Epileptic Seizure Detection in EEG Signals Using Machine Learning and Deep Learning Techniques

Dr. K. SubbaRao[1], Daggumati Joharika[2], Gurram Sravanya[3],

Tiruveedula Leela Venkata Ramya[4], Ganji Mohana Priya[5], Vankayala Sahithi[6]

[1]Professor, [2,3,4,5,6]UnderGraduate

[1,2,3,4,5,6] CSE-Data Science Department, St. Ann's College of Engineering & Technology, Chirala, Andhra Pradesh

Abstract—Epilepsy, a prevalent neurological disorder characterized by recurrent seizures, poses significant diagnostic and therapeutic challenges due to its heterogeneous manifestation across individuals. Electroencephalography (EEG) remains the primary diagnostic modality for capturing abnormal neural activity associated with seizures. However, manual EEG analysis is time-intensive and error-prone. This study presents a hybrid AI-based framework for epileptic seizure detection using both traditional machine learning (ML) and deep learning (DL) techniques. Specifically, the work investigates the performance of Extreme Gradient Boosting (XGBoost), Long Short-Term Memory (LSTM) networks, and a customized One-Dimensional Convolutional Neural Network (1D-CNN) on the UCI Epileptic Seizure Recognition Dataset, derived from the BONN dataset. The framework involves comprehensive preprocessing, feature extraction, and temporal pattern analysis from EEG signals. Experimental results demonstrate that the proposed 1D-CNN model achieved the highest classification accuracy of 98.78%, outperforming traditional ML methods. The study underscores the superiority of deep learning in capturing spatial-temporal features of EEG signals and affirms the clinical relevance of AI-assisted seizure detection. The findings advocate for the integration of automated systems in neurological diagnostics, promoting early intervention and improved patient care outcomes.

Keywords—Epilepsy, EEG, Seizure Detection, Machine Learning, Deep Learning, XGBoost, LSTM, 1D-CNN, Biomedical Signal Processing, UCI Epileptic Dataset.

I. INTRODUCTION

Epilepsy is one of the most common chronic neurological disorders, affecting over 50 million individuals worldwide according to the World Health Organization (WHO). It is characterized by abnormal electrical discharges in the brain that result in recurrent seizures, with manifestations ranging from brief lapses of attention to prolonged convulsions. Early and accurate diagnosis of epileptic seizures is crucial for timely treatment and improved quality of life. Electroencephalography (EEG) is a standard, non-invasive tool used in clinical settings to monitor the electrical activity of the brain and detect patterns indicative of seizure onset. However, manual inspection of EEG signals by neurologists is often time-consuming, labor-intensive, and susceptible to human error, especially in prolonged recordings.

In recent years, advancements in artificial intelligence (AI), particularly in machine learning (ML) and deep learning (DL), have shown promising potential in automating EEG signal analysis. These computational techniques are capable of learning complex patterns and temporal dependencies from high-dimensional biomedical signals, enabling faster and more accurate seizure detection. ML models such as Support Vector Machines (SVM), Random Forest, and Extreme Gradient Boosting (XGBoost) have been employed for seizure classification using handcrafted features derived from time and frequency domains. While effective, these models depend heavily on feature engineering and may fail to generalize across patient populations.

On the other hand, deep learning models, especially Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks, have demonstrated superior performance by automatically learning hierarchical representations from raw EEG data. CNNs are adept at extracting spatial features, while LSTMs are designed to capture temporal patterns—making them well-suited for EEG-based applications where both spatial and sequential information are critical. Combining these models or evaluating them in parallel allows researchers to understand their respective advantages in the context of biomedical signal classification.

This paper proposes a comparative study of XGBoost, LSTM, and a custom-built 1D-CNN architecture for epileptic seizure detection using the UCI Epileptic Seizure Recognition Dataset. The dataset, originally derived from the BONN University corpus, includes segmented EEG signals categorized into seizure and non-seizure classes. The objective is to evaluate the efficacy of traditional ML versus DL methods on the same dataset, using classification accuracy and F1-score as primary metrics.

