

Deep Learning Enabled Automatic Abnormal Eeg Identification Using Custom Recurrent Neural Network

Malathi P ^{#1}, Jothika P ^{*2}, Vennila R ^{#3}, V. Vijayakumari ^{#4}

*# Department Of Artificial Intelligence And Data Science, The Kavery Engineering College, Salem, Tamil Nadu, India
(Affiliated With Anna University)*

malathip28061969@gmail.com , ¹, *kajothi06@gmail.com* , ², *kaviyavennila3003@gmail.com* , ³

Abstract

Routine interpretation of electroencephalograms (EEGs) is labor-intensive, subjective, and difficult to scale in clinical practice. This study proposes a custom convolutional recurrent neural network (CRNN) for automatic detection of abnormal EEGs from minimally processed multi-channel recordings. The architecture integrates 1-D dilated convolutions, lightweight bidirectional recurrent units, and channel-wise attention to model local waveform irregularities and long-range temporal dependencies. Evaluated on the TUH Abnormal EEG Corpus, the method attained 92.0% accuracy, 89.0% sensitivity, and 0.95 AUC-ROC, while sustaining inference below 100 ms per 10-s segment.

Keywords—Abnormal EEG detection, CRNN, deep learning, TUH EEG Corpus, clinical decision support, recurrent neural network.

I. INTRODUCTION

Electroencephalography is one of the most widely used non-invasive modalities for assessing neurological disorders, especially epilepsy, encephalopathy, and diffuse cerebral dysfunction [1]. Because EEG offers excellent temporal resolution, it is well suited for identifying short-lived abnormalities that may not be visible in structural neuroimaging. Despite its clinical value, routine EEG assessment still relies largely on expert neurologists. As a result, interpretation remains time-consuming, expensive, and vulnerable to inter-observer variation. Earlier clinical investigations have shown that even under standardized review settings, visual EEG interpretation may differ across clinicians, thereby affecting diagnostic consistency and confidence [3].

The rapid growth of digital EEG repositories has increased the demand for reliable automated analysis. Conventional machine learning methods typically depend on handcrafted descriptors, including spectral features, wavelet coefficients, entropy measures, and statistical summaries. Although such approaches can perform reasonably well in restricted settings, they require substantial domain expertise and often struggle to represent the nonlinear and non-stationary characteristics of EEG signals. In contrast, recent deep learning models have reduced dependence on manual feature design by learning directly from time-series recordings or transformed signal representations [2], [4].

Even so, standalone convolutional neural networks (CNNs) tend to prioritize local waveform morphology over longer temporal structure, whereas standard recurrent neural networks (RNNs) are susceptible to overfitting when applied to long EEG sequences or diverse clinical recordings.

To address these limitations, this work introduces a custom convolutional recurrent neural network (CRNN) for automatic abnormal EEG identification. The proposed framework combines dilated one-dimensional convolutional layers for local temporal feature extraction with a lightweight bidirectional recurrent module for sequential modeling. In addition, a channel-wise attention mechanism is incorporated to emphasize diagnostically informative electrodes while reducing the influence of noisy or weakly relevant channels. The model is designed for end-to-end learning from minimally preprocessed EEG segments, thereby reducing reliance on handcrafted features and supporting practical clinical deployment.

This study makes three principal contributions. First, it presents a compact CRNN architecture that captures both local waveform characteristics and longer temporal dependencies without excessive parameter growth. Second, it develops a complete end-to-end pipeline for the TUH Abnormal EEG Corpus using standardized multi-channel inputs, artifact-aware preprocessing, augmentation, and regularized training. Third, it evaluates

the system in terms of both diagnostic effectiveness and deployment suitability, demonstrating strong classification performance, low-latency inference, and an interpretable output layer capable of generating abnormality heatmaps for clinical support.

II. LITERATURE SURVEY

EEG interpretation has long been central to epilepsy diagnosis and syndrome classification, yet its limitations in sensitivity and reproducibility have also been widely acknowledged. Smith [1] noted that routine scalp EEG remains clinically indispensable, but its utility is constrained by limited sampling, restricted sensitivity, and the possibility of misinterpretation when considered in isolation. Azuma *et al.* [3] further showed that structured interventions can improve interrater agreement, highlighting the continuing challenge of variability in expert review.

