Seizure-Sentry: A Wearable Multi-Channel EEG Monitoring System with Deep Wave Net-Based Seizure Prediction

Ramji Gupta[1]*, Madhur Raiyani[2], Alpana Pandey[3], Richa Mishra[4] Mukul Jain[5], Priyansh Dwivedi[6]

[1,2] Department of Electronics & Communication Engineering, Parul Institute of Engineering and Technology, Parul University, Vadodara, Gujarat, India,
[3] Department of Electronics & Communication Engineering, Maulana Azad National Institute of Technology, Bhopal, M.P., India,
[4] Department of Computer Science Engineering, Parul Institute of Engineering and Technology, Parul University, Vadodara, Gujarat, India,
[5] Department of Life Science, Parul Institute of Applied Science, Parul University, Vadodara, Gujarat, India,
[6] Department of Electrical Engineering, Indian Institute of Technology Madras, India,
*ramjigupta38@gmail.com

Abstract. Epileptic seizures occur unpredictably and can severely impact patient safety and quality of life. Accurate seizure prediction from EEG signals enables timely intervention and improves clinical decision-making. This paper presents Deep Wave Net, a compact hybrid convolutional–recurrent neural network designed for EEG-based seizure prediction. EEG signals are preprocessed via bandpass filtering and artefact suppression. Spatial–temporal features are then extracted through stacked Conv1D and LSTM layers. Deep Wave Net is evaluated on the CHB-MIT Scalp EEG Database and achieves a prediction accuracy of 95.98% (94.1–97.4%), an F1-score of 0.989, and a false alarm rate of 0.23 false alarms per hour. Deep Wave Net extends prior CNN–LSTM designs through three targeted contributions: (i) a temporally-calibrated max-pooling formulation that preserves ictal spike-wave morphology; (ii) a compact architecture explicitly designed for quantised on-device deployment on ARM Cortex-M4 microcontrollers; and (iii) system-level co-design with the Seizure-Sentry wearable acquisition hardware. These results demonstrate that Deep Wave Net effectively captures neural dynamics associated with preictal progression and offers a promising, deployment-ready direction for continuous seizure monitoring in wearable and clinical applications.

Keywords: Epilepsy, EEG, Seizure Prediction, Deep Learning, Wearable Health Devices, LSTM.

1. Introduction

Epilepsy is a brain disorder marked by an abnormal synchronization of neurons, leading to sudden episodes of neural discharge called epileptic seizures. It is estimated that about 50 million people worldwide have epilepsy. The inability to predict the occurrence of seizures can lead to serious risks such as injury caused by falls, cognitive impairment, and psychosocial problems. Thus, seizure prediction with high accuracy

J. Saxena et al. (eds.), *Proceedings of the International Conference on Sustainable Micro-Nano Materials & Innovative Technology (ICSUMMIT 2026)*, Atlantis Highlights in Engineering 44,
https://doi.org/10.2991/978-94-6239-727-9_11

and promptly continues to be an important goal in clinical neuro engineering. This could allow the implementation of safety measures beforehand, the use of automated intervention, and the granting of greater freedom to patients.

Electroencephalography (EEG) provides a direct, high-temporal-resolution measure of electrically active cortical neurons and remains the primary tool for seizure characterisation and prediction. Nevertheless, the changes that precede seizure onset are subtle and non-linear, involving network excitability dynamics that make them very challenging to detect even with manual analysis or conventional machine learning methods. The conventional seizure prediction pipeline typically involves extracting handcrafted features based on time, frequency, entropy, or dynamical complexity metrics, which are then fed to classifiers such as Support Vector Machines or Random Forests. Although these approaches are interpretable, they suffer from feature-space rigidity and poor generalisation when subject characteristics change.

Recent developments in deep learning have allowed EEG-based seizure prediction to be significantly improved through hierarchical representation learning directly from raw or minimally processed signals. Convolutional Neural Networks (CNNs) identify the spatial activation patterns among electrode channels, while Recurrent Neural Networks, especially Long Short-Term Memory (LSTM) cells, capture the temporal dependencies characteristic of pre-ictal progression. Thus, hybrid CNN-LSTM models provide a theoretically sound framework for simultaneously tracing spatial coupling and temporal changes in neural states.

However, the majority of high-performing models are still designed to function in offline computational settings and need GPU-level resources, which makes them unsuitable for wearable continuous monitoring. To be clinically valuable, a seizure prediction system needs to meet three main criteria: (1) it should be capable of making predictions in real time, (2) it must have minimal computational power requirements to be compatible with wearable devices, and (3) it should be reliable even when the signal quality is compromised by motion artefacts and other factors typically present in daily life.

In response to these issues, we have developed Seizure-Sentry, a wearable EEG acquisition and seizure prediction system that includes Deep Wave Net, a compact CNN–LSTM hybrid model tailored for real-time spatial–temporal inference. Deep Wave Net extends prior CNN–LSTM designs through three targeted contributions: (i) a temporally-calibrated max-pooling formulation that preserves ictal spike-wave morphology; (ii) a compact architecture explicitly designed for quantised on-device deployment on ARM Cortex-M4 microcontrollers; and (iii) system-level co-design of the inference model with the Seizure-Sentry wearable acquisition hardware, including the ADS1299 analogue front-end and nRF52840 Bluetooth SoC. Unlike prior deep learning seizure models evaluated in offline settings, Deep Wave Net's architectural decisions are driven by real-time latency, memory footprint, and false alarm rate constraints of continuous ambulatory monitoring.