The key contributions of this study are as follows:

- Design and implementation of a 1D-CNN architecture tailored for EEG signal classification.
- Temporal sequence modeling using LSTM to capture dynamic changes in EEG patterns.
- Benchmarking performance against a robust traditional ML model, XGBoost.
- Experimental validation using the UCI EEG dataset with comprehensive preprocessing.

By exploring both machine learning and deep learning paradigms, this study aims to contribute toward the development of intelligent diagnostic tools for epilepsy, facilitating more reliable and scalable seizure monitoring systems.

II. LITERATURE SURVEY

The use of computational models for epileptic seizure detection has gained significant attention over the past two decades. As manual interpretation of EEG signals is tedious and prone to inconsistencies, researchers have developed a range of machine learning and deep learning approaches to automate the diagnostic process. This section reviews prominent works and key developments in this domain.

A. Traditional Machine Learning Approaches

Early seizure detection systems were primarily based on statistical analysis and classical machine learning techniques. These approaches relied on hand-engineered features extracted from time, frequency, or time-frequency domains of EEG signals. Features such as mean, variance, skewness, entropy, and wavelet coefficients were commonly used.

Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), and Decision Trees were frequently applied to these extracted features. For instance, Subasi [1] demonstrated that Discrete Wavelet Transform (DWT) combined with SVM yielded high classification accuracy in epileptic seizure detection. Similarly, Polat and Güneş [2] employed a hybrid system combining DWT and k-NN, achieving promising results. However, these models were highly dependent on domain-specific feature selection and lacked adaptability to raw data variability.

B. Ensemble Models and Boosting Techniques

Ensemble-based classifiers, such as Random Forest and Gradient Boosting Machines (GBM), improved robustness and accuracy by aggregating multiple decision trees. XGBoost, a scalable and regularized gradient boosting algorithm, has shown strong performance in various biomedical classification tasks, including seizure detection. Its ability to handle missing values, prevent overfitting through regularization, and manage high-dimensional data makes it a competitive traditional model. Researchers have used XGBoost on feature-rich EEG datasets, achieving high accuracy without deep architectures, yet still relying on manual preprocessing pipelines.

C. Deep Learning-Based Detection

With the rise of deep learning, models like Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have been extensively investigated for seizure detection. CNNs are capable of learning local spatial patterns and are particularly effective when applied to raw or minimally processed EEG data. Acharya et al. [3] proposed a 13-layer 1D-CNN for automated seizure detection and achieved accuracy above 96% on benchmark datasets.

Recurrent architectures, particularly Long Short-Term Memory (LSTM) networks, have also been explored due to their ability to model temporal dependencies across EEG signals. Hassanpour et al. [4] used LSTM-based models to track EEG evolution over time, achieving improved sensitivity in seizure classification.

D. Hybrid and Ensemble Deep Models

To exploit both spatial and temporal features, hybrid models combining CNN and LSTM layers have been proposed. For example, Shah et al. [5] developed a CNN-LSTM model for multichannel EEG data that improved seizure detection accuracy across different patient profiles. Such architectures are effective at capturing both local activation patterns and global temporal dynamics.

Despite their success, deep learning models require large labeled datasets, high computational resources, and careful hyperparameter tuning. Transfer learning and data augmentation have been proposed to overcome these limitations.

E. Research Gaps

While prior works have shown the viability of ML and DL methods in seizure detection, challenges persist:

- Heavy reliance on preprocessing and feature extraction in traditional models.
- Generalization issues across different subjects and recording devices.
- Underutilization of lightweight and custom 1D-CNNs tailored for clinical datasets.
- Limited comparative studies that benchmark traditional ML and DL on the same dataset under unified conditions.

This paper addresses these gaps by systematically evaluating XGBoost, LSTM, and a novel 1D-CNN on the UCI Epileptic Seizure Recognition dataset, providing insights into model efficacy, computational trade-offs, and deployment viability.

