Patient-Adaptive Seizure Detection via Cluster-Based Zero-Shot and Fine-Tuned TCN on Pediatric EEG

Nafisa Binte Ghulam Kibria[1], Zareen Tasnim Pear[2], Md Kishor Morol[3], Ramisa Anjum Oishi[4], Tamzid Tanvi Alam[3], and Diya Rahman[5]

[1] Islamic University of Technology, Gazipur-1704, Bangladesh
nafisa21@iut-dhaka.edu
[2] Rajshahi University of Engineering & Technology, Rajshahi-6204, Bangladesh
pearzareen@gmail.com
[3] ELITE Research Lab, New York, USA
kishoremorol@gmail.com, tamzid.t.alam@gmail.com
[4] California Lutheran University, California, United States
roishi@callutheran.edu
[5] American International University-Bangladesh
drahman024@gmail.com

Abstract. Automated seizure detection from scalp electroencephalography (EEG) remains challenging because of substantial inter-patient variability, which limits the ability of population-trained models to generalize to unseen individuals. In this work, a cluster-based seizure detection framework is presented in which pediatric patients are grouped according to their ictal spectral fingerprints, and a dedicated LightweightTCN model is trained for each cluster in a leave-one-patient-out setting. The framework supports both zero-shot inference and patient-specific fine-tuning using as few as 10 labeled EEG windows. To ensure a rigorous evaluation, a strict hold-out protocol is enforced before any clustering, preprocessing, or training step, thus eliminating data leakage throughout the pipeline. The proposed lightweight TCN incorporates a trainable per-channel softmax gate for a per-channel gate intended to support spatial interpretability and squeeze-and-excitation modules for feature recalibration. Experiments conducted on all 24 patients of the CHB-MIT Scalp EEG Database demonstrate that zero-shot inference achieves a mean F1-score of 0.711 and an AUPRC of 0.869. With patient-specific fine-tuning using a 30% support set, the mean F1-score increases to 0.822 ($\Delta = +0.111$), sensitivity improves to 0.808 ($\Delta = +0.118$), and AUPRC reaches 0.905. Consistent improvements are also observed in MCC (0.571 $\rightarrow$ 0.683) and Cohen's Kappa (0.541 $\rightarrow$ 0.675). The largest improvement is achieved for patient chb15, where the F1-score increases from 0.065 to 0.809 ($\Delta = +0.745$) using 75 support windows. Furthermore, after applying a 30-second persistence filter, zero clinical false alarms are achieved for 16 out of 24 patients. Overall, the proposed framework demonstrates that robust cross-patient seizure detection can be achieved through cluster-aware modelling and lightweight patient-specific adaptation under a leakage-free evaluation setting.

Keywords: Seizure detection · EEG · Temporal convolutional network · Zero-shot learning · Channel attention · Paediatric epilepsy

1 Introduction

Epilepsy affects approximately 50 million people worldwide and is associated with a two-to-three times higher risk of premature death than the general population (21). For many patients, seizures remain refractory to medication, making continuous EEG monitoring important for diagnosis, management, and surgical planning (1).

Despite advances in deep learning, automated seizure detection remains challenging due to substantial inter-patient variability in seizure morphology, frequency content, spatial distribution, and duration. As a result, models often generalise poorly to unseen individuals, particularly across heterogeneous populations (6). While patient-specific models can achieve strong performance, they require labelled data from the target patient, limiting immediate clinical deployment. In contrast, generalised cross-patient models are easier to deploy but typically sacrifice accuracy.

To address this trade-off, two complementary paradigms have emerged. Zero-shot models are applied directly to unseen patients without adaptation, whereas few-shot fine-tuning adapts a pre-trained model using a small amount of labelled patient data. Together, these approaches enable immediate deployment while supporting progressive personalisation as additional EEG data become available. In this work, a cluster-based zero-shot and fine-tuned seizure detection framework is presented using the CHB-MIT paediatric EEG dataset under a strict hold-out evaluation protocol designed to eliminate data leakage. Patients are grouped according to their ictal spectral fingerprints using K-means clustering, and a dedicated LightweightTCN model is trained for each cluster using only peer-patient data. During evaluation, the corresponding cluster model is either applied directly in a zero-shot setting or adapted using a small support set from the target patient. Performance is evaluated exclusively on held-out query windows that are never used during training or preprocessing.

The main contributions of this work are as follows:

1. A strict per-patient hold-out protocol is introduced and applied before clustering, training, or scaler fitting, thereby preventing data leakage throughout the entire pipeline. All evaluation is performed on completely unseen query data.
2. A cluster-specific training strategy based on ictal spectral fingerprints is proposed, in which a dedicated LightweightTCN model is trained for each patient cluster. A cross-cluster regularisation mechanism is incorporated to improve robustness near cluster boundaries.
3. A lightweight depthwise-separable TCN with a trainable channel-attention gate is developed to learn interpretable spatial importance across EEG channels, while squeeze-and-excitation modules are used for feature recalibration.

4. A controlled comparison between zero-shot and fine-tuned seizure detection is conducted using identical query sets, demonstrating that fine-tuning with as few as 10 support windows improves the mean F1-score by +0.111 across 24 patients.
5. Extended clinical evaluation metrics, including AUPRC, MCC, Cohen's Kappa, and persistence-filtered false alarm rate (FAR), are reported. By applying a 30-second persistence filter, the mean FAR is reduced from 59.20 to 21.22 false alarms per hour.