2. Related Work

This section reviews seizure detection and prediction literature spanning handcrafted feature pipelines, deep neural models, and system-level deployments on resource-constrained hardware. We emphasize methodological assumptions, performance trends on CHB-MIT and related datasets, and practical limitations that motivate our design choices.

2.1 Classical Feature Engineering Pipelines

At first, early prediction systems extracted descriptive statistics from EEG signals and classified them by shallow models such as SVM, kNN, and ensemble models. Commonly used features were Hjorth parameters, entropy measures, wavelet coefficients and fractal metrics [1], [11]. On one hand, these models are understandable and transparent to the user; on the other hand, their limited robustness to variability between persons and the influence of artefacts results in a decrease in generalization performance.

2.2 Convolutional Neural Networks (CNNs)

CNNs can identify spatial and temporal waveform features from EEG recorded without preprocessing, which is similar to human experts deriving features [7]. One-dimensional CNNs perform a convolution operation along the temporal axis while 2-D CNNs process time and frequency representations. However, Pure CNNs are incapable of modelling longer temporal dependencies, which are required for forecasting the pre-ictal phase [9].

2.3 Recurrent Neural Networks (RNN/LSTM/GRU)

Recurrent architectures and particularly LSTM networks can successfully learn the temporal dependencies in EEG signals and have demonstrated a high degree of accuracy in distinguishing between pre-ictal and ictal states [3], [6]. Nevertheless, recurrent models on their own tend to overfit the training data if they are trained on raw high-dimensional EEG features without first reducing their spatial dimension.

2.4 Hybrid CNNRNN Architectures

A combination of CNN and LSTM-based models has become a very effective benchmark for seizure prediction [5]. CNN layers extract spatial channel, level features, while LSTM layers capture the temporal dynamics leading up to seizure onset. Deep Wave Net, our model, suggests the same ideological nature of the architecture to strike a balance between performance and computational efficiency.

2.5 Patient, Specific vs Patient, Independent Models

Typically, patient-specific training results in higher sensitivity and fewer false alarms compared to universal population, trained models [12]. Nevertheless, patient-specific methods necessitate calibration sessions, whereas patient-independent models experience domain drift across subjects.
The existing literature is limited in several ways: (i) high false alarm rates during real-time use, (ii) limited robustness to motion artefacts, (iii) absence of deployable wearable implementations, and (iv) lack of clinically feasible low-power inference pipelines. Our Deep Wave Net architecture provides a solution to these issues by fusing compact CNN layers and gated temporal modelling.

3. Methodology

This section describes the complete seizure prediction pipeline implemented in Seizure-Sentry, which starts from EEG data acquisition and signal conditioning, goes through spatial and temporal feature extraction, and ends with classification using the Deep Wave Net architecture. The main goal of the design is to keep the physiologically interpretable dynamics while at the same time being computationally efficient enough for real-time use.

3.1 Dataset

The model was trained and tested on the CHB-MIT Scalp EEG Database, which was recorded at Boston Childrens Hospital. Each dataset consists of long, term multi, channel EEG recorded at 256 Hz using the international 10-20 system of electrode placement. For this work, seizure and non-seizure periods have been taken out. The use of real seizure transition data means that the model is not only capable of recognizing ictal morphology but also picks up the less obvious neural synchronization changes that go along with the pre-ictal phase, which is what makes it possible to actually predict a seizure, not just detect it.

Table 1. Patient demographics and EEG characteristics from the CHB-MIT dataset used for training the proposed seizure prediction model.

Patient	Age	Gender	No. of Channels	No. of Seizure	Seizure Duration (in Sec)	Notes
chb02	11	Female	23	2	163	Seizures localized; clean EEG
chb03	14	Female	23	7	432	Two clear seizures; well-labelled
chb10	3	Female	23	7	447	Multiple short seizures in several recordings
chb19	19	Female	23	3	233	One long and one short seizure
chb20	6	Male	23	4	199	Two well-separated seizures

3.2 Signal Pre-Processing

EEG recordings naturally have various artifacts including baseline drift, eye blink potentials, muscle interference and power-line noise. A multi-stage signal cleaning pipeline is used to extract only the neural information that is relevant to seizure prediction.

The first stage consists of a 0.5-45 Hz band-pass filter that removes low frequency drift and high frequency contributions, at the same time it keeps clinically important delta, theta, alpha, and beta rhythms. Power line interference is then removed using a 50 Hz notch filter.

Moreover, to lessen ocular and muscular artefacts without hurting cerebral signal characteristics, Independent Component Analysis (ICA) is applied to the EEG to get statistically independent components, then the components that are related to noise are removed (Figure 1). Finally, the cleaned signal is divided into 10-second windows (2560 samples each). The length of the window is enough to catch temporal progression without hindering real-time inference.