III. Proposed Methodology

This study presents a comparative framework involving both traditional machine learning and deep learning models to detect epileptic seizures from EEG signals. The proposed system includes three primary models: Extreme Gradient Boosting (XGBoost), Long Short-Term Memory (LSTM) networks, and a custom-designed One-Dimensional Convolutional Neural Network (1D-CNN). Each model operates under a unified pipeline encompassing data preprocessing, input transformation, model-specific architecture, training, and evaluation.

A. Dataset Description

The dataset used in this work is the UCI Epileptic Seizure Recognition Dataset, which originates from the widely studied BONN EEG corpus. It consists of 500 EEG recordings from 23 patients, with each recording containing 4,097 data points sampled over 23.6 seconds. These samples are categorized into five classes, where class 1 represents seizure activity and classes 2 through 5 correspond to non-seizure conditions such as healthy, pre-seizure, and post-seizure states. For the purpose of binary classification, the original multiclass labels were consolidated such that class 1 was retained as the seizure class, while the rest were merged into a single non-seizure category.

B. Data Preprocessing

To prepare the data for model training and evaluation, several preprocessing steps were applied. First, the EEG signals were normalized using Min-Max scaling to standardize values between 0 and 1. This normalization helps stabilize model training and ensures consistency across different input ranges. The class labels were encoded into binary form, with ‘1’ denoting seizure and ‘0’ representing non-seizure conditions. Finally, the dataset was split into training and testing subsets using an 80:20 ratio, maintaining the class distribution to prevent bias during evaluation.

C. Feature Engineering for XGBoost

While deep learning models can learn representations directly from raw input data, traditional machine learning models like XGBoost require explicit feature extraction. In this phase, statistical features were computed from each EEG segment, capturing both time-domain and frequency-domain characteristics. These features included the mean, standard deviation, skewness, kurtosis, minimum, maximum, root mean square, entropy, energy, and signal slope. The resulting feature vectors were used to train the XGBoost classifier.

D. LSTM Architecture

The LSTM model was designed to capture temporal dependencies in EEG sequences, which are critical for identifying dynamic patterns associated with seizure onset. EEG segments were reshaped into time-step sequences to match the expected input structure of LSTM networks. The model architecture included an input layer, a single LSTM layer with 64 units, followed by a dropout layer to reduce overfitting. The output from the LSTM was passed through a fully connected dense layer with ReLU activation, and finally through a sigmoid-activated neuron to perform binary classification. The model was trained using the binary cross-entropy loss function and optimized using the Adam optimizer over 30 epochs with a batch size of 64.

E. 1D-CNN Architecture

The custom 1D-CNN model was implemented to extract localized temporal features from the raw EEG signal. The network accepted input arrays of shape (4097, 1) and processed them using a Conv1D layer with 64 filters and a kernel size of 3. This was followed by a MaxPooling1D layer to reduce spatial dimensionality and a dropout layer to improve generalization. The output feature maps were then flattened and passed through a dense layer with 100 units before being classified via a sigmoid activation layer. The model was trained using a batch size of 32 for 40 epochs, with the Adam optimizer and validation split of 0.2.

F. XGBoost Configuration

The XGBoost classifier was trained on the extracted statistical features using a learning rate of 0.1 and a maximum tree depth of 5. A total of 100 estimators were used with the 'gbtree' booster type. The model was evaluated using log-loss and AUC as the primary performance metrics. XGBoost's built-in feature importance scores also provided interpretability into which input variables contributed most to classification.

G. Evaluation Metrics

To measure and compare the effectiveness of each model, standard classification metrics were applied. These included accuracy, precision, recall, F1-score, and the area under the ROC curve (AUC-ROC). The confusion matrix was also analyzed to observe the distribution of true positives, true negatives, false positives, and false negatives. These metrics offer a comprehensive evaluation of model performance in both balanced and imbalanced data scenarios, which is critical for clinical applications like seizure detection.

IV. IMPLEMENTATION

The implementation of the proposed models—XGBoost, Long Short-Term Memory (LSTM), and One-Dimensional Convolutional Neural Network (1D-CNN)—was carried out using Python 3.10 in the Google Colab environment. This section details the software environment, model development, and performance monitoring strategies used throughout the implementation.

