ICU - EEG Pattern Detection by a Convolutional Neural Network

Degano, Giulio; Quintard, Hervé; Kleinschmidt, Andréas Karl; Francini, Nikita; Sarbu, Oana E.; De Stefano, Pia

How to cite

DEGANO, Giulio et al. ICU - EEG Pattern Detection by a Convolutional Neural Network. In: Annals of clinical and translational neurology, 2025. doi: 10.1002/acn3.70164

This publication URL: https://archive-ouverte.unige.ch/unige:187419
Publication DOI: 10.1002/acn3.70164

Annals of Clinical and Translational Neurology

RESEARCH ARTICLE OPEN ACCESS

ICU-EEG Pattern Detection by a Convolutional Neural Network

Giulio Degano[1,2] | Hervé Quintard[1] | Andreas Kleinschmidt[2] | Nikita Francini[1,2] | Oana E. Sarbu[1,2] | Pia De Stefano[1,2]

[1]Neuro-Intensive Care Unit, Department of Intensive Care, University Hospital of Geneva, Geneva, Switzerland | [2]Neurology Unit, Department of Clinical Neurosciences, University Hospital of Geneva, Geneva, Switzerland

Correspondence: Giulio Degano (giulio.degano@hug.ch)

Received: 28 February 2025 | **Revised:** 7 July 2025 | **Accepted:** 21 July 2025

Funding: The authors received no specific funding for this work.

Keywords: deep learning | electroencephalography | intensive care units | seizures

ABSTRACT

Objective: Patients in the intensive care unit (ICU) often require continuous EEG (cEEG) monitoring due to the high risk of seizures and rhythmic and periodic patterns (RPPs). However, interpreting cEEG in real time is resource-intensive and heavily relies on specialized expertise, which is not always available. This study introduces a lightweight convolutional neural network (CNN) to automatically detect key EEG patterns, including seizures and RPPs.

Methods: We classified time–frequency spectrograms of EEG data from the Harmful Brain Activity Classification challenge, including 1950 patients. We tested our model on a subset of this dataset and a small independent cohort of ICU patients with epileptic seizures from the Geneva University Hospital.

Results: Our model showed good performance metrics on the open-source data with an AUROC score of 93% for SZ, 91% for lateralized PD, 94% for generalized PD, 87% for lateralized RDA, 89% for generalized RDA, and 88% for others. The evaluation with the Geneva University Hospital dataset also demonstrated strong temporal detection capabilities, showing an false positive rate (FPR) of 22%, 20%, and 21% at 50 s, 30 s, and 20 s before seizure onset, and a true positive rate (TPR) of 76%, 84%, and 89% at 20 s, 30 s, and 50 s after seizure onset.

Interpretation: This study presents a lightweight CNN model capable of detecting critical EEG patterns in ICU patients with minimal preprocessing. Moreover, the model's design provides reliable detection of ICU-EEG epileptic patterns shortly after their onset. These features underscore the model's potential to enhance timely EEG monitoring in resource-limited and advanced clinical contexts.

1 | Introduction

Patients in the intensive care unit (ICU) are often challenging to monitor clinically, as both the underlying pathological picture and sedative medications can obscure clinical signs [1, 2]. Furthermore, the majority of seizures and status epilepticus occurring in the ICU are nonconvulsive and go unnoticed without continuous EEG (cEEG) monitoring [3]. In this context, cEEG has emerged as a valuable tool for detecting seizures and other harmful epileptiform EEG findings, including rhythmic and periodic patterns (RPPs) [4], while also providing pivotal insight into patient management and prognosis [5–7]. However, cEEG interpretation requires specialized expertise that is not always available, especially in real time, and inter-rater reliability can vary substantially across different EEG patterns [8]. Moreover, the resources needed for continuous monitoring and rapid interpretation are substantial, making the current model difficult to sustain [9, 10].

There is, therefore, an urgent need for methods that can automatically detect harmful EEG patterns suggestive of seizures or clinical deterioration without sacrificing accuracy. Artificial intelligence (AI) holds promise in addressing this gap by enabling faster, more consistent EEG interpretation, reducing the demand on human experts [11, 12].

Recent research has focused on seizure detection [11, 12], on recognition of interictal discharges [13] or abnormal recordings [14], with most of these models exploiting the precision given by convolutional neural networks [15] (CNN). However, few models have been optimized to recognize key ICU-specific EEG patterns, such as RPPs, electrographic seizures, and status epilepticus as defined by the 2021 American Clinical Neurophysiology Society EEG terminology [4], which have both diagnostic and therapeutic implications in critically ill patients [16–18].