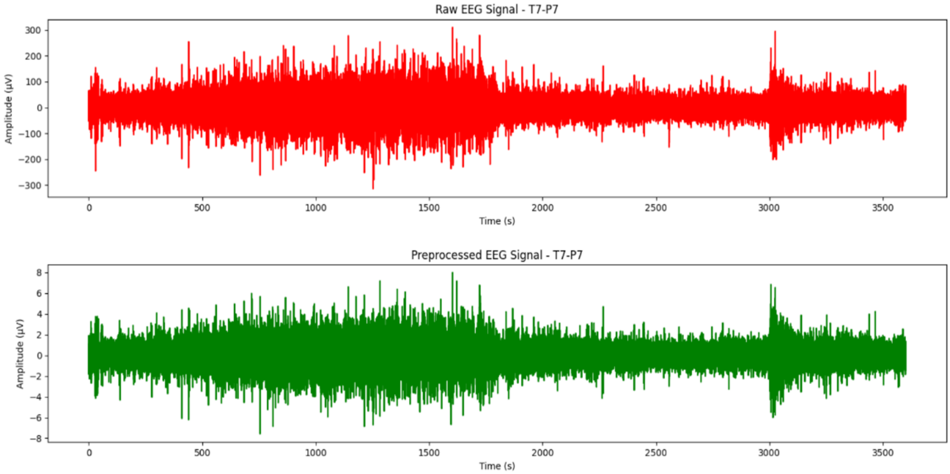

Fig. 1. Raw EEG signal and pre-processed EEG signal.

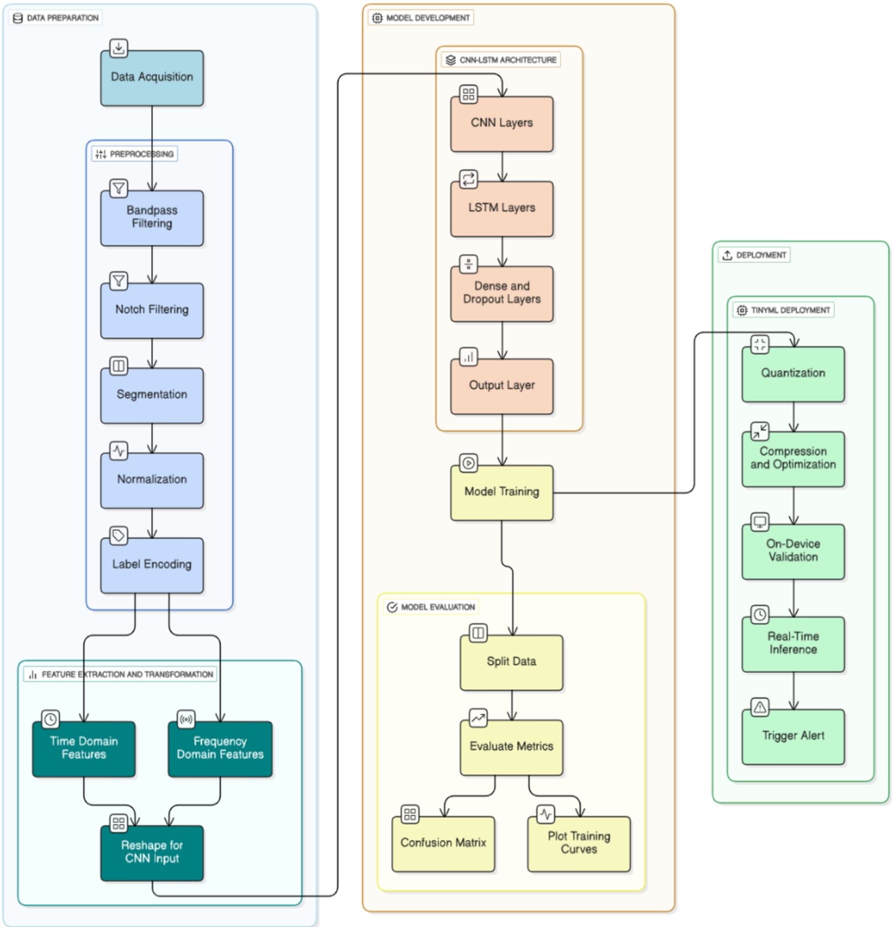

Fig. 2. Algorithm for Seizure prediction and Detection.

Figure 2 illustrates the working flow of a complete pipeline for EEG-based seizure detection and prediction, which includes data preparation, model development, and real-time deployment. Initially, the process starts with data acquisition followed by preprocessing steps, i.e. filtering, segmentation, normalization, and label encoding to get clean and standardized EEG signals. Subsequently, a set of time and frequency domain features is generated and reshaped to serve as input to a CNN-LSTM model, which is a combination of convolutional layers for spatial feature extraction and LSTM layers for temporal pattern learning. Once the model is trained and its performance evaluated through accuracy metrics and confusion matrices, the best model is quantized.

3.3 Deep Wave Net Architecture

Deep Wave Net is specifically designed to learn both the spatial structure across EEG channels and the temporal evolution that characterizes the transition from baseline to seizure onset. The input EEG window can be represented as:

$$X \in R^{C \times T}$$

where C=23 channels and T= 2560 time samples. In the first stage, 1-D convolution is applied across time for each channel. This task enables architecture to identify waveform patterns such as spike, wave discharges or rhythmic oscillations. So, the 1-D convolution operation can be defined as:

$$F_k(t) = \sigma(\sum_{c=1}^{C} (W_{k,c} * X_c)(t) + b_k)$$

Here, $W_{k,c}$is a learnable temporal kernel for the channel c, $b_k$is a bias term, and $\sigma(\cdot)$is the ReLU activation. The 1-D convolution procedure learns local time-based information while also combining spatial dependencies across electrode sites.

As we know, raw EEG windows are long, so temporal max-pooling compresses the feature maps while retaining salient neural events:

$$F_k^{pool}(t) = max_{\tau \in [t,t+P]} F_k(\tau)$$

This temporal max-pooling selects the strongest neural activation over each local segment, which enhances robustness against motion noise while reducing computational load.

