochs, are then concatenated into 6-second sequences. File boundaries and seizure boundaries are enforced to ensure all sequences are time-consecutive. The inception-style blocks enable us to evaluate both shorter events with smaller kernels and full sequential patterns with larger kernels [26]. For the train set, due to the imbalance of the ictal to interictal samples, the ictal sequences were generated with a 3-second overlap. For the test patient, the ictal sequences have a 5-second overlap, resulting in a shorter detection latency from the onset of the seizure. The training set is downsampled across all interictal files to achieve a 5:1 interictal-to-ictal sample ratio, to reduce overfitting on interictal sequences. The test set is also downsampled across all interictal files to have a 1:1 ratio, ensuring that all ictal samples are captured while providing balanced recall and F1 scores without bias toward interictal samples. The inputs are standardized using Scikit-learn's *StandardScaler*. The evaluation was performed using LOOCV, with the model trained on three patients and tested on the remaining patient.

The inputs are initially processed using a layer normalization to standardize feature activations. Afterwards, the first inception-block concatenates four convolutional layers with 1, 2, 3, and 6 kernels each, 64 units for each layer, and the ReLU activation function, to evaluate multiple temporal patterns throughout the sequence. The second inception block concatenates two convolutional layers with 1 and 3 kernels, 64 units, and the ReLU function to extract features from the first block. Both blocks also include a dropout layer of 0.25 to reduce overfitting. Afterwards, a global average pooling layer is applied, and the result is then fed to a Dense layer with 128 units and a ReLU activation function. A dropout layer of 0.3 is used to further reduce overfitting, and then fed to a Dense layer with 1 unit and the sigmoid activation function to output the probability that the sequence is ictal. A threshold of 0.50 was utilized, meaning the final output must be 0.50 or higher to be considered ictal.

Early stopping was also implemented, given that the dataset is relatively small and prone to overfitting.

E. Postprocessing

Although typically not used for seizure detection due to the importance of short latency for onset detection, our use of a 5-second overlapping sliding window may warrant postprocessing to reduce mislabeled events, as the high overlap mitigates the effects of high detection latency. The *k-of-n* method with no overlap was utilized, with $k = 3$ and $n = 5$, meaning for five consecutive sequences, if three or more sequences were labeled ictal, the set of sequences were labeled ictal. Using overlapping sliding windows, this method evaluated five new epochs of input data, except for the first and last five sequences, which were evaluated over ten epochs in total.

IV. Results

A. Adaptive Chirplet Transform Parameters

With the high spikes of activity during ictal periods, a suitable and accurate set of parameters for the ACT dictionary and calculation was necessary. The order of the approximation was crucial to avoid overestimating the signal, and the length of each epoch was essential to account for sudden amplitude changes. Our results were compared using the RMSE value of the approximation compared to the original signal.

The parameters t_c and f_c represent the time and frequency centers, respectively, and the parameter c represents the chirp rate, all of which influence the reconstruction of the approximate signal. As the epoch size changed, the upper t_c bound was altered to account for the duration. As the bound increased from 256 to 1280, it no longer captured the rapid changes in

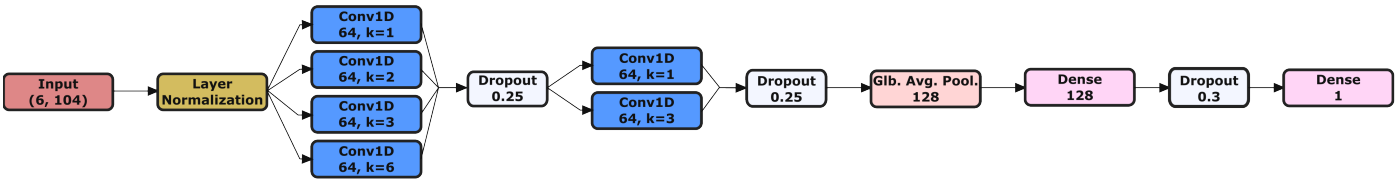

Fig. 3: CNN model utilizing Inception-style blocks with 1-D convolutional layers. This diagram is based on the model which includes RMSE as a feature.

the signal during ictal samples, so a final upper bound of 1280 was utilized.

Similarly, as the upper bound of f_c increased, it was capable of capturing more of the frequency changes in the ictal period. Seizures are often captured at frequencies up to 50 Hz, so the upper bound of f_c was set to 45 Hz to capture this range.