Among these components, the central contribution is the cluster-specific training and leakage-free zero-shot/fine-tuning comparison protocol; the TCN, SE blocks, and channel gate are adopted from prior work and combined within this protocol.

2 Related Work

Automated epileptic seizure detection using electroencephalography (EEG) has evolved from classical signal processing methods to deep learning approaches. However, key challenges remain, including inter-subject variability, limited data availability, poor generalisation, and real-time processing constraints. Moreover, persistent gaps in the literature include (i) data leakage in cross-subject evaluation, (ii) the scarcity of lightweight models for edge deployment, and (iii) limited comparisons between zero-shot and personalised fine-tuned models under strictly leak-free protocols. This work addresses these gaps within a unified framework.

2.1 Seizure Detection on CHB-MIT: Benchmarking and Leakage

The CHB-MIT dataset is a commonly used benchmark for paediatric seizure detection, yet reported performance varies significantly across studies. Many high-performing methods adopt subject-dependent evaluation, where training and testing are performed on the same patient, often reporting accuracies of 97–99%, which may not reflect realistic clinical deployment. Under stricter evaluation protocols, performance drops substantially. Ali et al. (2) reported sensitivities of 72.6% and 75.3% using five-fold and leave-one-out cross-validation, respectively, after accounting for class imbalance, cross-subject variability, and event-level evaluation. Similarly, Shyu et al. (20) observed a reduction in F1-score from 69.3% in subject-dependent settings to 37.3% in subject-independent evaluation. Andrade et al. (4) further showed that alarm-based evaluation reduces accuracy from about 80% to nearly 50% under leave-one-patient-out protocols, partly due to overlapping windowing effects. Overall, these findings suggest that commonly used evaluation pipelines may lead to overestimated real-world performance due to preprocessing overlap, data imbalance, and per-window scoring.

2.2 Temporal Convolutional Networks and Lightweight Design

Temporal Convolutional Networks (TCNs) are widely used for EEG analysis due to their ability to capture long-range dependencies with efficient parallel computation. Several studies report strong performance in patient-specific seizure

detection using TCN-based models such as Rao et al. (17), Huang et al. (13), and Alotaibi et al. (3), but these approaches rely on prior patient-specific data and do not generalise well to unseen subjects. To improve efficiency, Lim et al. (15) proposed a lightweight TCN using depthwise-separable and grouped convolutions, reducing parameters within a federated learning framework. However, cross-patient generalisation and zero-shot evaluation remain less explored in existing lightweight designs.

2.3 Channel Attention and Interpretability

Channel attention mechanisms have been widely used to model spatial variability in EEG signals. Methods such as Baghdadi et al. (5) and Hu et al. (12) show that attention-based architectures can improve performance while enabling channel-level interpretability. Complementary explainability methods such as SHAP-based analysis by Raab et al. (16) and saliency-based approaches by Einizade et al. (9) further highlight task-relevant brain regions. Despite these advances, prior work has not extensively studied the relationship between learned channel gating mechanisms and post-hoc attribution methods within a unified framework.

2.4 Cross-Subject Domain Adaptation and Zero-Shot Strategies

Domain adaptation approaches aim to reduce inter-subject variability in EEG signals. Methods such as Cui et al. (8) and Cui et al. (7) have been proposed to minimise distribution discrepancies between source and target patients. Adversarial and kernel-based techniques by Wang et al. (22) and Jemal et al. (14) further improve cross-subject performance. However, these methods generally require access to target-patient data during adaptation, limiting applicability in zero-shot settings. Recent work shows that minimal fine-tuning on target patients can improve performance, as shown by Shafiezadeh et al. (19), but results are not directly compared with zero-shot inference under identical conditions. As a result, the benefit of personalisation remains unclear in strictly controlled settings.

Motivated by these limitations, this study proposes a leak-free evaluation framework, introduces a cluster-aware lightweight modelling strategy, and provides a controlled comparison between zero-shot and personalised fine-tuned models under identical evaluation conditions. The following sections describe the proposed methodology, including the dataset setup, model architecture, and evaluation protocol in detail.

3 Methodology

Figure 1 illustrates the overall framework of the proposed patient-adaptive seizure detection pipeline. The methodology consists of EEG preprocessing, spectral fingerprint-based patient clustering, cluster-specific LightweightTCN training,