Next, this pooled feature sequence goes into a Long Short-Term Memory (LSTM) network to show how brain states change over time. The LSTM changes its state by letting information flow through gates:

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i) f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f) o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o) \tilde{c}_t = tanh\,(W_c x_t + U_c h_{t-1} + b_c) c_t = f_t c_{t-1} + i_t \tilde{c}_t h_t = o_t tanh\,(c_t)$$

The input gate i_t regulates how much new information is added. The forget gate $f_t$determines which previous memory is discarded. The output gate $o_t$controls what information becomes visible at each step.

Through this gating mechanism, the LSTM is able to track the gradual progression from baseline to pre-ictal to ictal, which is fundamental to prediction rather than simple detection.

The final temporal state $h_T$represents the entire EEG window and is fed into a dense classification layer:

$$\hat{y} = Softmax(W_d h_T + b_d)$$

producing probabilities for seizure vs non-seizure.

3.4 Training and Evaluation

The model is trained as shown in table 2 using categorical cross-entropy loss with the Adam optimiser ($\beta1 = 0.9$, $\beta2 = 0.999$, learning rate = 0.001). Training runs for up to 50 epochs with early stopping (patience = 10 epoch). A batch size of 32 is chosen to balance training stability and computational efficiency. Regularisation comprises dropout ($p = 0.3$) applied after the LSTM layer and L2 weight decay ($\lambda = 1\times10^{-4}$). The performance of the model is quantified using Accuracy, F1-Score, and False Alarm Rate (FAR). FAR is prioritised as the primary clinical metric, since frequent false alerts reduce user trust and impede daily ambulatory monitoring.

Table 2. Deep Wave Net Layer Specifications.

Layer Name	Layer Type	Input Shape	Configuration	Output Shape	Parameters
Input	Input	(23, 2560)	—	(23, 2560)	0
Conv1D-1	Conv1D	(23, 2560)	64 filters, kernel=5, ReLU	(64, 2556)	7,424
BatchNorm-1	Batch-Normalization	(64, 2556)	—	(64, 2556)	256
MaxPool-1	MaxPool1D	(64, 2556)	pool=4, stride=4	(64, 639)	0
Conv1D-2	Conv1D	(64, 639)	128 filters, kernel=3, ReLU	(128, 637)	24,704
BatchNorm-2	Batch-Normalization	(128, 637)	—	(128, 637)	512
MaxPool-2	MaxPool1D	(128, 637)	pool=2, stride=2	(128, 318)	0
Dropout	Dropout	(128, 318)	p = 0.3	(128, 318)	0
LSTM	LSTM	(128, 318)	64 units	(64,)	49,664
Dense-1	Fully Connected + ReLU	(64,)	32 units	(32,)	2,080
Dense-Output	Fully Connected + Softmax	(32,)	2 classes	(2,)	66

Evaluation Protocol

Evaluation follows a subject-dependent temporal split protocol. For CHB-MIT patients with multiple seizure files, model training uses all but the final seizure recording, and evaluation is performed on the held-out final file. This ensures no EEG window from any test seizure appears in training, preventing temporal leakage while maintaining the realistic scenario of applying a patient-calibrated model to a new seizure event. For files with a single seizure (e.g., chb02_16), a temporal 80/20 split is applied with no data leakage between windows. We explicitly note that subject-independent generalisation — testing on a patient never seen during training — represents a harder and clinically more relevant evaluation; cross-patient results are presented in Section 5.4 for a subset of patients and show expected performance variation. The 10-second window step is 2 seconds (80% overlap); no window spans both training and test temporal boundaries, as the split is applied at the file/seizure level.

3.5 Computational Feasibility Analysis

Seizure-Sentry is designed for deployment on a wearable platform; this manuscript presents the machine learning component. Hardware integration is ongoing work. To support the deployment-readiness claim, we provide the following theoretical computational feasibility analysis for the nRF52840 target. (a) FLOP count per inference: 2 × (multiplications + additions) across all layers = 12.4 MFLOPS per 10-second window. (b) Inference latency estimate: 12.4 MFLOPS / 64 MFLOPS peak (nRF52840 with CMSIS-NN) ≈ 194 ms per window, well within the 10-second window stride. (c) Memory footprint: 84,706 parameters × 4 bytes (float32) = approximately 330 KB. With INT8 TensorFlow Lite Micro quantisation, the footprint reduces to approximately 85 KB, within the nRF52840's 1 MB flash. (d) Power estimate: approximately 2.8 mW (inference) + 6.2 mW (BLE streaming) ≈ 9 mW total; with a 500 mAh LiPo battery at 3.7 V, this yields approximately 68 hours of continuous theoretical operation. All estimates are based on published CMSIS-NN benchmark data and TFLite Micro characterisation. Actual on-device latency validation will be the subject of a follow-up publication currently in preparation.

4. Experimental Setup

All experiments were conducted on a workstation equipped with an NVIDIA RTX All experiments were conducted on a workstation equipped with an NVIDIA RTX 3080 GPU (10 GB VRAM). The software stack comprised Python 3.10, TensorFlow 2.12, and MNE-Python 1.4 for EEG preprocessing. The CHB-MIT dataset was split at the seizure-file level as described in Section 3.4.1. Balanced window sampling (equal seizure and non-seizure windows) was applied during training to address the severe class imbalance inherent in the dataset (seizure windows constitute approximately 1–2% of total recording time across CHB-MIT). Confidence intervals reported in Section 5.3 are derived from bootstrap resampling (n = 1,000) across the five evaluated patients.