In this work, we propose a lightweight CNN based on the EfficientNet [19] architecture that processes spectral EEG data with minimal preprocessing. This model aims to accurately detect seizures, periodic discharges (PD), and rhythmic delta activity (RDA), while reducing computational costs compared to typical 2D CCN models [17]. Additionally, we explore the network's internal representations to assess its ability to provide timely and accurate EEG interpretations, critical for prompt therapeutic interventions.

2 | Methods

2.1 | Cohort and Annotation

The cohort of patients was extracted from the Harmful Brain Activity Classification (HBAC) challenge in Kaggle (https://kaggle.com/competitions/hms-harmful-brain-activity-classification). The data contained the EEG traces of 1950 patients collected by the authors of the challenge. The EEG segments used in this competition were annotated by experts from the Critical Care EEG Monitoring Research Consortium (CCEMRC). A team of annotators reviewed each EEG segment. The experts examined 50-s EEG samples along with corresponding spectrograms covering a 10-min window, and they labeled the central 10 s of each segment. The EEG data outside the 10-s windows was shown to annotators in order to give them context regarding the specifics of the pattern. They were tasked with identifying six specific patterns of interest: seizure (SZ), generalized PD (GPD), lateralized PD (LPD), lateralized RDA (LRDA), generalized RDA (GRDA), or "Other" (which includes patterns of artifacts).

As a further generalization step of the classifier's performance, we retrospectively selected 15 subjects recorded from the ICU at the Geneva University Hospital (HUG) who suffered from one or multiple epileptic seizures. We chose seizure over the other EEG patterns because they provide more precise time-locked annotations with specific onset. Annotation was performed by board-certified ICU-EEG neurologists (PDS) and board-certified EEG technologists (NF, OES), ensuring accuracy in the data used for testing, and all onset and offset times were re-reviewed by PDS, NF, or OES with controversies resolved through consensus meetings.

2.2 | Data Selection Ad Preprocessing

A selection of the training data was first performed on the dataset by selecting only the entries that contained at least 20 or more votes. This procedure allowed filtering only the data with higher quality of annotation and obtaining an expert distribution of probabilities that had enough variability [15]. We then selected the central 50 s of the EEG pattern that also included the window of annotation of the raters. Two montages were created and analyzed: a standard double banana bipolar montage extracted across the 19 EEG channels and a transverse montage computed to further enhance the lateralization of EEG patterns. The 50-s EEG data segment was then processed to prepare it for analysis. For each chain generated by the montage, a spectrogram was computed across the respective channels with a window size of 128 time points, a hop parameter at 30% of the window length, and a frequency range between 0 and 20 Hz. Fast Fourier Transform (FFT) window length was set to 1024 for spectral analysis. The resulting spectrogram of each bipolar channel with dimension 128 (time points) × 256 (frequencies) was averaged across the respective chain. All the parameters for the spectral preprocessing were selected as a compromise between computational efficiency and resolution. For example, the window size was chosen to control the dimensionality of the spectral representations across the eight chains and consequently the input size to the CNN. Similarly, hop size and frequency range were set to reduce processing time, as they did not impact CNN decoding performance in preliminary analyses. This process thus generated, for each participant, a total of eight spectrograms (four for each chain in the double banana montage and four for each chain in the traversal montage). Finally, each spectrogram was concatenated to generate a single image used as input for the decoding.

2.3 | Architecture

To classify the preprocessed EEG traces in the dataset, we implemented a set of CNNs (Figure 1A). We adopted the choice of deep learning methods as they have been shown to be superior to classical machine learning techniques in combination with hand-crafted features [20]. For this work, we chose an EfficientNet-B0 network which is a lightweight CNN and a variant of the EfficientNet family designed to be a more balanced network, offering a good trade-off between model size, computational efficiency, and accuracy [19]. After a convolutional layer, the architecture is composed of seven stages of Mobile Inverted Bottleneck Convolution (MBConv). The first one uses a single expansion ratio (MBConv1) then the following six stages have a constant expansion ratio fixed at six (MBConv6). The last layers of the network are a fully connected layer with SoftMax activation with six dimensions (one for each class). Overall, the number of parameters and the computational time for the training phase are drastically reduced compared to recent CNNs used for EEG pattern classification like U-Nets [17] making it suitable even for low-powered laptops used in hospital settings.