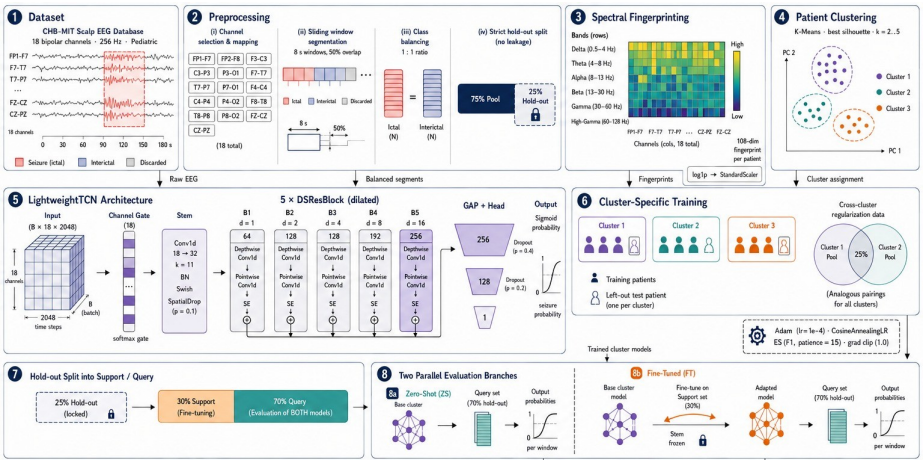

Fig. 1. Overview of the proposed patient-adaptive seizure detection framework. The pipeline includes EEG preprocessing, spectral fingerprint extraction, patient clustering, cluster-specific LightweightTCN training, and parallel zero-shot and fine-tuned evaluation branches.

and two parallel evaluation branches for zero-shot and fine-tuned inference. In addition, the framework incorporates strict hold-out splitting to prevent data leakage and includes interpretability analysis through channel-wise attention and saliency visualization.

3.1 Dataset and Preprocessing

Experiments are conducted on the CHB-MIT Scalp EEG Database (10), which contains continuous EEG recordings from 24 paediatric epilepsy patients collected at Boston Children's Hospital. Recordings follow the International 10-20 system (a standard scalp electrode placement scheme) and are sampled at $f_s = 256\,\text{Hz}$. The dataset includes 198 annotated seizures across approximately 664 hours of EEG data, with substantial inter-patient variability in seizure characteristics.

We use 18 bipolar EEG channels (voltage difference between adjacent electrode pairs) consistently available across all patients: FP1-F7, F7-T7, T7-P7, P7-O1, FP1-F3, F3-C3, C3-P3, P3-O1, FP2-F4, F4-C4, C4-P4, P4-O2, FP2-F8, F8-T8, T8-P8, P8-O2, FZ-CZ, and CZ-PZ.

Raw EEG signals are converted from volts to microvolts by multiplying by 10^6. No additional filtering or artefact removal is applied.

EEG recordings are segmented using an 8-second sliding window ($W = 2048$ samples) with 50% overlap. A window is labelled ictal if more than 50% of its samples correspond to seizure activity, and interictal if no seizure samples are present. Ambiguous windows are discarded. To reduce class imbalance, interictal

Table 1. Per-patient ictal window counts (three-column layout, 24 patients total). Interictal windows are downsampled to match the ictal count.

Patient	Ictal	Total	Patient	Ictal	Total	Patient	Ictal	Total
chb01	114	228	chb09	70	140	chb17	75	150
chb02	44	88	chb10	112	224	chb18	80	160
chb03	103	206	chb11	202	404	chb19	61	122
chb04	96	192	chb12	251	502	chb20	75	150
chb05	141	282	chb13	137	274	chb21	52	104
chb06	40	80	chb14	44	88	chb22	52	104
chb07	82	164	chb15	506	1012	chb23	107	214
chb08	231	462	chb16	24	48	chb24	237	474

windows are randomly downsampled to match the number of ictal windows for each patient. Table 1 summarises the final window counts.

3.2 Strict Hold-Out Protocol

To prevent data leakage, each patient's balanced dataset is divided into a pool split (75%) and a strict hold-out split (25%) using stratified sampling. The pool split is used only for spectral fingerprint extraction, scaler fitting, clustering, and model training, while the hold-out split is reserved exclusively for evaluation and fine-tuning.

Let $\mathcal{D}_p = \mathcal{D}_p^{pool} \cup \mathcal{D}_p^{ho}$ denote the dataset of patient p, where $\mathcal{D}_p^{pool} \cap \mathcal{D}_p^{ho} = \emptyset$. The StandardScaler is fitted only on pool data and applied to hold-out data without refitting, ensuring complete separation between training and evaluation. For fine-tuning experiments, the hold-out split is further divided into a support set (30%) and a query set (70%). Zero-shot and fine-tuned performance are both evaluated on the same query windows.

3.3 Spectral Fingerprinting and Patient Clustering

To capture inter-patient seizure variability, patients are grouped according to their ictal spectral characteristics. For each patient, a spectral fingerprint $\mathbf{f}_p \in \mathbb{R}^{B \times C}$ is extracted from ictal windows in the pool split, where $B = 6$ frequency bands and $C = 18$ EEG channels. Band power is computed using Welch's method (24) across the delta, theta, alpha, beta, gamma, and high-gamma bands. The resulting 108-dimensional feature vector is log-transformed and standardised using a StandardScaler fitted only on pool data.

Patients are then clustered using K-means with $k \in 2, 3, 4, 5$. Silhouette scores of 0.379, 0.285, 0.233, and 0.237 were obtained for $k = 2, 3, 4$, and 5, respectively; therefore, $k = 2$ was selected, yielding two clusters containing 11 and 13 patients. Principal component analysis (PCA) is used to visualise the resulting cluster structure (Fig. 2).