The chirp rate served as the most critical feature, as this parameter dictated how fast the chirplet could change frequencies. During ictal periods, frequencies could rapidly oscillate, which could only be captured with higher chirp rates. After thorough testing, the optimal balance of accuracy and computational expense was found to be a chirp rate magnitude of 10.

Testing with various orders also proved challenging, as it was difficult to balance between computational efficiency and accuracy. Overall, a fifth-order approximation yielded the best results.

Fig. 4 shows the comparison of the raw EEG and the approximation from the ACT, with our final parameters for the dictionary set as:
t_c: (0, 256, 64)
f_c: (0.5, 45, 0.5)
c: (-10, 10, 0.5)
Order = 5
Epoch = 1 second

B. Ictal State Classification

The CNN model represented in Figure 3 was tested with and without RMSE as an input feature. Table I presents the average results across all four patients for the model with RMSE as a feature and for the model without it. The two most important metrics for seizure state detection are recall and F1-score, as they provide insight into ictal sequences and the model's ability to classify them effectively relative to interictal states. Ictal states are dangerous if misclassified, and an emphasis on false negatives is crucial. Surprisingly, including RMSE as an input feature has reduced our metrics, and the model performs better without RMSE as an input; therefore, RMSE will not be included in our final comparison and for future work. Postprocessing appeared to improve the model's results marginally. Table II presents the results of the model with postprocessing applied.

Since postprocessing increases classification latency due to the larger number of windows required, real-time deployments should omit postprocessing and rely on the raw outputs of each sequence. However, our model is limited by the available data,

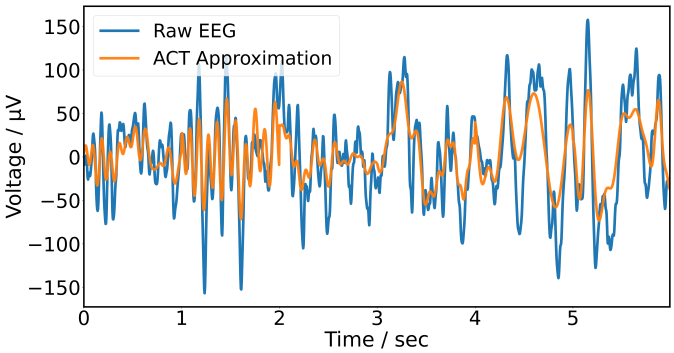

(a) Ictal period; RMSE = 21 μV

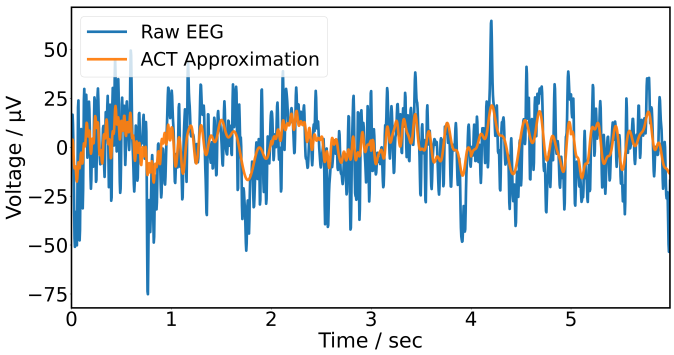

(b) Interictal period; RMSE = 14 μV

Fig. 4: Comparison of six epochs of filtered EEG data from channel FP1–F3 against the approximation from the ACT for ictal and interictal intervals, measured in quantigraphic units.

and it would be beneficial to develop a per-patient model for deployment. This is because seizures vary significantly across patients, making it challenging to identify a generalized pattern in all patients [24].

RMSE Included	AUC	Accuracy	Precision	Recall	F1
Yes	88.02%	79.14%	82.34%	84.21%	81.08%
No	88.42%	79.51%	82.42%	84.54%	81.58%

TABLE I: Model performance metrics and comparison between including and excluding RMSE as an input feature

RMSE Included	AUC	Accuracy	Precision	Recall	F1
Yes	87.89%	78.14%	81.93%	85.59%	81.22%
No	89.41%	79.81%	81.85%	88.45%	83.26%

TABLE II: Model performance metrics and comparison between including and excluding RMSE as an input feature after postprocessing

C. Power Spectral Density Comparison

A common feature previously used for ictal state detection is the power spectral density (PSD) of the Delta, Theta, Alpha, and Beta frequency bands. As such, we have modified the CNN to utilize the PSD values as inputs, further evaluating the potential improvements that the ACT may hold in seizure detection. The average for all four patients has an AUC of 0.7333, an accuracy of 66.28%, a precision of 72.91%, a recall of 68.33%, and an F1 score of 68.09%. As shown, the ACT model drastically outperforms the PSD model, with 16.21% and 13.49% improvements in recall and F1 score, respectively, indicating that the ACT model is significantly better at detecting ictal states than the PSD model.