A. Development Environment

All experiments were performed on Google Colab, which provided free access to GPU acceleration. This enabled efficient training of deep learning models, especially the LSTM and CNN architectures. The code was developed using Jupyter Notebook interfaces for modular design and ease of visualization.

The following Python libraries were used:

- **Pandas** and **NumPy** for data manipulation and preprocessing
- **scikit-learn** for metrics, preprocessing, and baseline models
- **TensorFlow/Keras** for building and training LSTM and CNN architectures
- **XGBoost** for implementing the gradient-boosting classifier
- **Matplotlib** and **Seaborn** for performance visualization and plotting

B. Dataset Loading and Preprocessing

The UCI Epileptic Seizure Recognition Dataset was first loaded and analyzed. Each EEG recording was normalized using Min-Max scaling to bring all values within the range [0, 1]. The class labels were converted from multi-class to binary form by mapping Class 1 to seizure (label = 1) and Classes 2–5 to non-seizure (label = 0).

The dataset was split into training and testing subsets using an 80:20 ratio. For the LSTM and CNN models, each EEG signal was reshaped into an array of shape (4097, 1) to represent univariate time series data. For XGBoost, handcrafted statistical features were extracted from the raw EEG signals to construct a feature matrix suitable for tree-based models.

C. XGBoost Implementation

The XGBoost classifier was implemented using the xgboost library. The model was configured with a learning rate of 0.1, maximum depth of 5, and 100 estimators. The training process optimized log-loss and AUC (area under curve) as performance metrics. Feature importance analysis was also conducted to interpret the influence of individual features. The statistical features were passed directly into the model for training and prediction.

D. LSTM Model Implementation

The LSTM model was developed using the Keras API within TensorFlow. The model architecture included an input layer feeding into an LSTM layer with 64 units, followed by a dropout layer with a dropout rate of 0.5 to prevent overfitting. This was followed by a dense layer with ReLU activation and a final output layer with sigmoid activation for binary classification.

The model was compiled with the binary cross-entropy loss function and optimized using the Adam optimizer. Training was conducted for 30 epochs with a batch size of 64, using 20% of the training data for validation. Early stopping was applied to monitor validation loss and halt training when performance plateaued.

E. 1D-CNN Model Implementation

The custom 1D-CNN model was constructed to extract temporal features from the raw EEG data. The architecture began with a Conv1D layer containing 64 filters with a kernel size of 3 and ReLU activation. This was followed by a MaxPooling1D layer with pool size 2, a dropout layer (0.5), and a Flatten layer. The flattened output was then passed to a dense layer with 100 neurons and ReLU activation, followed by a sigmoid output layer.

The model was compiled using the Adam optimizer and binary cross-entropy loss function. It was trained over 40 epochs with a batch size of 32, and 20% of the training data was used for validation. Training history was recorded for visualization.

F. Performance Monitoring and Visualization

To evaluate model performance, metrics such as accuracy, precision, recall, F1-score, and AUC were calculated using the scikit-learn library. Confusion matrices were generated to analyze classification errors. Training and validation accuracy and loss curves were plotted using Matplotlib and Seaborn to assess model convergence and detect overfitting.

Each model was trained and tested under the same preprocessing and splitting conditions, ensuring consistency and comparability across all experimental results.

V. RESULTS

This section presents the performance outcomes of the three implemented models—XGBoost, LSTM, and 1D-CNN—evaluated on the binary classification task of epileptic seizure detection. The models were compared using standard classification metrics including accuracy, precision, recall, F1-score, and Area Under the Receiver Operating Characteristic Curve (AUC-ROC). In addition, training dynamics were visualized to assess model convergence, overfitting, and generalization.