Research on automated abnormal EEG identification advanced considerably with the availability of large clinical datasets. Lopez *et al.* [2] examined early machine learning and deep learning methods for abnormal adult EEG classification and showed that learned representations can outperform manually engineered pipelines. This direction was extended in Lopez’s later thesis [4], where deeper 1-D CNN-RNN architectures and data augmentation were applied to the TUH corpus. However, the recurrent components were still prone to overfitting, particularly when longer recordings were used. A major

catalyst for this area of work was the TUH EEG Corpus introduced by Obeid and Picone [5], which provided a large and heterogeneous clinical dataset suitable for benchmarking realistic EEG classification systems.

Developments in sequence modeling from adjacent domains have also informed EEG analysis. Kiral-Kornek *et al.* [6] demonstrated that deep temporal architectures can support clinically meaningful seizure prediction under real-time operating conditions. Wang and Oates [7] reported that transforming time series into image-based representations may improve classification performance in certain settings, although such preprocessing introduces additional complexity and may obscure native signal characteristics. Foundational advances in optimization and sequence learning, including Adam optimization [8], GRU-based encoder-decoder models [9], and LSTM-based sequence modeling [10], have significantly shaped modern biomedical signal processing architectures.

Despite these developments, several gaps remain. First, many existing methods continue to rely heavily on handcrafted features or computationally intensive signal transformations. Second, architectures with strong temporal modeling capacity often incur large parameter counts and an increased risk of overfitting. Third, relatively few studies achieve a satisfactory balance among clinical robustness, computational efficiency, and interpretability. These limitations motivate the CRNN framework proposed in this paper.

TABLE I: LITERATURE COMPARISON

Ref.	Method / Focus	Data / Scope	Major Strength	Key Limitation
[1]	Clinical EEG review for epilepsy	Routine scalp EEG	Strong diagnostic relevance	Low sensitivity; sampling limitations
[3]	Interrater reliability improvement	Clinical EEG interpretation	Improves consistency	Still expert-dependent
[2]	Automated abnormal EEG detection	Adult EEG; early DL/ML	Reduced manual effort	Limited temporal depth
[4]	1-D CNN-RNN on TUH	TUH Abnormal EEG Corpus	Better feature learning; augmentation	Recurrent overfitting on long sequences
[5]	TUH EEG Corpus release	Large clinical EEG repository	Enables scalable benchmarking	Heterogeneous recordings require harmonization
[6]	Deep temporal seizure prediction	Long-term EEG	Real-time feasibility	Mainly patient-specific setting
[7]	Time-series imaging + CNN	General time-series classification	Effective learned image features	Added preprocessing overhead

As summarized in Table I, prior studies clearly establish both the importance of automation and the shortcomings of current architectures. The proposed CRNN is intended to

bridge the gap between shallow feature-based systems and computationally heavy hybrid models by improving

generalization while preserving temporal modeling capability.

III. METHODOLOGY / APPROACH

A. Dataset and Preprocessing

The experiments were carried out using the TUH Abnormal EEG Corpus, a subset of the Temple University Hospital EEG repository [5]. This corpus contains both normal and abnormal clinical EEG recordings collected under real hospital conditions. To ensure consistency across samples, recordings were mapped to a common 19-channel montage whenever feasible, and sessions with severe channel mismatches or corrupted metadata were removed during quality assurance. After harmonization, a subset comprising more than 2,000 recordings was retained for model development and evaluation.

Most recordings were acquired at a sampling rate of 250 Hz. Each EEG session was divided into overlapping 10-s windows to expand the effective training set and provide more temporal context than protocols limited to the first minute of recording. Segment-level predictions were subsequently aggregated to produce the recording-level label. Preprocessing was intentionally kept minimal yet clinically meaningful. A 0.5-40 Hz bandpass filter was applied to suppress baseline drift and high-frequency noise, independent component analysis was used to reduce common ocular and muscle artifacts, and per-channel z-score normalization was employed to stabilize the input distribution. For a channel signal x_c , normalization was defined as

$$\hat{x}_c = \frac{x_c - \mu_c}{\sigma_c + \epsilon},$$

where μ_c and σ_c represent the mean and standard deviation of channel c , respectively.