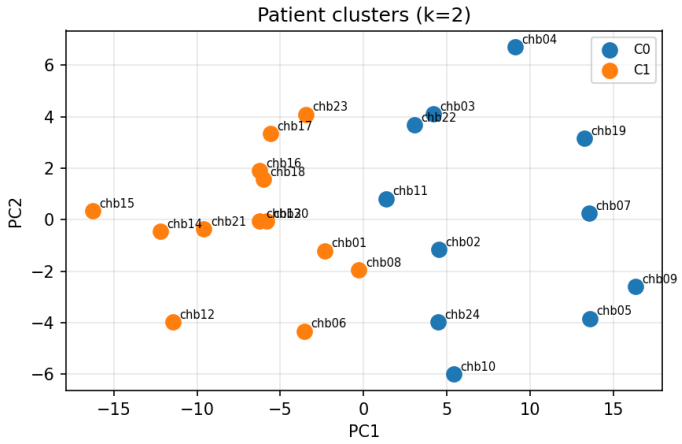

Fig. 2. PCA visualisation of patient spectral fingerprints clustered using K-means ($k = 2$). Blue and orange points represent Cluster 0 and Cluster 1, respectively.

Cluster centroids are estimated using data from the entire patient cohort, including the target patient. Consequently, the evaluation avoids query-level leakage but does not constitute a fully patient-independent deployment setting, where centroids would be fixed prior to the arrival of unseen patients.

3.4 Lightweight TCN Architecture

The proposed model, called LightweightTCN, is a lightweight temporal convolutional network designed for efficient seizure detection from multi-channel EEG signals. The architecture contains four main components: a learned channel gate, a stem convolution layer, multiple dilated residual blocks with squeeze-and-excitation attention, and a fully connected classification head.

The stem begins with a single Conv1D layer using a kernel size of 11 and 32 output channels. This is followed by batch normalisation, Swish activation, and spatial dropout. The main network consists of five depthwise-separable residual blocks with dilation factors $d \in \{1, 2, 4, 8, 16\}$ and channel sizes $\{64, 128, 128, 192, 256\}$. The increasing dilation factors expand the receptive field without significantly increasing the number of parameters. This enables the model to learn both short seizure spikes and longer rhythmic seizure patterns within an 8-second EEG segment.

The classification head uses global average pooling followed by two fully connected layers (256→128→1) with Swish activation and dropout rates of 0.4 and 0.2. The complete model contains approximately 1.2 million parameters, making it suitable for resource-limited devices and real-time applications.

Learned Channel Gate Seizure patterns vary across patients, meaning that some EEG channels may be more important than others depending on the pa-

tient. To handle this variability, a learned channel gate is used as the first stage of the model. A trainable scalar w_c is assigned to each of the $C = 18$ EEG channels. During inference, the gate computes a softmax-normalised weight for each channel:

$$g_c = \frac{\exp(w_c)}{\sum_{c'=1}^{C} \exp(w_{c'})} \quad c = 1, \ldots, C \tag{1}$$

The input EEG signal is then scaled using these weights:

$$\tilde{x}_{c,t} = g_c \cdot x_{c,t} \tag{2}$$

where $x_{c,t}$ represents the signal amplitude of channel c at time t. The softmax operation ensures that all channel weights are positive and sum to one, making them easy to interpret as channel importance scores.

Unlike fixed or manually selected channels, the learned gate is fully differentiable and trained together with the rest of the network. Since each cluster model is trained using different groups of patients, the gate learns cluster-specific channel importance patterns that reflect the common seizure activity of that patient group.

Squeeze-and-Excitation Blocks Each depthwise-separable residual block includes a squeeze-and-excitation (SE) module (11) to improve feature-level channel attention. Let $\mathbf{F} \in \mathbb{R}^{C \times T}$ represent the feature map produced by the convolution layers. The SE module first applies global average pooling across time:

$$z_c = \frac{1}{T} \sum_{t=1}^{T} F_{c,t} \tag{3}$$

The pooled features are then passed through a two-layer bottleneck structure:

$$\mathbf{s} = \sigma(W_2\, \delta(W_1\, \mathbf{z})) \tag{4}$$

where δ represents the Swish activation function, σ is the sigmoid function, and $W_1 \in \mathbb{R}^{r \times C}$, $W_2 \in \mathbb{R}^{C \times r}$ with reduction ratio $r = C/8$.

Finally, the generated attention weights are used to scale the feature channels:

$$\tilde{F}_{c,t} = s_c \cdot F_{c,t} \tag{5}$$

This mechanism allows the network to reduce the influence of less useful feature channels while strengthening seizure-related features. The SE module therefore provides feature-level attention, which complements the input-level attention provided by the learned channel gate.

3.5 Cluster-Specific Training

Instead of using a single global model or a fully personalised model per patient, we adopt a cluster-specific training strategy that balances generalisation and specialisation. For each target patient p, a cluster model $\mathcal{M}_p$ is trained using the pool splits of peer patients within the same cluster, while excluding patient p to avoid data leakage.

To improve robustness, an additional $\alpha = 0.25$ fraction of randomly sampled windows from out-of-cluster patients is added during training. This cross-cluster regularisation helps the model generalise better to patients near cluster boundaries. Training sets contain 1,599-1,950 windows, depending on cluster size, with 15% used for validation.