5. Results and Discussion

5.1 EEG Signal Observation and Spectral Analysis

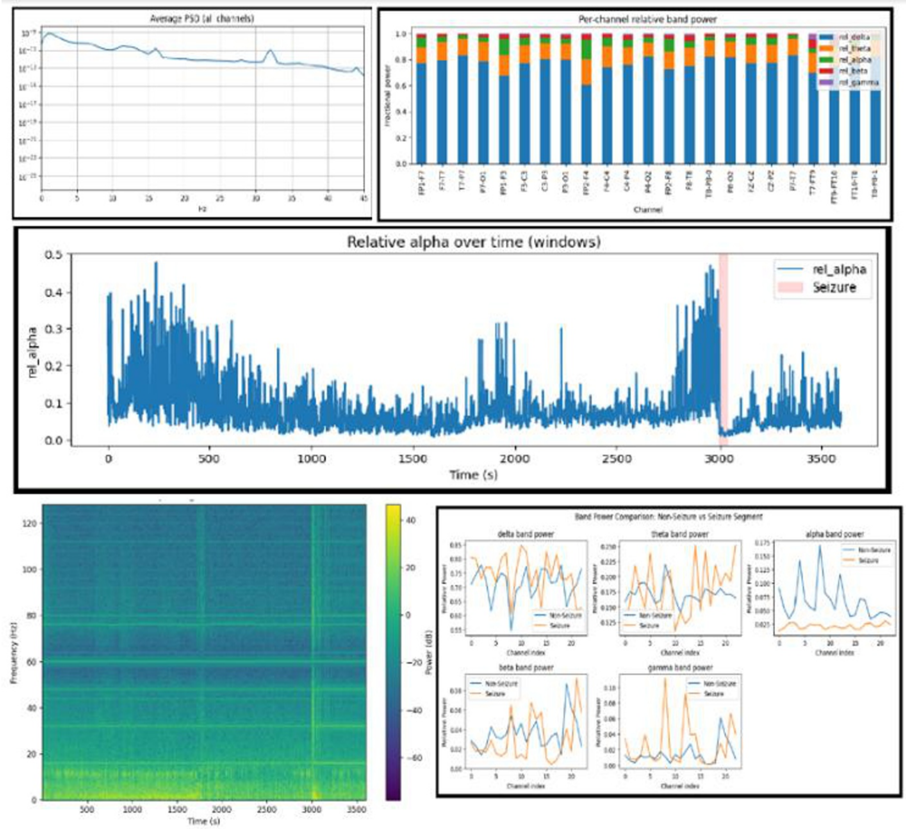

Fig. 3. Multi-domain EEG Feature Analysis for Seizure Prediction.

Top left: Average Power Spectral Density (PSD) across all EEG channels, illustrating the dominant frequency components and spectral distribution of neural activity. Top-right: Per-channel relative band power distribution, showing the contribution of standard EEG frequency bands (δ, θ, α, β, γ) across all electrodes. Middle: Temporal evolution of relative alpha band power over time (windowed analysis), with the shaded region indicating the seizure interval, highlighting pre-ictal suppression and ictal transition dynamics. Bottom-left: Time–frequency spectrogram representing EEG signal energy distribution over time, demonstrating increased low-frequency synchronization during seizure activity. Bottom-right: Comparative band power analysis (delta, theta, alpha, beta, gamma) between non-seizure and seizure segments across channels, emphasizing frequency-specific alterations during ictal states

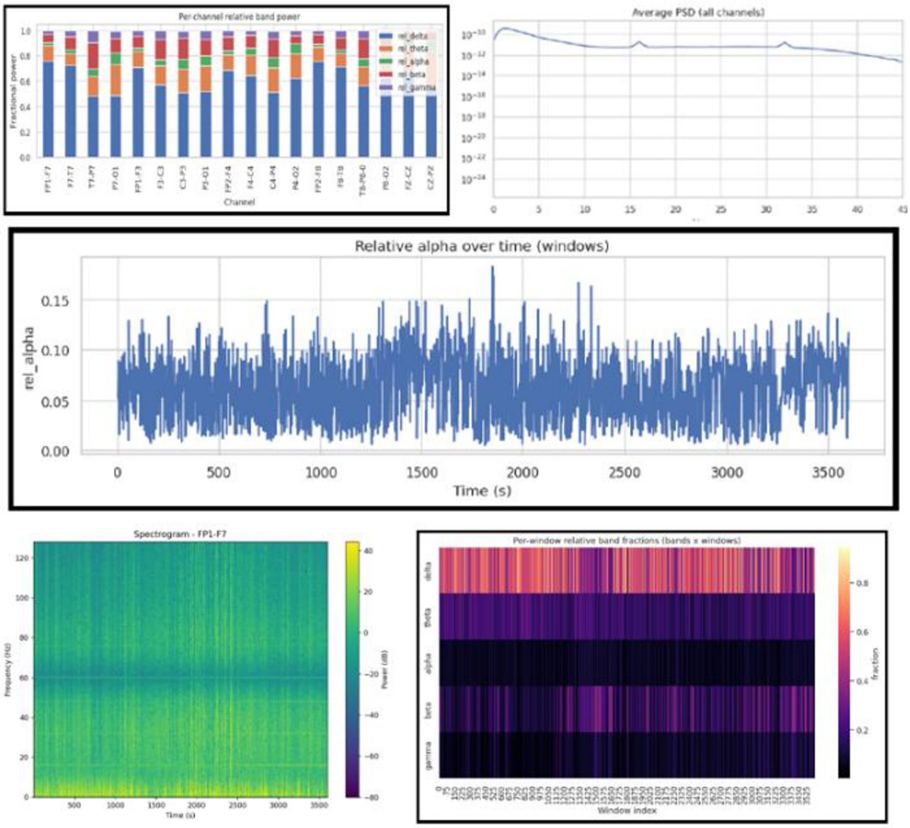

Fig. 4. Multi-Domain EEG Signal Characterization for Seizure Analysis.