To enhance robustness, training-time augmentation included random time shifting, amplitude perturbation, and low-level Gaussian noise injection. These augmentations were deliberately restrained so that pathological characteristics remained intact while variability across recording conditions was better represented.

The overall processing workflow is illustrated in Fig. 1.

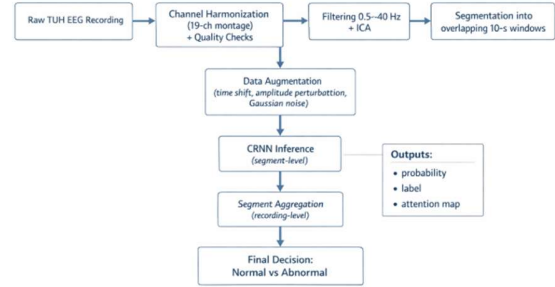


Fig. 1. End-to-end workflow of the proposed abnormal EEG identification system.

B. Proposed CRNN Architecture

The proposed CRNN operates on a multi-channel EEG segment $X \in \mathbb{R}^{C \times T}$, where $C = 19$ and T denotes the number of time samples. The first stage of the network comprises stacked one-dimensional dilated convolution blocks. These layers learn short-duration waveform signatures, including spikes, sharp waves, rhythmic discharges, and transient bursts, while progressively expanding the receptive field. For the l -th convolution block, the feature map is expressed as

$$H^{(l)} = \phi \left(W^{(l)} *_{d_l} H^{(l-1)} + b^{(l)} \right),$$

where $*_{d_l}$ denotes convolution with dilation factor d_l , $W^{(l)}$ and $b^{(l)}$ are learnable parameters, and $\phi(\cdot)$ denotes batch normalization followed by ReLU activation. Max-pooling is introduced selectively to reduce temporal resolution and suppress redundant activations.

Following convolutional encoding, the feature maps are reorganized as a temporal sequence and passed to a lightweight bidirectional gated recurrent module inspired by gated sequence modeling [9], [10]. For each time step t , the recurrent update is defined as

$$\begin{aligned} z_t &= \sigma(W_z x_t + U_z h_{t-1}), r_t = \sigma(W_r x_t + U_r h_{t-1}), \\ \tilde{h}_t &= \tanh(W_h x_t + U_h(r_t \odot h_{t-1})), \\ h_t &= (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t, \end{aligned}$$

where z_t and r_t denote the update and reset gates, respectively. Bidirectional processing allows the model to use contextual information from both preceding and succeeding states within a segment, which is beneficial for identifying diffuse or evolving abnormalities.

A channel-wise attention module is then applied to highlight informative electrodes. Let s_c denote a summary

statistic for channel c . The corresponding attention weight α_c is computed as

$$\alpha_c = \text{softmax}(W_a s_c + b_a),$$

and the reweighted representation becomes $\tilde{H}_c = \alpha_c H_c$. This mechanism enhances interpretability and directs the model toward clinically meaningful channel groups.

Dense skip connections between recurrent blocks facilitate feature reuse and improve gradient propagation. Dropout in the range of 0.3-0.5, together with L2 regularization, is used to mitigate overfitting. The final prediction is obtained through a fully connected layer followed by sigmoid activation:

$$\hat{y} = \sigma(W_o z + b_o),$$

where \hat{y} denotes the probability that the EEG segment is abnormal.

The architecture is summarized in Fig. 1.




Fig. 1. Architecture of the proposed CRNN-based abnormal EEG classification framework.

C. Training Strategy and Experimental Setup

The model was implemented in Python using PyTorch, while MNE-Python was employed for EEG loading and preprocessing. Training was performed on an Intel i7 workstation equipped with 16-32 GB RAM and an NVIDIA RTX-class GPU. Optimization was carried out with Adam [8] using an initial learning rate of 10^{-3} , mini-batch training, early stopping, and five-fold cross-validation. The loss function was binary cross-entropy with an L2 penalty term:

$$\mathcal{L} = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})] + \lambda \| \theta \|_2^2.$$

Performance was assessed using accuracy, sensitivity, specificity, F1-score, AUC-ROC, and inference time per 10-s segment. Baseline comparisons were conducted against a CNN-only model, a recurrent-only model, and a conventional CNN-BiGRU hybrid.