2.4 | Training

The classifier was trained by using a stratified group fivefold strategy (Figure 1B). Moreover, the learning rate was changed across

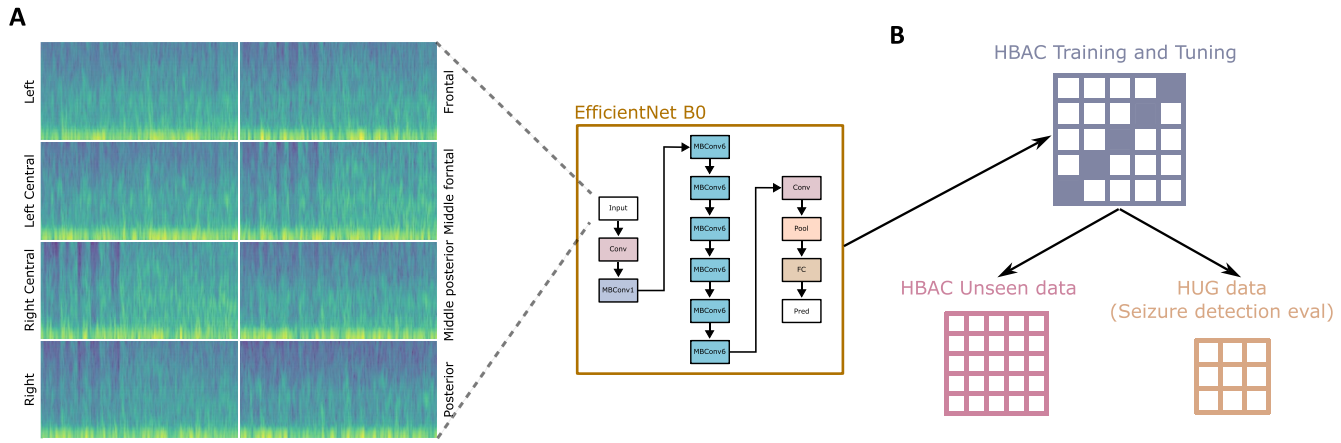

FIGURE 1 | EEG decoding pipeline. (A) The EEG spectrogram of the double banana and transverse chain were first extracted to create a single image or input and in the orange box the EfficientNet architecture used for the decoding of the EEG data. (B) The EfficientNet was trained on the open-source dataset in a fivefold cross-validated fashion. Then, the model was tested by assessing its performance in EEG pattern detection on a help out set (pink figure). Finally, we assessed the temporal performance of the model on our Geneva University Hospital (HUG) seizure dataset.

training epochs using a scheduler (LearningRateScheduler in TensorFlow). We thus optimized the network by setting a decaying learning rate function of $10^{-\text{EpochNum}}$ with an initial warm-up of three epochs fixed at 0.001. The final models were trained on a total of 1400 subjects with 6431 different EEG patterns (19% SZ, 17% LPD, 11% GPD, 12% GRDA, 8% LRDA, 33% other).

2.5 | Objective Function During Training

The objective function was selected to minimize the probability distribution across classes defined by expert annotations. We used Kullback–Leiber (KL) divergence to minimize the distance between this latter probability and the SoftMax output of the EfficientNet model. The KL divergence, also known as relative entropy, can be expressed mathematically as follows:

$$D_{\text{KL}}(P \parallel Q) = \sum_{x} P(X) \log\left(\frac{P(X)}{Q(X)}\right) \tag{1}$$

where $P(X)$ is the ground truth or the expert probability distribution, and Q(X) is the output of the network. We used Adam for stochastic optimization during the training phase of the EfficientNet [21].

2.6 | Testing Set in HBAC—Soft Labels

After the training, we tested the models' performance on the held out data of the HBAC dataset. The testing set scores were computed by averaging across the SoftMax output of the five ensembles of models obtained from cross-validation.

To assess the model's performance on the unseen data in HBAC data, we computed the receiver-operating characteristic (ROC) and precision-recall (PR) curves for each class in the annotation labels. The area under the ROC (AUROC) and PR (AURC) curves was also computed to assess the final score of the models across different threshold levels. A 95% confidence interval was extracted via bootstrapping with 5000 resamples.