All models use the LightweightTCN architecture and are trained with the Adam optimiser ($\eta = 10^{-4}$), binary cross-entropy loss, cosine annealing learning rate scheduling ($\eta_{min} = 10^{-6}$), and gradient clipping at norm 1.0. Training runs for up to $E_{max} = 30$ epochs with early stopping patience of 15 epochs based on validation F1-score. To avoid patient imbalance, each patient's contribution is capped at 267 windows.

3.6 Zero-Shot and Fine-Tuned Evaluation Protocol

During evaluation, the cluster model $\mathcal{M}_p$ is applied directly to the held-out query set of patient p without using target-patient data for training, normalisation, or threshold selection. This is referred to as zero-shot inference.

Cluster assignment requires at least one previously recorded and labeled seizure event. Therefore, "zero-shot" in this study denotes the absence of patient-specific training rather than the absence of prior seizure observations. Zero-shot performance serves as the baseline before adaptation.

For patient-specific adaptation, the hold-out data are divided into a support set (30%, minimum 10 windows) and a query set (70%). The support set is used for fine-tuning, while evaluation is performed exclusively on the query set.

During fine-tuning, the stem convolution layer is frozen, while the remaining layers are updated using learning rates of $\eta_{ft} = 2 \times 10^{-5}$ for intermediate layers and $\eta_{head} = 10^{-4}$ for the classification head. Training is performed for up to $E_{ft} = 15$ epochs with early stopping (patience = 8), except when the support set contains fewer than 16 samples. This strategy enables rapid adaptation to patient-specific seizure patterns while preserving the general representations learned during cluster training.

4 Evaluation Protocol

A fixed decision threshold of $\tau = 0.35$ is applied to the sigmoid output in both zero-shot and fine-tuned settings. The threshold was chosen to prioritise sensitivity, reflecting the clinical importance of minimising missed seizures. Performance is evaluated on each patient's hold-out query set using accuracy (ACC), precision (PREC), sensitivity (SEN), specificity (SPEC), negative predictive value

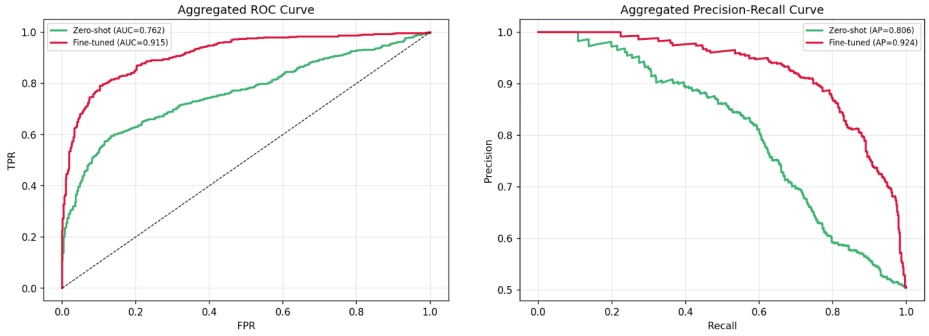

Fig. 3. Aggregated ROC and Precision-Recall (PR) curves across all 24 patients for zero-shot and fine-tuned seizure detection models.

(NPV), F1-score (F1), area under the receiver operating characteristic curve (AUC), area under the precision–recall curve (AUPRC), Matthews correlation coefficient (MCC), and Cohen's kappa. Post-hoc threshold calibration using F1-optimal, FAR-constrained, and Youden-J criteria on support sets did not improve aggregate results (mean F1: 0.825→0.811; FAR: 60.0→87.7/hr), indicating that calibration from small patient-specific support sets is unstable and that a fixed threshold is a reasonable choice.

To obtain clinically meaningful false alarm estimates, a 30-second persistence filter is applied. Because overlapping EEG windows can inflate per-window false alarm rates, an alarm is generated only when four consecutive windows exceed the threshold. With 8-second windows and a 4-second stride, this corresponds to approximately 30 seconds of sustained abnormal activity, providing a more realistic estimate of clinically actionable false alarms.

5 Results & Discussion

5.1 Zero-Shot, Fine-Tuned, and Clinical FAR Results

Table 2 summarises the per-patient performance on the strict hold-out split. In the zero-shot setting, the cluster model was directly applied to unseen patients without using target-patient data. For fine-tuning, 30% of the hold-out data (minimum 10 windows) was used as a support set to adapt the model before evaluation on the remaining query set.

Figure 3 presents the aggregated ROC and Precision-Recall curves across all 24 patients for both the zero-shot and fine-tuned models.Fine-tuning substantially improves discriminative performance, increasing the ROC-AUC from 0.762 to 0.915 and the PR-AUC/AP from 0.806 to 0.924. The improved PR curve further demonstrates enhanced seizure sensitivity while maintaining stable precision across different recall levels. These improvements align with the observed gains in F1-score, sensitivity, and MCC, indicating that patient-specific adap-

Table 2. Per-patient zero-shot (ZS), fine-tuned (FT), and clinical FAR results on the strict hold-out split. Raw FAR and Clinical FAR (after persistence filtering) are reported for the fine-tuned model. Hard patients (chb13–15) are marked with †.