With the complexity and unpredictable behavior of neurons during a seizure, it is evident that advanced signal processing techniques, such as the ACT, are imperative for accurately detecting seizure states.

D. Computational Runtime

For wearable options, the system must be capable of detecting seizures as early as possible to administer the necessary treatment immediately upon onset. While we have demonstrated that our model can effectively process EEG data and infer results through it successfully, it is still crucial to ensure a feasible computational runtime for potential future implementations.

The processing and inference of the ACT were performed on a NVIDIA 5070 GPU. Table III shows the average epoch processing time for each patient along with the average inference time for the model for each sequence. As we can see, our runtime for all patients is approximately 0.07 seconds, indicating that our model can effectively process and diagnose signals at a high rate. With future efforts to implement processing and inference on hardware, there is potential to rapidly and accurately diagnose seizure onsets through the ACT.

Patient	Avg. Epoch Runtime / μs	Avg. Inference Runtime / μs	Tot. Runtime / μs
chb01	71110.40	2.90	71113.30
chb05	70197.47	2.30	70199.77
chb08	70694.65	1.30	70695.95
chb10	71324.62	2.59	71327.21

TABLE III: ACT Processing and model inference runtimes across all four patients

V. Conclusion

This work shows a promising approach to accurately detecting seizure onset, with an average recall of 84.54% and an F1 score of 81.58% across all four patients, with an average computation time of 70834.06 μs. Reducing the total number of channels analyzed reduced dimensionality and computational expense, while also demonstrating potential practical implementations by utilizing common electrode locations for wearable systems. Our work has effectively shown the possible application of the ACT in seizure detection systems.

VI. Future Work

Our current results are constrained by the lack of patient-level data, as the available dataset contains limited information for each individual. A majority of patients have only ictal data recorded, with little to no interictal or preictal data. Therefore, we applied strict exclusion criteria to ensure each patient had a reasonable balance of samples from each class. For future work, we aim to collect our own data, subject to the requirements for improved data collection. Seizures also exhibit distinct patterns for each individual, so a patient-specific system would be preferable, given that more data are available for each patient and that sufficient samples are available for both ictal and interictal periods. There will be emphasis placed on patient-specific models to ensure reliable deployment.

Compared with scalp EEG data, intracranial EEG (iEEG) data often yield better seizure detection results due to the higher signal quality and improved localization of signals within the brain. Analysis and training on iEEG data rather than scalp EEG data could yield a stronger model.

Hardware implementations and evaluations of our work are also worth exploring. Our data was processed using a desktop GPU, which is impractical for a portable system. As such, we aim to evaluate the potential of a mobile, hardware-efficient system.

Lastly, we aim to extend our work into seizure prediction. As it currently stands, detection can help reduce further complications from secondary seizures, but this does not solve the issues caused by lead seizures. With a predictive model, we can intervene to prevent seizures outright or identify the necessary support at the onset. We intend to next predict the preictal state of a seizure with the ACT.

References

[1] World Health Organization. (2024) Epilepsy. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/epilepsy

[2] Epilepsy Foundation. Drug resistant epilepsy. [Online]. Available: https://www.epilepsy.com/treatment/medicines/drug-resistant-epilepsy

[3] Y. Zhang, Y. Savaria, S. Zhao, G. Mordido, M. Sawan, and F. Leduc-Primeau, "Tiny cnn for seizure prediction in wearable biomedical devices," in *2022 44th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2022, pp. 1306–1309.

[4] P. Busia, G. Leone, A. Matticola, L. Raffo, and P. Meloni, "Wearable epilepsy seizure detection on fpga with spiking neural networks," *IEEE Transactions on Biomedical Circuits and Systems*, pp. 1–11, 2025.

[5] D. Sopic, A. Aminifar, and D. Atienza, "e-glass: A wearable system for real-time detection of epileptic seizures," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2018, pp. 1–5.

[6] B. Huang, A. Abtahi, and A. Aminifar, "Lightweight machine learning for seizure detection on wearable devices," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–2.