Finally, we assessed model calibration to measure how well predicted probabilities match actual outcomes [22]. The Expected Calibration Error (ECE) is computed as

$$\text{ECE} = \sum_{m=1}^{M} \left(\frac{|B_m|}{n}\right) * \left|\text{acc}\left(B_m\right) - \text{conf}\left(B_m\right)\right| \tag{2}$$

where n is the number of samples, B_m is the probabilities bins, and $\text{acc}\left(B_m\right)$ and $\text{conf}\left(B_m\right)$ are the accuracy and confidence within each bin respectively.

2.7 | Testing Set in HBAC—Hard Labels

Since the model outputs six different probabilities, we implemented a classifier to convert these probabilities into a single decision or hard label, aligning with the standard reporting format used by clinicians. To achieve this, we fit a linear support vector classifier (SVC) during CNN training, using the SoftMax probabilities from each of the five models as input and the consensus labels as output. The SVC was set with an L2 penalty, a regularization parameter of $C = 1$, and was trained using a one-versus-rest classification strategy. For the hard labels classification, we computed the true positive rate (TPR) and false positive rate (FPR) for each class as:

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \tag{3,4}$$

2.8 | Embedding Exploration

To gain a deeper understanding of the network's performance, particularly in distinguishing RPPs, we conducted an embedding exploration analysis. The idea behind this analysis was to examine how the network internally organizes EEG patterns—specifically, whether similar patterns are grouped (clustered) together in its learned representation space. To do this, we assessed the Euclidean distances between pairs of classes at the network's logit layer. In other words, we evaluated how close different EEG patterns are to each other before the network computes class

probabilities. This helped reveal which EEG patterns the network considers similar or different and highlighted potential confusions or overlaps. Further, qualitative analysis was computed by extracting the 2D Isomap representation of the embeddings by extracting the second-to-last representation layer (after pooling).

We tested the pairwise distances in a mixed two-way ANOVA (distances of the classes against the representations of GPD, GRDA, LRDA, and LPD) with the between factor being the two classes of PD and two classes of RDA, while the within factor being the five EfficientNets trained on each fold. Post hoc pairwise testing was computed to assess the difference in distance across all the models. Multiple comparisons were taken into account with the default Tukey method of the emmeans package [23]. All the statistics were computed in R studio (R version 4.4.0).

2.9 | Temporal Classification of Seizures in the HUG Dataset

To evaluate the performance of EfficientNet over time, we tested its capacity to assign the pattern to the seizure class using epileptic seizures from 15 patients in the HUG dataset. We focused on seizures as a pattern of interest due to their distinct onset. This allowed us to evaluate the delay of the network's estimates with respect to a clear onset marked by a consensus of clinicians. Detection classification was made by hard labeling the output probabilities of the EfficientNet at 50, 30, and 20 s before and after the seizure onset.

To assess the network's ability to calibrate for a specific pattern, we analyzed the performance of both an SVC classifier and a fine-tuned seizure alarm detector. The alarm detector was calibrated during cross-validation by testing 10 different threshold levels, allowing us to evaluate fluctuations in TPR and FPR across fivefold. The threshold was selected by balancing TPR and FPR, with a priority given to sensitivity due to its clinical significance in seizure detection. Ultimately, the threshold with the highest TPR that did not exceed 20% FPR across all fivefold was chosen.

Finally, we inspected the activation map of the EfficientNet model used for classifying the different EEG patterns. To do that, we employed Gradient-weighted Class Activation Mapping (Grad-CAM) [24]. Here, Grad-CAM enabled us to visualize the model's weights after convolution by highlighting the regions of the input spectrogram that are important for classifying the EEG pattern. These maps provided us with a qualitative understanding of how the model interprets the temporal and spectral features of EEG data as well as produces a map for the localization of the seizures.

3 | Results

3.1 | Testing Set in HBAC

The performance of the HBAC testing set can be seen in Figure 2. The sample size was composed of 918 patterns from 201 subjects (19% SZ, 18% LPD, 11% GPD, 11% GRDA, 8% LRDA, 33% other). The EfficientNet model had an AUROC score of 93.52 (94.18, 92.86)% for SZ, 91.48 (92.19, 90.77)% for LPD, 93.67 (94.31, 93.02)% for GPD, 87.00 (87.93, 86.06)% for LRDA, 89.15 (89.89, 88.42)% for GRDA, and 88.16 (88.83, 87.49)% for other. The score for the AUPR was instead of 83.62 (84.96, 82.28)% for SZ, 73.94 (75.79, 72.10)% for LPD, 72.50 (74.59, 70.41)% for GPD, 29.36 (31.26, 27.46)% for LRDA, 52.58 (55.02, 50.15)% for GRDA, and 78.71 (79.92, 77.50)% for other. Model calibration showed an ECE score of 2.70%.