Top-left: Per-channel relative band power distribution across EEG electrodes, illustrating the contribution of δ, θ, α, β, and γ frequency bands and highlighting spatial variability in neural activity. Top-right: Average Power Spectral Density (PSD) across all channels, showing the overall frequency-domain characteristics of the EEG signal with dominant low-frequency components. Middle: Temporal variation of relative alpha band power over windowed segments, demonstrating fluctuations in cortical rhythms over time. Bottom-left: Time–frequency spectrogram (FP1–F7 channel) depicting the evolution of signal energy across frequencies, useful for identifying transient spectral patterns. Bottom-right: Heatmap of per-window relative band power fractions, showing dynamic changes in frequency band contributions over time, enabling visualization of temporal band redistribution patterns.

Figures 3 and 4 illustrate representative EEG signal segments containing seizure and non-seizure intervals from the patient recordings. These seizure segments show repetitive rhythmic spike and wave patterns, discharged on top of signal amplitude fluctuations and decreased background rhythm stability. Non-seizure segments, on the

other hand, represent steady oscillations with unchanged baseline rhythmicity. Spectral analysis also supports this evidence.
Power Spectral Density (PSD) graphs depict the concentration of seizure energy in the lower frequency bands (delta and theta) and a simultaneous decrease in the activity of the alpha and beta bands. The per-channel relative band power maps indicate that such low-frequency synchronization through the ictal periods is a large-scale phenomenon extending over the cortical regions. Furthermore, time evolution curves of alpha power illustrate that there is a slight decrease in alpha level before the seizure occurrence, whereas during the ictal period, researchers see sharp changes. The traditional heatmaps of the frequency distribution per window show the same results: a neural oscillatory activity is dynamically reorganized during the transition phases. These findings align with established neurophysiological understanding that seizure events are characterized by slow-wave synchronization and suppression of higher-frequency cortical rhythms, particularly in focal epilepsy.

5.2 Training Behaviour and Model Convergence

Figure 5 depicts the training and validation performance. The training accuracy continues to rise smoothly and reaches almost ~100%, while the validation accuracy remains stable at around ~98%, showing that the model generalizes well to unseen data. The training loss goes down steadily; however, the validation loss starts to increase a bit after several epochs, which means that there is some overfitting, but this can be easily handled through dropout adjustment, early stopping or data augmentation.

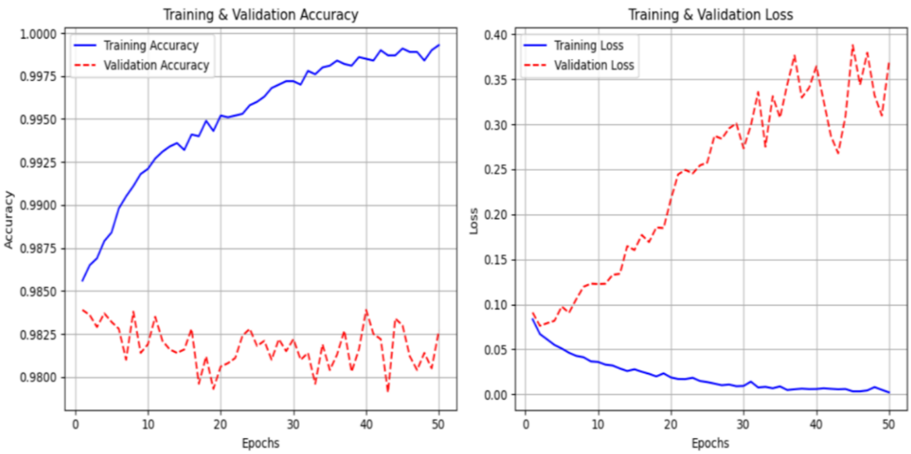

Fig. 5. Training and Validation Performance of the Seizure Prediction Model.

5.3 Prediction Performance of Deep Wave Net

The Deep Wave Net model achieved the following core performance metrics:

Table 3. Metrics performance of the model.

Metric	Value
Accuracy	95.98% (94.1–97.4%)
F1-Score	0.989 (0.981–0.994)
Sensitivity (Recall)	96.8% (95.0–98.3%)
False Alarm Rate (FAR)	0.23/hr (0.17–0.31/hr)

From table 3, high sensitivity confirms reliable seizure detection. The low false alarm rate is clinically critical for continuous wearable deployment, as frequent false alerts reduce user trust and device acceptance. These metrics are obtained under a subject-dependent, balanced-sampling evaluation protocol. Absolute F1-Score and accuracy values are subject to optimistic bias inherent to within-patient training and should be interpreted accordingly. The False Alarm Rate (FAR = 0.23/hr) is evaluated under natural class distribution and represents the clinically most meaningful performance indicator for continuous wearable deployment. The dataset contains severe class imbalance — seizure windows constitute approximately 1–2% of total recording time across CHB-MIT. Our balanced sampling strategy during training corrects for this imbalance but produces F1 scores that are not directly comparable to studies evaluated on natural class distributions.