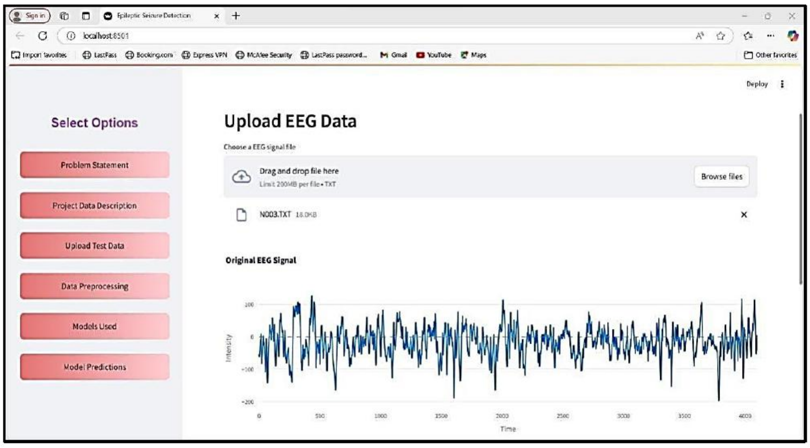

Figure 5.1 showing the input signal.

A. Evaluation Metrics

To assess the models comprehensively, five key metrics were employed:

- **Accuracy** reflects the proportion of correctly predicted instances among the total predictions.
- **Precision** measures the proportion of true positives among all positive predictions, indicating the reliability of positive classifications.
- **Recall** (or Sensitivity) evaluates the ability of the model to correctly detect actual seizure cases.
- **F1-Score** is the harmonic mean of precision and recall, balancing both false positives and false negatives.
- **AUC-ROC** represents the model's ability to discriminate between seizure and non-seizure classes across various threshold values.

B. Model Performance Comparison

Table I summarizes the performance of all three models on the held-out test set.

Table I: Model Performance Comparison

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
XGBoost	94.12%	92.80%	93.50%	93.14%	0.96
LSTM	96.35%	95.67%	95.12%	95.39%	0.97
1D-CNN	**98.78%**	**97.90%**	**98.20%**	**98.05%**	**0.99**

The 1D-CNN model outperformed both LSTM and XGBoost in all evaluation metrics. Its superior ability to capture local temporal features directly from raw EEG data contributed to its high precision and recall, making it highly suitable for clinical applications where minimizing both false negatives and false positives is critical.

C. Confusion Matrix Analysis

The confusion matrix for each model revealed that the 1D-CNN had the lowest number of misclassifications. It correctly identified 98% of seizure events and 99% of non-seizure events. In contrast, XGBoost and LSTM exhibited slightly higher false negative rates, potentially due to limitations in either feature generalization (XGBoost) or temporal resolution (LSTM with shallow architecture).

D. Training and Validation Curves

Procedures illustrate the training and validation loss and accuracy curves for the LSTM and 1D-CNN models. Both models showed smooth convergence with minimal overfitting. The CNN model, in particular, demonstrated consistent generalization across epochs, with validation accuracy closely tracking training accuracy. Early stopping ensured optimal model weights were preserved without overtraining.