Patient	C	ZS-F1	FT-F1	ΔF1	ZS-AUC	FT-AUC	FT-MCC	Raw FAR	Clinical FAR	Reduction
chb01	1	0.976	0.974	−0.001	0.997	0.997	0.951	0.00	0.00	∞
chb02	0	1.000	1.000	+0.000	1.000	1.000	1.000	0.00	0.00	∞
chb03	0	0.857	0.947	+0.090	0.889	0.944	0.892	24.32	50.00	0.5×
chb04	0	0.774	0.812	+0.038	0.934	0.931	0.652	52.94	0.00	∞
chb05	0	0.676	0.873	+0.197	0.968	0.984	0.735	108.00	108.00	1.0×
chb06	1	0.600	0.727	+0.127	0.680	0.840	0.408	180.00	0.00	∞
chb07	0	0.968	0.968	+0.000	1.000	1.000	0.933	31.03	64.29	0.5×
chb08	1	0.835	0.833	−0.002	0.883	0.905	0.659	87.80	65.85	1.3×
chb09	0	0.722	1.000	+0.278	0.955	1.000	1.000	0.00	0.00	∞
chb10	0	0.974	0.974	+0.000	0.997	0.995	0.951	0.00	0.00	∞
chb11	0	0.933	0.959	+0.026	0.994	0.997	0.916	25.35	0.00	∞
chb12	1	0.196	0.696	+0.500	0.714	0.819	0.371	151.69	40.91	3.7×
chb13†	1	0.250	0.577	+0.327	0.433	0.530	0.101	220.41	37.50	5.9×
chb14†	1	0.222	0.222	+0.000	0.278	0.444	−0.192	150.00	0.00	∞
chb15†	1	0.065	0.809	+0.745	0.720	0.872	0.630	70.79	30.34	2.3×
chb16	1	0.444	0.444	+0.000	0.639	0.639	0.192	75.00	0.00	∞
chb17	1	0.933	0.933	+0.000	0.989	0.989	0.860	66.67	0.00	∞
chb18	1	0.727	0.783	+0.055	0.939	0.918	0.688	0.00	0.00	∞
chb19	0	1.000	0.952	−0.048	1.000	1.000	0.909	0.00	0.00	∞
chb20	1	0.667	0.667	+0.000	0.764	0.885	0.570	0.00	0.00	∞
chb21	1	0.769	0.824	+0.054	0.938	0.922	0.630	112.50	112.50	1.0×
chb22	0	0.933	1.000	+0.067	1.000	1.000	1.000	0.00	0.00	∞
chb23	1	0.973	0.973	+0.000	0.978	0.983	0.949	0.00	0.00	∞
chb24	0	0.567	0.769	+0.203	0.745	0.858	0.577	64.29	0.00	∞
Mean		0.711	0.822	+0.111	0.851	0.894	0.683	59.20	21.22	—

tation enhances seizure detection capability while maintaining stable precision and specificity without introducing excessive false positives.

To estimate clinically realistic false alarm rates, a 30-second persistence filter was applied, where an alarm was triggered only if four consecutive windows exceeded the decision threshold. This reduced the mean false alarm rate from 59.20 to 21.22 false alarms per hour. Sixteen of the 24 patients achieved a clinical FAR of zero after filtering. However, persistence filtering did not uniformly reduce false alarms across all patients; in some cases, isolated false-positive windows were consolidated into alarm events, resulting in unchanged or slightly increased FAR. Median clinical FAR (0.0/hr) is substantially lower than the mean (21.22/hr), indicating the mean is driven by a few high-FAR patients (e.g., chb05, chb21); cases where filtering failed to reduce FAR involved sustained rather than isolated false positives, which persistence filtering cannot suppress.

5.2 Hard Patient Analysis

Three patients (chb13, chb14, and chb15) were identified as hard cases because of poor zero-shot performance (2). Among them, chb15 showed the largest improvement after fine-tuning, with F1-score increasing from 0.065 to 0.809, and FT-MCC reaching 0.630 . Patient chb13 also improved, although its MCC remained low, indicating continued class confusion.

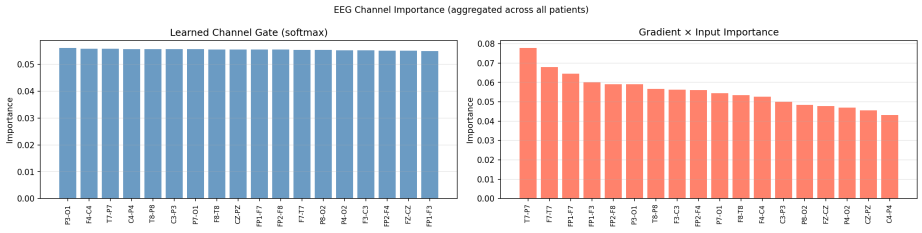

Fig. 4. Comparison between learned channel gate weights and Gradient × Input saliency-based EEG channel importance aggregated across all patients.

In contrast, chb14 showed no improvement after fine-tuning, likely because of its very small hold-out set and near-random AUC value. Hard patients also produced elevated raw false alarm rates, ranging from 70.79/hr (chb15) to 220.41/hr (chb13). However, persistence filtering substantially reduced the FAR for chb12 and chb13, suggesting that many false positives occurred as short isolated bursts.