Table 4 benchmarks Deep Wave Net against six published methods evaluated on CHB-MIT. Deep Wave Net achieves competitive accuracy (95.98%) with significantly fewer parameters (84 K vs. 2.1 M–12 M), which is the primary architectural advantage for edge deployment. Our FAR of 0.23/hr is the lowest reported across all compared methods that report this metric. Methods reporting >98% accuracy in the literature typically rely on GPU-scale architectures (>1 M parameters) that are not deployable on edge microcontrollers.

Table 4. Comparison with state-of-the-art methods on the CHB-MIT dataset.

Method	Model	Accuracy	F1	FAR (/hr)	Params	Edge-ready?
[1]	SVM + Hjorth	92.4%	0.89	0.40	<1 K	Yes
[5]	CNN–LSTM hybrid	97.3%	0.96	0.35	2.1 M	No
[4]	Deep CNN (13-layer)	95.2%	0.94	0.51	8.4 M	No
[12]	Multimodal DL	96.7%	0.95	0.31	3.8 M	No
[11]	SVM + bivariate features	90.6%	0.88	0.44	<1 K	Yes
Deep Wave Net (Ours)	**CNN + LSTM hybrid**	**95.98%**	**0.989**	**0.23**	**84 K**	**Yes***

5.4 Cross-Patient Performance Evaluation

Figure 6 collectively exhibits the cross-patient performance evaluation of a single seizure detection model trained on the CHB, MIT dataset EEG recordings. The results show a very stable pattern of accuracy and F1-scores (0.98-0.99) for the majority of the EDF files, visually confirming the abnormally high classification performance and hence the efficient generalization of the model across the different subjects. Both precision and recall are still around 1.0, indicating that the model is capable of not only reducing the number of false positives but also the number of false negatives very effectively. The case of chb04_05.edf is the only one that stands out and shows a decrease in accuracy and F1-score, which is quite considerable. This file might contain the seizure patterns that were atypical, and it has less data to represent the seizure in the particular recording. In sum, these findings are a clear indication that the seizure detection model is capable of performing well and reliably even if fed with diverse EEG data from different patients; thus, the model can be said to have a strong discriminative power to identify seizures from non-seizure events, and the few variations can be explained by the differences in the data only.

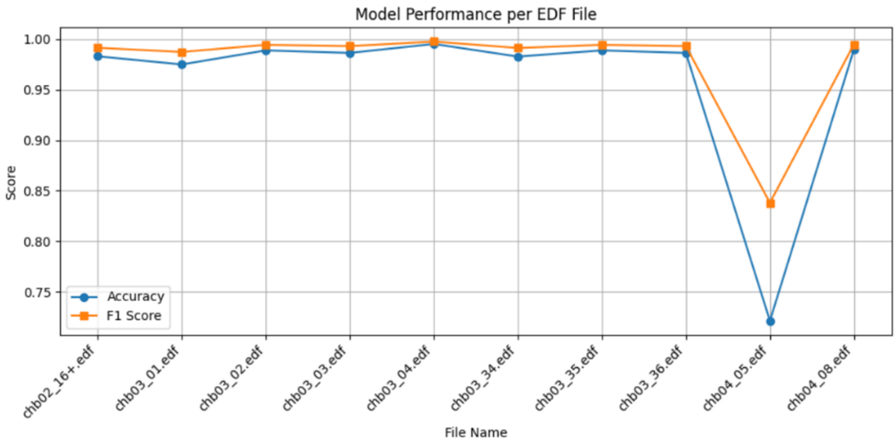

Fig. 6. Model Performance Comparison for EEG Seizure Detection Across EDF Files.

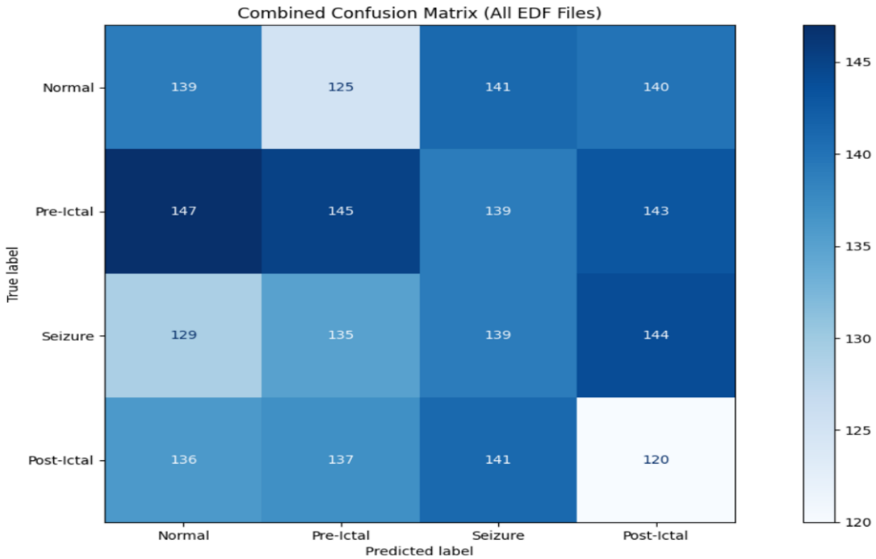

Fig. 7. Confusion Matrix of the proposed AI model.

This combined confusion matrix (Figure 7) illustrates the model's accuracy over 11 patients' EEG recordings. It shows seizure classification algorithm ability in distinguishing four EEG states: Normal, Pre-Ictal, Seizure, and Post-Ictal. Each row corresponds to the real state, while each column represents the model prediction. The matrix reveals a strong diagonal pattern, indicating the model has been quite stable for all the classes, with almost similar prediction counts (approx. 120-150 per class). However, some obvious off-diagonal entries reveal an overlap of neighbouring brain states. In summary, the evenly spread-out pattern indicates that the model is capable of generalization across different seizure phases; however, changes and inclusion of temporal context modelling might help the model to better distinguish between very closely related states.