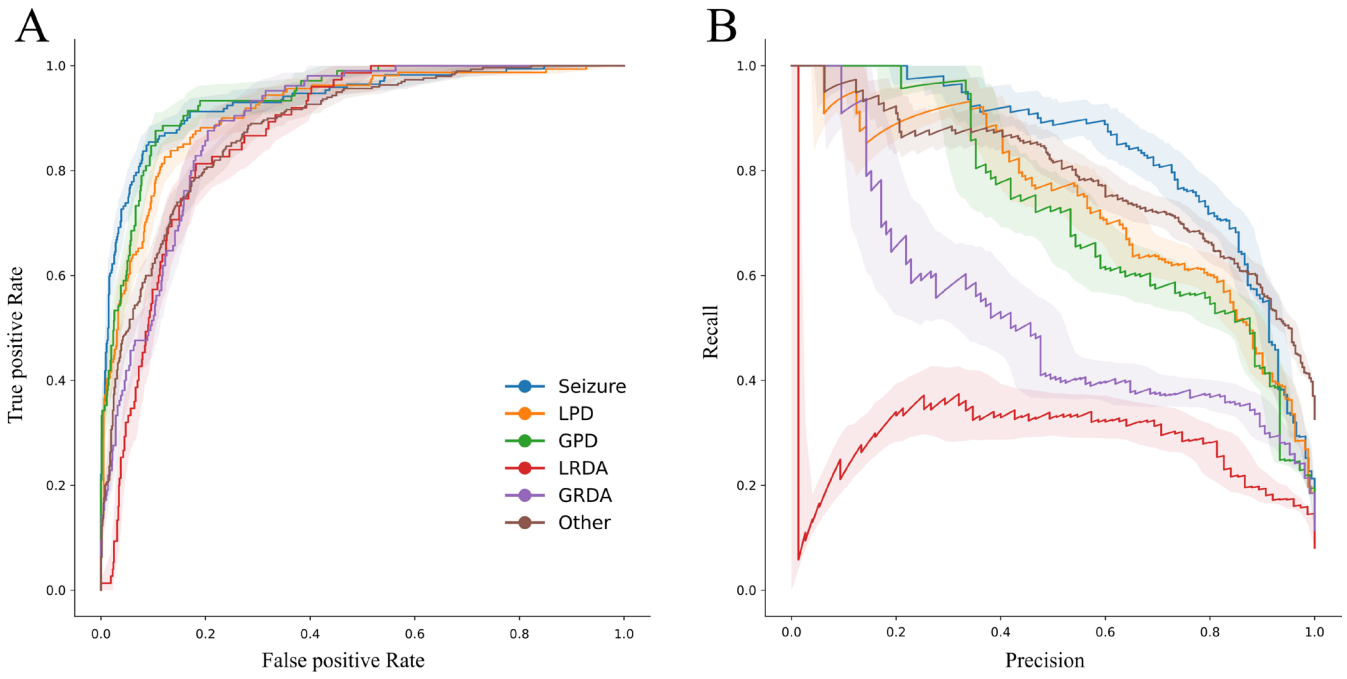

FIGURE 2 | Receiver-operating characteristic (ROC, panel A) and Precision-recall (PR, panel B) curves of the decoded classes. Color represents each classified class. Shaded areas are the confidence intervals of the ROC curves at 95%.

The hard labeling via SVC of the model probabilities produced instead a TPR of 75% for SZ, 68% for LPD, 64% for GPD, 55% for LRDA, 22% for GRDA, and 75% for others, and a FPR of 6% for SZ, 7% for LPD, 3% for GPD, 11% for LRDA, 1% for GRDA, and 14% for other.

3.2 | Embedding Exploration

We show that the embedding of the network successfully clustered the classes across its dimensions, as all embedding distances show an effect of annotation ($p < 0.001$ for all four ANOVAs, Figure 3A). The post hoc pairwise comparisons also show that the network has a good ability to discriminate pattern type (PD vs. RDA, $p < 0.001$ for all comparisons), but struggles to clearly represent lateralization of the patterns correctly. Indeed, while a small clustering can be found when assessing the pairwise distances between GPD and LPD ($p = 0.019$) from LRDA embeddings, none of the other patterns seem to show a significant encoding of the spatial information.