5.3 Channel Importance Analysis

Two methods were used to analyse EEG channel importance. The learned channel gate assigned nearly uniform weights across all 18 channels (range 0.0556–0.0561), suggesting that the model relies on broad spatial information rather than a few specific channels.

Gradient×Input saliency analysis produced more distinct patterns. T7-P7 showed the highest saliency score, followed by F7-T7 and FP1-F7, highlighting the left temporal region as the most informative area for seizure detection. This observation is consistent with the prevalence of temporal lobe epilepsy in the CHB-MIT dataset.

Figure 4 compares the learned channel gate weights with the Gradient × Input saliency-based channel importance scores aggregated across all patients. The near-uniform distribution of the gate weights suggests that the channel gate primarily functions as a regularisation mechanism during training rather than strongly prioritising specific electrodes. In contrast, the saliency rankings reveal that only a subset of EEG channels contributes most significantly during inference, indicating that these electrodes carry stronger discriminative information for seizure detection.

5.4 Failure cases.

Three patients (chb13, chb14, chb15) show poor zero-shot performance, and chb14 is entirely unresponsive to fine-tuning ($F1 = 0.222$, $MCC = -0.192$ before and after adaptation). The AUC of 0.278 for chb14 under zero-shot inference is below chance, indicating the cluster prior actively miscalibrates predictions for this patient. With only 22 hold-out windows and 10 support windows, the fine-tuning signal is too weak to overcome this miscalibration. These cases suggest

that a minimum data requirement should gate whether fine-tuning is attempted, and that an outlier detection mechanism at cluster assignment time would be clinically valuable.

5.5 State-of-the-arts Comparison

Table 3 presents prior cross-patient seizure detection results for contextual reference only, as differences in evaluation protocols, reported metrics, and class balancing strategies make direct numerical comparisons unreliable. Therefore, the main conclusions of this study are based on the internal comparison between the zero-shot and fine-tuned results in Table 2. Using a strictly leakage-free per-patient hold-out protocol, the proposed LightweightTCN achieves an F1-score of 82.2% and an AUC of 89.4% after fine-tuning, while still obtaining a zero-shot F1-score of 71.1% without using any target-patient data.

Table 3. Comparison with cross-patient seizure detection methods on CHB-MIT.

Reference	Method	ACC (%)	SEN (%)	F1 (%)	AUC (%)	Protocol
Rukhsar & Tiwari (18)	Lightweight Convolution Transformer	96.31	96.82	96.32	–	Pooled 75/25, segment
Zhang et al. (25)	Spiking Neural Network	–	90.46	–	96.86	LOO + fine-tune, segment
Ali et al. (2)	Random Forest	–	75.34	–	–	LOO, event-level
Wang et al. (23)	Privacy-preserving UDA	94.81	–	–	–	Cross-patient, segment
Ours (zero-shot)	LightweightTCN	–	–	71.1	85.1	Per-patient holdout, segment
Ours (fine-tuned)	LightweightTCN	–	–	82.2	**8**9.4	Per-patient holdout, segment

6 Conclusion & Future Work

In this work, a cluster-based seizure detection framework was presented for paediatric EEG that combines spectral fingerprinting, lightweight temporal convolutional networks, and patient-specific fine-tuning under a strictly leakage-free evaluation protocol. Patients were grouped into two clusters based on their ictal spectral characteristics, and a dedicated LightweightTCN model was trained for each cluster using only peer-patient data in a leave-one-patient-out setting.The proposed model, with approximately 1.2 million parameters, demonstrated that competitive seizure detection performance can be achieved within a resource-efficient architecture suitable for edge and real-time deployment. The trainable per-channel softmax gate and squeeze-and-excitation modules provided complementary spatial and feature-level attention.

Despite these contributions, several limitations remain. The framework was evaluated only on the CHB-MIT dataset, and generalisation to other EEG databases has not been verified. Patient chb14 remained unresponsive to fine-tuning, suggesting the need for data quality checks and outlier detection during cluster assignment. Furthermore, no ablation study was performed to isolate the contributions of the clustering and attention modules. The relatively small

patient cohort and limited size of some per-patient hold-out sets may also introduce variability in the reported performance estimates, which was not quantified through confidence intervals or repeated-run analyses. Future work will investigate adaptive clustering, transformer-based temporal modelling, component-wise ablations, statistical robustness analyses, and multi-dataset validation to improve generalisation and reliability.

Bibliography

[1] Ahmad, I., Wang, X., Javeed, D., et al.: Automated detection of epileptic seizures using machine learning: a systematic literature review. IEEE Access **10**, 27336–27361 (2022)

[2] Ali, E., Angelova, M., Karmakar, C.: Epileptic seizure detection using chb-mit dataset: The overlooked perspectives. Royal Society open science **11**(5) (2024)

[3] Alotaibi, Y., Al-Luhaybi, M., Alharbi, J., Majrashi, B., Alharbi, Z., Alhindi, M.: Hybrid cnn-tcn modeling of eeg temporal signals for epileptic seizure prediction. In: 2025 International Conference on Advanced Machine Learning and Data Science (AMLDS). pp. 371–377. IEEE (2025)