D. Prototype Deployment Interface

To facilitate translational use, a lightweight clinical dashboard was developed through a web-based interface for EEG upload, inference execution, and report visualization. The home page supports secure file submission and recording selection, whereas the results page presents the predicted label, confidence score, and attention-based heatmap. Such a prototype is useful for evaluating how the model may fit into practical hospital workflows.

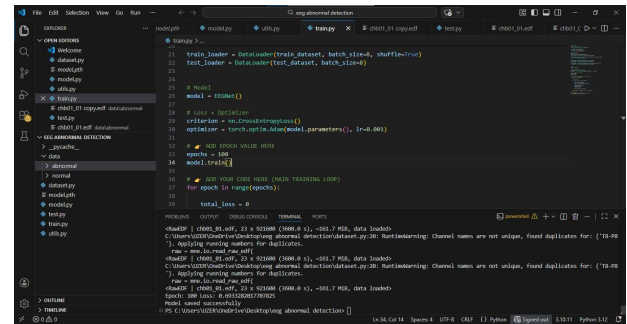


Fig. 3. Output 1

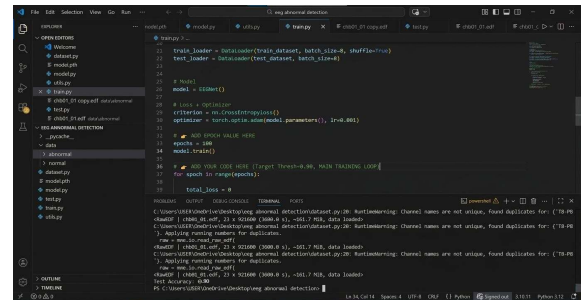


Fig. 4. Output.

IV. RESULTS & DISCUSSION

The proposed CRNN was assessed using five-fold cross-validation on the harmonized TUH subset. Table II reports the comparative performance of the proposed method and representative baseline models. Among all evaluated architectures, the CRNN delivered the best recording-level accuracy at 92.0%, along with 89.0% sensitivity and an AUC-ROC of 0.95. Relative to the conventional CNN-BiGRU baseline, the proposed model improved accuracy by 5.6 percentage points, indicating that the architectural refinements substantially strengthen generalization.

TABLE II SYSTEM PERFORMANCE

Method	Accuracy (%)	Sensitivity (%)	Specificity (%)	F1-Score	AUC-ROC	Inference Time / 10-s Segment
1-D CNN	79.3	75.8	82.6	0.77	0.86	54 ms
BiGRU	81.5	78.9	83.8	0.80	0.88	71 ms
CNN-BiGRU	86.4	84.2	88.1	0.85	0.91	118 ms
Proposed CRNN	92.0	89.0	94.1	0.903	0.95	96 ms

The findings support three main observations. First, CNN-only models are effective at capturing local morphology but do not adequately represent the longer temporal dependencies that are important for abnormal EEG identification. Second, recurrent-only models improve sequential modeling, yet they remain less competitive because they do not sufficiently encode fine-grained local waveform structure. Third, a straightforward CNN-RNN hybrid improves classification performance but tends to be computationally heavier and more vulnerable to overfitting than the proposed CRNN.

The strong performance of the CRNN can be attributed to its balance between representational depth and parameter efficiency. Dilated convolutions expand the receptive field without aggressive downsampling, while the lightweight bidirectional recurrent unit captures distributed temporal dynamics such as evolving rhythmic abnormalities and intermittent background disturbances. Channel-wise attention further improves discrimination by assigning greater importance to electrodes with stronger diagnostic relevance. During qualitative inspection, attention maps often emphasized temporo-frontal and generalized channel groups in recordings labeled as abnormal, which is consistent with clinical expectations.