3.3 | Temporal Classification in the HUG Dataset

Due to poor EEG technical quality, one patient was excluded, and we analyzed the EEG recordings of 14 patients (11 males; age: 59 ± 18) and a total of 135 epileptic seizures (mean seizure per patient: 9.64 ± 10.26).

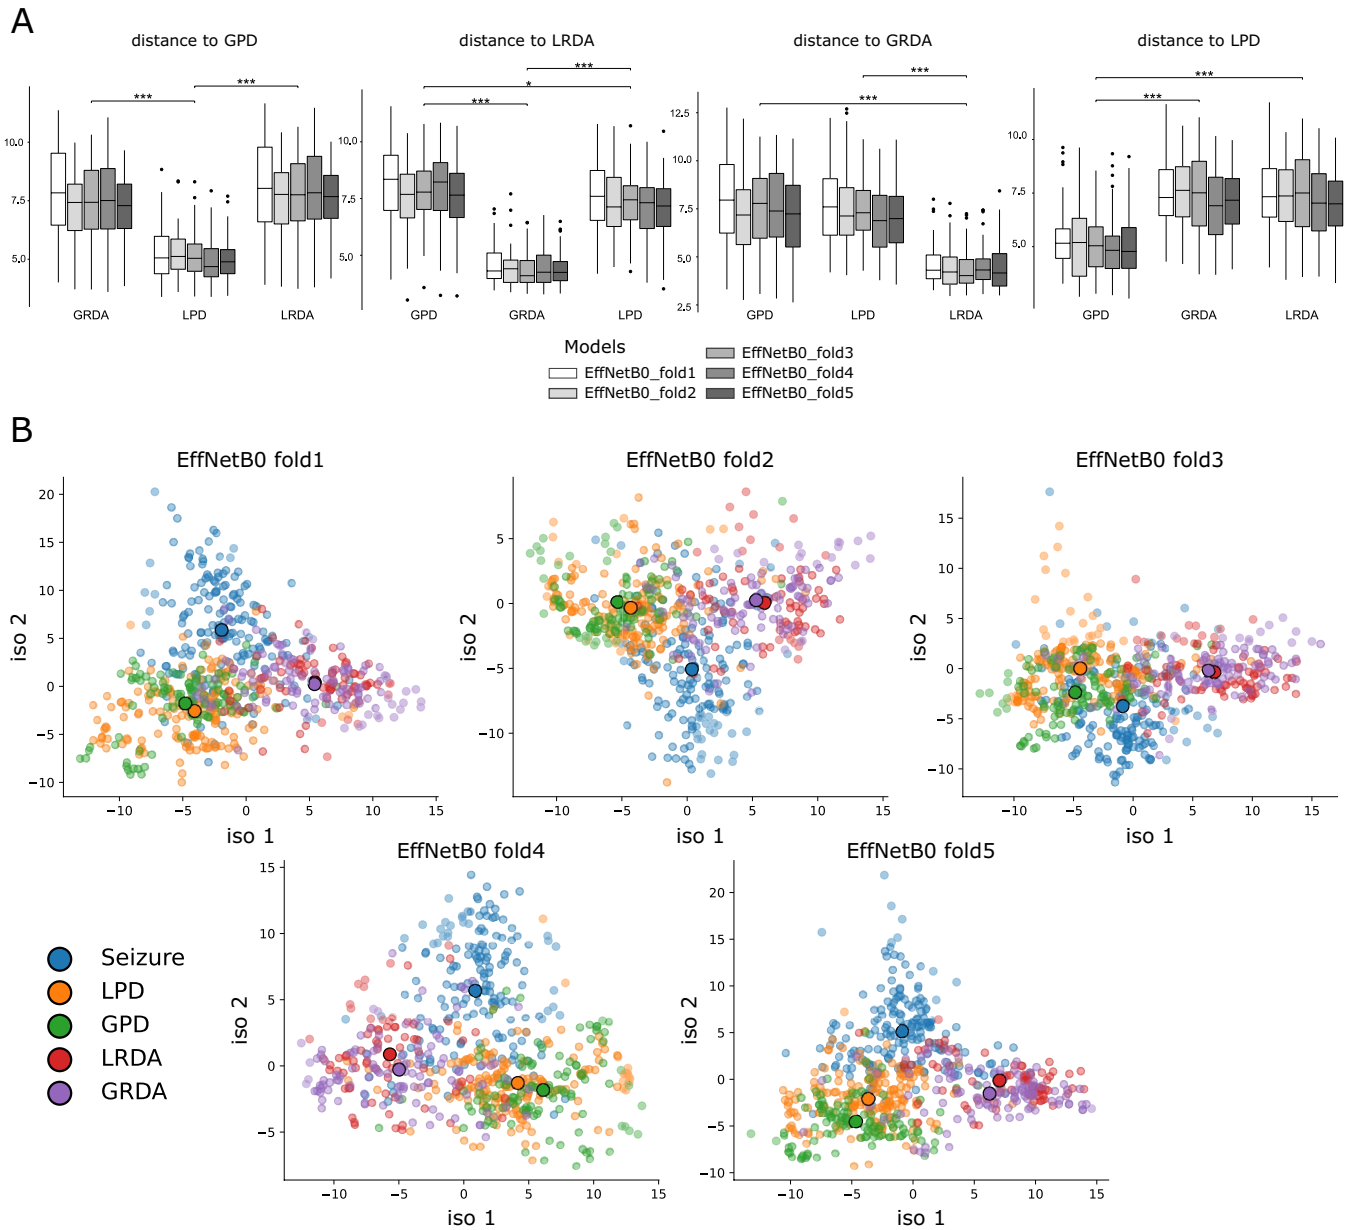

FIGURE 3 | (A) Distances across the logit embeddings between the three EEG clinical patterns. Distances were computed in a 1 vs. all fashion for each class in the set. Shading of grays corresponds to the distances of within each model embeddings. (***$p < 0.001$, *$p < 0.05$) (B) 2D Isomap representations of the embeddings from the five EfficientNet models. These embeddings were extracted from the penultimate layer of each network, after pooling. Overlapping areas correspond to EEG patterns where the network exhibits greater decision uncertainty across two or more classes. Colored circles correspond to the centroid of each class (Euclidean center of each distribution).

The SVC had an FPR of 4%, 1%, and 4% at −50, −30, and −20 s and a TPR of 42%, 60%, and 74% at 0, 30, and 50 s, respectively. The threshold of the calibrated detector was set at the probability of 14% of the EfficientNet seizures SoftMax output (Figure 4A). The fine-tuned seizure alarm detector had thus an FPR of 22%, 20%, and 21% at −50, −30, and −20 s and a TPR of 76%, 84%, and 89% at 20, 30, and 50 s, respectively (Figure 4B).

In Figure 5, the Grad-CAM analysis shows the focus of attention of the EfficientNet models on the time–frequency representations of four different EEG samples. A clear temporal correlation emerges between the model's focus and seizure onsets, regardless of the different certainty levels of the model. Specifically, all four spectrograms show that the area of attention is predominantly in the neighborhood of the annotation, accurately identifying its temporal position, with occasional overlap with the annotated onset.

4 | Discussion

While most prior research on automatic EEG annotation has concentrated mainly on identifying epileptic activity [12–14, 17], few studies have offered a comprehensive assessment of broader, clinically relevant EEG patterns [25] especially as defined by current guidelines [16, 17]. In the present study, we tested a novel lightweight CNN architecture to detect common EEG patterns in ICU patients. Our study demonstrated that a lightweight CNN could produce highly reliable results across standardized EEG patterns [4], with temporal performance enabling rapid detection of epileptic activity.