[4] Andrade, I., Teixeira, C., Pinto, M.: On the performance of seizure prediction machine learning methods across different databases: the sample and alarm-based perspectives. Frontiers in Neuroscience **18**, 1417748 (2024)

[5] Baghdadi, A., Fourati, R., Aribi, Y., Daoud, S., Dammak, M., Mhiri, C., Chabchoub, H., Siarry, P., Alimi, A.: A channel-wise attention-based representation learning method for epileptic seizure detection and type classification. Journal of Ambient Intelligence and Humanized Computing **14**(7), 9403–9418 (2023)

[6] Boonyakitanont, P., et al.: A review of feature extraction and performance evaluation in epileptic seizure detection using EEG. Biomedical Signal Processing and Control **57**, 101702 (2020)

[7] Cui, X., Cao, J., Lai, X., Jiang, T., Gao, F.: Cluster embedding joint-probability-discrepancy transfer for cross-subject seizure detection. IEEE Transactions on Neural Systems and Rehabilitation Engineering **31**, 593–605 (2022)

[8] Cui, X., Wang, T., Lai, X., Jiang, T., Gao, F., Cao, J.: Cross-subject seizure detection by joint-probability-discrepancy-based domain adaptation. IEEE Transactions on Instrumentation and Measurement **72**, 1–13 (2023)

[9] Einizade, A., Nasiri, S., Mozafari, M., Sardouie, S.H., Clifford, G.D.: Explainable automated seizure detection using attentive deep multi-view networks. Biomedical Signal Processing and Control **79**, 104076 (2023)

[10] Guttag, J.: CHB-MIT Scalp EEG Database. PhysioNet (Jun 2010). https://doi.org/10.13026/C2K01R, https://doi.org/10.13026/C2K01R, version 1.0.0

[11] Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)

[12] Hu, Z., Yi, R., Li, Y., Cui, F., Chen, Z., Chen, S., Liu, H., Zhou, H., Wan, Y., Feng, Y.: Multi-scale kernel and electrode attention network for eeg-based epileptic seizure detection. Computers in Biology and Medicine **205**, 111585 (2026)

[13] Huang, L., Zhou, K., Chen, S., Chen, Y., Zhang, J.: Automatic detection of epilepsy from eegs using a temporal convolutional network with a self-attention layer. BioMedical Engineering OnLine **23**(1), 50 (2024)
[14] Jemal, I., Abou-Abbas, L., Henni, K., Mitiche, A., Mezghani, N.: Domain adaptation for eeg-based, cross-subject epileptic seizure prediction. Frontiers in neuroinformatics **18**, 1303380 (2024)
[15] Lim, Z.Y., Pang, Y.H., Ooi, S.Y., Khoh, W.H., Chew, Y.J.: Neurofed-lighttcn: Federated lightweight temporal convolutional networks for privacy-preserving seizure detection in eeg data. Applied Sciences **15**(17), 9660 (2025)
[16] Raab, D., Theissler, A., Spiliopoulou, M.: Xai4eeg: spectral and spatio-temporal explanation of deep learning-based seizure detection in eeg time series. Neural Computing and Applications **35**(14), 10051–10068 (2023)
[17] Rao, S., Liu, M., Huang, Y., Yang, H., Liang, J., Lu, J., Niu, Y., Wang, B.: Anchoring temporal convolutional networks for epileptic seizure prediction. Journal of Neural Engineering **21**(6), 066008 (2024)
[18] Rukhsar, S., Tiwari, A.K.: Lightweight convolution transformer for cross-patient seizure detection in multi-channel eeg signals. Computer Methods and Programs in Biomedicine **242**, 107856 (2023)
[19] Shafiezadeh, S., Duma, G.M., Mento, G., Danieli, A., Antoniazzi, L., Del Popolo Cristaldi, F., Bonanni, P., Testolin, A.: Calibrating deep learning classifiers for patient-independent electroencephalogram seizure forecasting. Sensors **24**(9), 2863 (2024)
[20] Shyu, K.K., Huang, S.C., Lee, L.H., Lee, P.L.: Less parameterization inception-based end to end cnn model for eeg seizure detection. Ieee Access **11**, 49172–49182 (2023)
[21] Thijs, R.D., Surges, R., O'Brien, T.J., Sander, J.W.: Epilepsy in adults. The Lancet **393**(10172), 689–701 (2019)
[22] Wang, S., Feng, H., Lv, H., Nie, C., Feng, W., Peng, H., Zhang, L., Zhao, Y.: Cross-subject seizure detection via unsupervised domain-adaptation. International Journal of Neural Systems **34**(10), 2450055 (2024)
[23] Wang, S., Lv, H., Feng, H., Peng, H., Feng, W., Nie, C., Zhao, Y.: Cross patient seizure detection via unsupervised domain adaptation based on uncertainty estimation. Applied Intelligence **55**(12), 879 (2025)
[24] Welch, P.: The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. IEEE Transactions on audio and electroacoustics **15**(2), 70–73 (1967)
[25] Zhang, Z., Xiao, M., Ji, T., Jiang, Y., Lin, T., Zhou, X., Lin, Z.: Efficient and generalizable cross-patient epileptic seizure detection through a spiking neural network. Frontiers in Neuroscience **17**, 1303564 (2024)