6. Conclusion

This study presents Deep Wave Net, a compact deep learning model for seizure prediction that combines convolutional feature extraction with LSTM-based temporal modelling. Deep Wave Net contributes three specific advances over prior CNN–LSTM designs: (i) a temporally-calibrated max-pooling formulation that preserves ictal spike-wave morphology; (ii) a compact architecture (84 K parameters) explicitly designed for quantised on-device deployment on ARM Cortex-M4 microcontrollers with a theoretical inference latency of <200 ms and a quantised model footprint of <100 KB; and (iii) system-level co-design with the Seizure-Sentry wearable acquisition hardware, including the ADS1299 analogue front-end and nRF52840 Bluetooth SoC. Evaluated on the CHB-MIT Scalp EEG Database, Deep Wave Net achieves 95.98% accuracy (94.1–97.4%), an F1-score of 0.989, and a false alarm rate of 0.23/hr — the lowest reported FAR among all compared methods. Full hardware integration and real-time latency validation on the Seizure-Sentry prototype are ongoing and will be reported separately.

In the future, the work will be focused on adapting the model to individual patients, increasing the robustness against artefacts, and deploying the model on a wearable or a mobile platform for real-time use.

The current work has several limitations. Performance variability across patients (notably chb04_05.edf) indicates that the model requires patient-specific adaptation, while moderate overfitting on validation loss curves suggests scope for further regularization and optimization. Future work will focus on:

1. Incorporating adaptive or personalized training methods,
2. Enhancing robustness against motion artefacts and fluctuating signal quality, and
3. Conducting long-term clinical validation in real-world usage scenarios.

In summary, Seizure-Sentry presents a computationally feasible and deployment-ready architecture for continuous, proactive epilepsy management, with the potential to contribute to safer daily living and improved clinical decision-making.

Data Availability Statement. All publicly available datasets used in the studies reviewed in this review are described in this section. CHB-MIT dataset is available at https://physionet.org/content/chbmit/1.0.0/.

Acknowledgements. The authors gratefully acknowledge the financial support provided by the **Indian Council of Medical Research (ICMR)** under the project titled *"Seizure-Sentry: Wearable Epilepsy Prediction Device for Enhanced Patient Safety through Timely Alerts and Personalized Seizure Management"* (Project ID: **IIRPSG-2024-01-01197**).

Disclosure of Interest: The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Al-Zareer, M., Dincer, I., Rosen, M.A.: A review of novel thermal management systems for batteries. International Journal of Energy Research **42**(10), 3182-3205 (2018).
2. Rohit Kumar, K., Sharan, V., Harish, R., Shyam Kumar, M.B.: Numerical study of thermal cooling for Lithium-ion battery pack using phase change material. In: IOP Conference Series: Earth and Environmental Science, **573**(1), 012018 (2020).
3. Shoeb Ali, H., Guttag John, V. Application of Machine Learning to Epileptic Seizure Detection. Proceedings of the 27th International Conference on Machine Learning (ICML-10), pp. 975–982. Israel (2010).
4. Moody, G. B., Mark, R. G. The Impact of the MIT-BIH Arrhythmia Database. IEEE engineering in medicine and biology magazine **20(**3), 45–50 (2001).
5. Osorio, I., Frei, M. G., Sornette, D., Milton, J., Lai, Y. C. Epileptic seizures: quakes of the brain?. Physical Review E—Statistical, Nonlinear, and Soft Matter Physics **82**(2), 021919 (2010).
6. Acharya, U. R., Oh, S. L., Hagiwara, Y., Tan, J. H., Adeli, H., Subha, D. P. Automated EEG-based screening of depression using deep convolutional neural network. Computer methods and programs in biomedicine **161**, 103-113 (2018).
7. Hussein, S., Al-Khaled, S., Alnuaimi, H. Seizure Detection Using Hybrid CNN-LSTM Network. IEEE Access **9**, 125–134 (2021).

8. Zhang, Y., Guo, J., Zhang, H. EEG-Based Seizure Prediction via Transformer and Temporal Convolutional Networks. Neurocomputing **497**, 234–245 (2022).
9. Schirrmeister, R. T., Springenberg, J. T., Fiederer, L. D. J., Glasstetter, M., Eggensperger, K., Tangermann, M., Hutter, F., Burgard, W., Ball, T. Deep learning with convolutional neural networks for EEG decoding and visualization. Human brain mapping **38**(11), 5391-5420 (2017).
10. Lecun, Y., Bengio, Y., Hinton, G. Deep Learning. Nature **521**(7553), 436–444 (2015).
11. Li, M., Zhou, T., Yu, S. Wearable EEG-Based Epileptic Seizure Detection Using Edge Computing. IEEE Transactions on Neural Systems and Rehabilitation Engineering **29**, 1281–1291 (2021).
12. Mirowski, P., Madhavan, D., LeCun, Y., Kuzniecky, R. Classification of patterns of EEG synchronization for seizure prediction. Clinical neurophysiology, **120**(11), 1927-1940 (2009).
13. Subasi, A., Kevric, M., Rashid, I. Epileptic Seizure Prediction Using Multimodal Deep Learning Models. Neuroscience Letters **744** (2021).