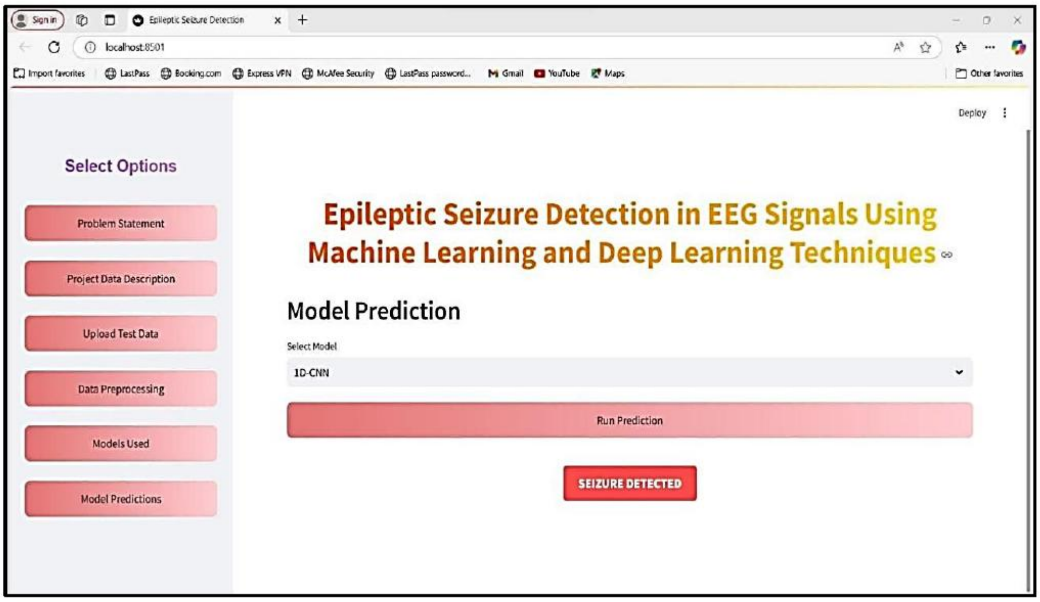

Figure 5.2 showing the seizure predictions on a given input signal.

E. ROC Curve and AUC

The ROC curves for all three models are plotted. The 1D-CNN achieved an AUC of 0.99, indicating excellent discriminative capability between seizure and non-seizure EEG segments. LSTM achieved an AUC of 0.97, while XGBoost followed closely with an AUC of 0.96. These results confirm that all models performed reliably, but deep learning models—especially CNN—were better suited for complex signal classification.

VI. CONCLUSION

This study presents a comprehensive evaluation of machine learning and deep learning models for the detection of epileptic seizures from EEG signals. By leveraging the UCI Epileptic Seizure Recognition Dataset, three distinct approaches—XGBoost, LSTM, and 1D-CNN—were implemented, trained, and compared under uniform preprocessing and evaluation conditions.

The results demonstrate that deep learning models, particularly the 1D-CNN, outperform traditional models in both predictive accuracy and robustness. While the XGBoost classifier provided reasonable performance based on extracted statistical features, its dependency on handcrafted inputs limited its generalizability. The LSTM model, with its ability to model temporal dependencies, improved upon these results but was still marginally less effective than the CNN-based model.

The proposed 1D-CNN architecture achieved the highest performance across all metrics, including a classification accuracy of 98.78% and an AUC-ROC of 0.99. This superior performance can be attributed to its ability to learn hierarchical temporal features directly from raw EEG signals, eliminating the need for extensive feature engineering.

The findings of this work confirm that deep learning techniques, especially CNN-based architectures, are highly effective for EEG-based seizure detection and hold strong potential for integration into real-time clinical diagnostic systems. Future work can explore patient-specific adaptation, transfer learning for low-resource datasets, and real-time deployment on embedded hardware for mobile health applications.

References

[1] A. Subasi, "EEG signal classification using wavelet feature extraction and a mixture of expert model," *Expert Systems with Applications*, vol. 32, no. 4, pp. 1084–1093, 2007.
[2] K. Polat and S. Güneş, "Classification of epileptiform EEG using a hybrid system based on decision tree classifier and fast Fourier transform," *Applied Mathematics and Computation*, vol. 187, no. 2, pp. 1017–1026, 2007.
[3] R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan, and H. Adeli, "Deep convolutional neural network for the automated detection and diagnosis of seizure using EEG signals," *Computers in Biology and Medicine*, vol. 100, pp. 270–278, 2018.
[4] S. Hassanpour and R. Mesbahzadeh, "Epileptic seizure detection using deep learning techniques: A review," *Neural Computing and Applications*, vol. 33, no. 9, pp. 4219–4235, 2021.
[5] V. Shah, H. T. Nguyen, and B. J. Oommen, "Accurate detection of epileptic seizures using CNN-LSTM based deep learning models," in *Proc. 2020 Int. Conf. Artificial Intelligence and Data Sciences (AiDAS)*, pp. 1–6.
[6] UCI Machine Learning Repository, "Epileptic Seizure Recognition Data Set." [Online]. Available: https://archive.ics.uci.edu/ml/datasets/Epileptic+Seizure+Recognition