[7] R. Zanetti, A. Aminifar, and D. Atienza, "Robust epileptic seizure detection on wearable systems with reduced false-alarm rate," in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, 2020, pp. 4248–4251.

[8] A. Shoeb and J. Guttag, "Application of machine learning to epileptic seizure detection," in *ICML'10: Proceedings of the 27th International Conference on International Conference on Machine Learning*, 2010, pp. 975–982.

[9] Z. Wang and P. Mengoni, "Seizure classification with selected frequency bands and eeg montages: a natural language processing approach," *Brain Inf*, vol. 9, no. 11, 2022.

[10] A. Emami, N. Kunii, T. Matsuo, T. Shinozaki, K. Kawai, and H. Takahashi, "Seizure detection by convolutional neural network-based analysis of scalp electroencephalography plot images," *NeuroImage: Clinical*, vol. 22, p. 101684, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2213158219300348

[11] Y. Chung, A. Cho, H. Kim, and K. Kim, "Single-channel seizure detection with clinical confirmation of seizure locations using chb-mit dataset," *Frontiers in Neurology*, vol. 15, 2024.

[12] P. Fergus, D. Hignett, A. Hussain, D. Al-Jumeily, and K. Abdel-Aziz, "Automatic epileptic seizure detection using scalp eeg and advanced artificial intelligence techniques," *BioMed Research International*, vol. 2015, no. 1, p. 986736, 2015. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1155/2015/986736

[13] J. Prasanna, M. S. P. Subathra, M. A. Mohammed, R. Damaševičius, N. J. Sairamya, and S. T. George, "Automated epileptic seizure detection in pediatric subjects of chb-mit eeg database—a survey," *Journal of Personalized Medicine*, vol. 11, no. 10, 2021. [Online]. Available: https://www.mdpi.com/2075-4426/11/10/1028

[14] X. Cao, S. Zheng, J. Zhang, W. Chen, and G. Du, "A hybrid cnn-bi-lstm model with feature fusion for accurate epilepsy seizure detection," *BMC Med Inform Decis Mak*, vol. 25, no. 6, 2025.

[15] N. Kumar and S. Mann, "Gpu-accelerated chirplet transform: Scalable runtime profiling and analysis," to appear in *Proc. IEEE Int. Conf. on Sensing Technology (ICST)*, 2025.

[16] S. Grewal and J. Gotman, "An automatic warning system for epileptic seizures recorded on intracerebral eegs," *Clinical Neurophysiology*, vol. 116, no. 10, pp. 2460–2472, Oct. 2005.

[17] S. Mann and S. Haykin, "Adaptive chirplet: an adaptive generalized wavelet-like transform," in *SPIE, 36th Annual International Symposium on Optical and Optoelectronic Applied Science and Engineering*, vol. 31, no. 6, Jun. 1991, pp. 1243–1256.

[18] ——, "The chirplet transform: physical considerations," *IEEE Transactions on Signal Processing*, vol. 43, no. 11, pp. 2745–2761, 1995.

[19] S. Mann, N. Kumar, J. P. Bicalho, M. Sibai, and C. Leaver-Preyra, "Adaptive chirplet transform-based sleep state detection," in *2025 IEEE International Conference on Consumer Electronics (ICCE)*, 2025, pp. 1–6.

[20] S. Mann and S. Haykin, "The chirplet transform: A generalization of gabor's logon transform," *Vision Interface*, vol. 91, pp. 205–212, 1991.

[21] A. Bhargava and S. Mann, "Adaptive chirplet transform-based machine learning for p300 brainwave classification," in *2020 IEEE-EMBS Conference on Biomedical Engineering and Sciences (IECBES)*, 2021, pp. 62–67.

[22] J. Cui and W. Wong, "The adaptive chirplet transform and visual evoked potentials," *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 7, pp. 1378–1384, 2006.

[23] J. Guttag, "Chb-mit scalp eeg database (version 1.0.0)," PhysioNet, 2010, rRID: SCR_007345.

[24] A. Shoeb, "Application of machine learning to epileptic seizure onset detection and treatment," Ph.D. thesis, Harvard University–MIT Division of Health Sciences and Technology, MIT, Cambridge, MA, 2009.

[25] A. Goldberger, L. Amaral, L. Glass, J. Hausdorff, P. C. Ivanov, R. Mark, J. E. Mietus, G. B. Moody, C. 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, rRID: SCR_007345. [Online]. Available: https://physionet.org

[26] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1–9.