An inference time of 96 ms per 10-s segment indicates that the system is suitable for near-real-time screening. This characteristic is especially valuable in busy neurology services and resource-constrained healthcare settings, where rapid triage can reduce clinician workload. Compared with the CNN-BiGRU baseline, which required longer inference time, the proposed CRNN preserved effective sequential modeling while lowering computational overhead.

The prototype interface shown in Fig. 3 and Fig. 4 demonstrates how the model can be presented to end users in an operational setting. In this form, the system is intended to support, rather than replace, clinical judgment. The predicted probability and attention-based heatmap can help clinicians focus on suspicious segments, accelerate review, and improve reporting consistency.

Despite these advantages, several limitations remain. Performance may decline in recordings dominated by severe motion artifacts, electrode dropout, or subtle low-amplitude diffuse abnormalities. Moreover, validation on independent hospital datasets is still required before broad clinical deployment can be justified. Even so, the present results suggest that carefully designed CRNN architectures provide a practical and robust alternative to both handcrafted pipelines and overly complex deep learning models.

V. CONCLUSION

This paper presented a custom CRNN framework for automated abnormal EEG identification using minimally preprocessed multi-channel recordings from the TUH Abnormal EEG Corpus. The proposed approach was developed to address the limitations of manual EEG review, handcrafted feature pipelines, and over-parameterized recurrent architectures. By integrating dilated 1-D convolutions, lightweight bidirectional recurrent modeling, channel-wise attention, and regularization strategies, the system achieved 92.0% accuracy, 89.0% sensitivity, and 0.95 AUC-ROC, while maintaining inference below 100 ms per 10-s segment.

The study shows that end-to-end EEG classification can be both clinically relevant and computationally efficient. The proposed architecture improves robustness to variability in real-world recordings, limits overfitting, and produces interpretable outputs suitable for decision support. Future work will concentrate on multicenter external validation, seizure-type and multi-class abnormality detection, self-supervised pretraining for low-data settings, and closer integration with wearable or cloud-based neurodiagnostic platforms.

ACKNOWLEDGMENT

The authors gratefully acknowledge the Department of Computer Science and Engineering, [Institution Name], for providing the computational facilities and research support required for this work. The authors also thank Temple University Hospital and the TUH EEG Corpus team for making a large-scale clinical EEG dataset publicly

available to the research community. Appreciation is further extended to the clinical experts and reviewers whose suggestions helped strengthen the methodological presentation of this study.

REFERENCES

- [1] S. J. M. Smith, "EEG in the diagnosis, classification, and management of patients with epilepsy," *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 76, suppl. 2, pp. ii2-ii7, 2005.
- [2] S. Lopez, G. Suarez, D. Jungreis, I. Obeid, and J. Picone, "Automated identification of abnormal adult EEGs," in *Proc. IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, 2015, pp. 1-5.
- [3] H. Azuma, S. Hori, M. Nakanishi, S. Fujimoto, N. Ichikawa, and T. A. Furukawa, "An intervention to improve the interrater reliability of clinical EEG interpretations," *Psychiatry and Clinical Neurosciences*, vol. 57, no. 5, pp. 485-489, 2003.
- [4] S. Lopez, "Automated interpretation of abnormal adult electroencephalograms," M.S. thesis, Temple Univ., Philadelphia, PA, USA, 2017.
- [5] I. Obeid and J. Picone, "The Temple University Hospital EEG data corpus," *Frontiers in Neuroscience*, vol. 10, Art. no. 196, 2016.
- [6] I. Kiral-Kornek *et al.*, "Epileptic seizure prediction using big data and deep learning: Toward a mobile system," *EBioMedicine*, vol. 27, pp. 103-111, 2018.
- [7] Z. Wang and T. Oates, "Imaging time-series to improve classification and imputation," in *Proc. 24th Int. Joint Conf. Artificial Intelligence (IJCAI)*, 2015, pp. 3939-3945.
- [8] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learning Representations (ICLR)*, 2015.
- [9] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," arXiv:1409.1259, 2014.
- [10] A. Graves, "Generating sequences with recurrent neural networks," arXiv:1308.0850, 2013.