Our network showed strong results on a testing dataset reaching an AUROC score of 93% for SZ, 91% for LPD, 94% for GPD, 87% for LRDA, 89% for GRDA, and 88% for the "other" class, which importantly includes artifacts activity. Similarly, the AUPR score was above 70% for all the classes except RDA where it performed more poorly, especially for LRDA (30%, see below discussion on embeddings). While using only spectro-temporal information, our results are comparable with those of Jing and colleagues [17] and the SPaRCNet (seizures, periodic, and rhythmic pattern Continuum) [17] except for PR values of LRDAs (70% in the SPaRCNet vs. 30% in our model). The comparison with the other benchmark, SCORE-AI (Standardized Computer-based Organized Reporting of EEG–Artificial Intelligence) [16], is more challenging. While SCORE-AI is currently the most externally validated model [26], it is applied to standard EEGs, including those of healthy individuals. It explicitly excludes EEGs from critical care, making direct comparison with our ICU-based model difficult. Additionally, differently from SPaRCNet and our model, SCORE-AI uses terminology that does not fully align with that of the American Clinical Neurophysiology Society [4]. Our embedding analysis provides valuable insights into the model's ability to distinguish between RPPs and to identify lateralization. However, while the model could effectively differentiate between RDA and PD, it faced challenges in capturing the nuances of lateralization, particularly with RDA. Not surprisingly, this is indeed consistent with the low PR values obtained during decoding. This limitation mirrors our clinical experience, and it probably depends on the origin and representation of RDA and PD patterns: Delta rhythm mainly reflects subcortical dysfunction [27, 28] or epileptiform discharges in mesio-temporal regions [29]; and tends to have a more spread scalp distribution. As GRDA is a very frequent finding in ICU "encephalopathic" patients, especially frontally predominant [30]; LRDA is less frequent, involving mostly the temporal lobe [31]. LPD, on the contrary, has cortical, more superficial generators and is a frequent finding in ICU patients, with diverse precise localization frequently observed (occipital, frontal, temporal, and parietal). In

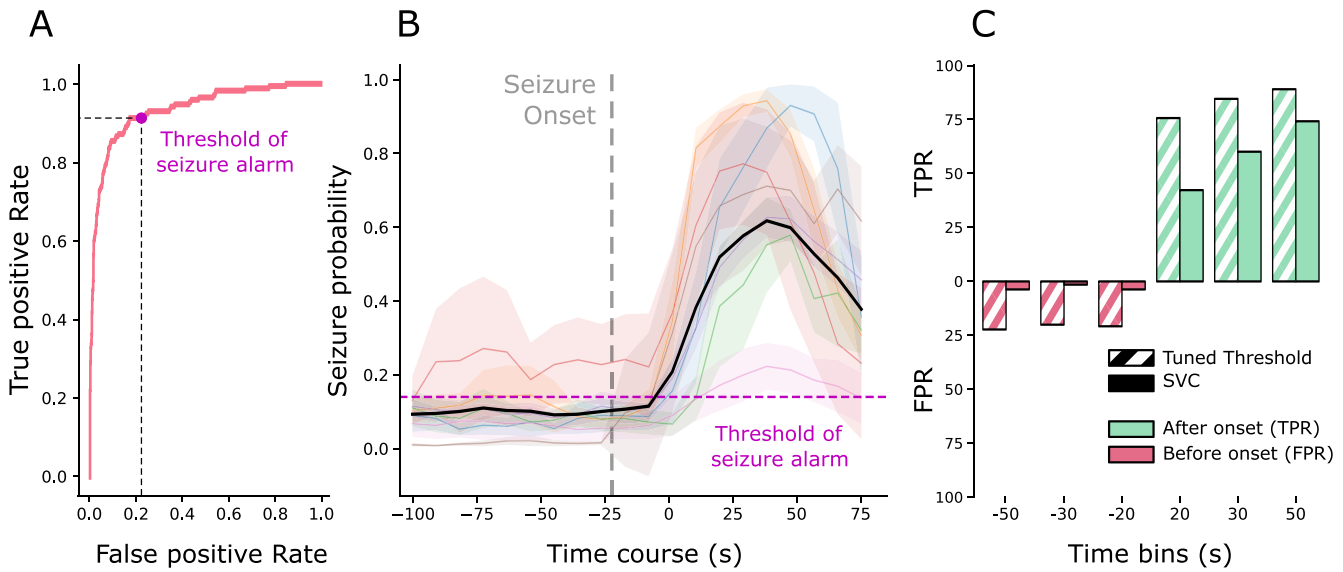

FIGURE 4 | (A) FPR and TPR of the seizure optimized threshold overlaid on the ROC curve of seizure class. (B) Probability of seizure across time before and after seizure onset annotation. Black line represents the total average across all the seizures. Each colored line corresponds to a different subject. Horizontal dashed line corresponds to the threshold of the tuned seizure alarm. For visualization purposes, we selected only the subject with at least six seizures during the recording session. Shaded areas are 2 SEM. Vertical gray dash line corresponds to the onset of the seizure. (C) Bar plot of the true positive rate (TPR) and false positive rate (FPR) before and after seizure onset. The dashed bar corresponds to the results of the tuned alarm detector. The filled bar is the SVC classifier performance.

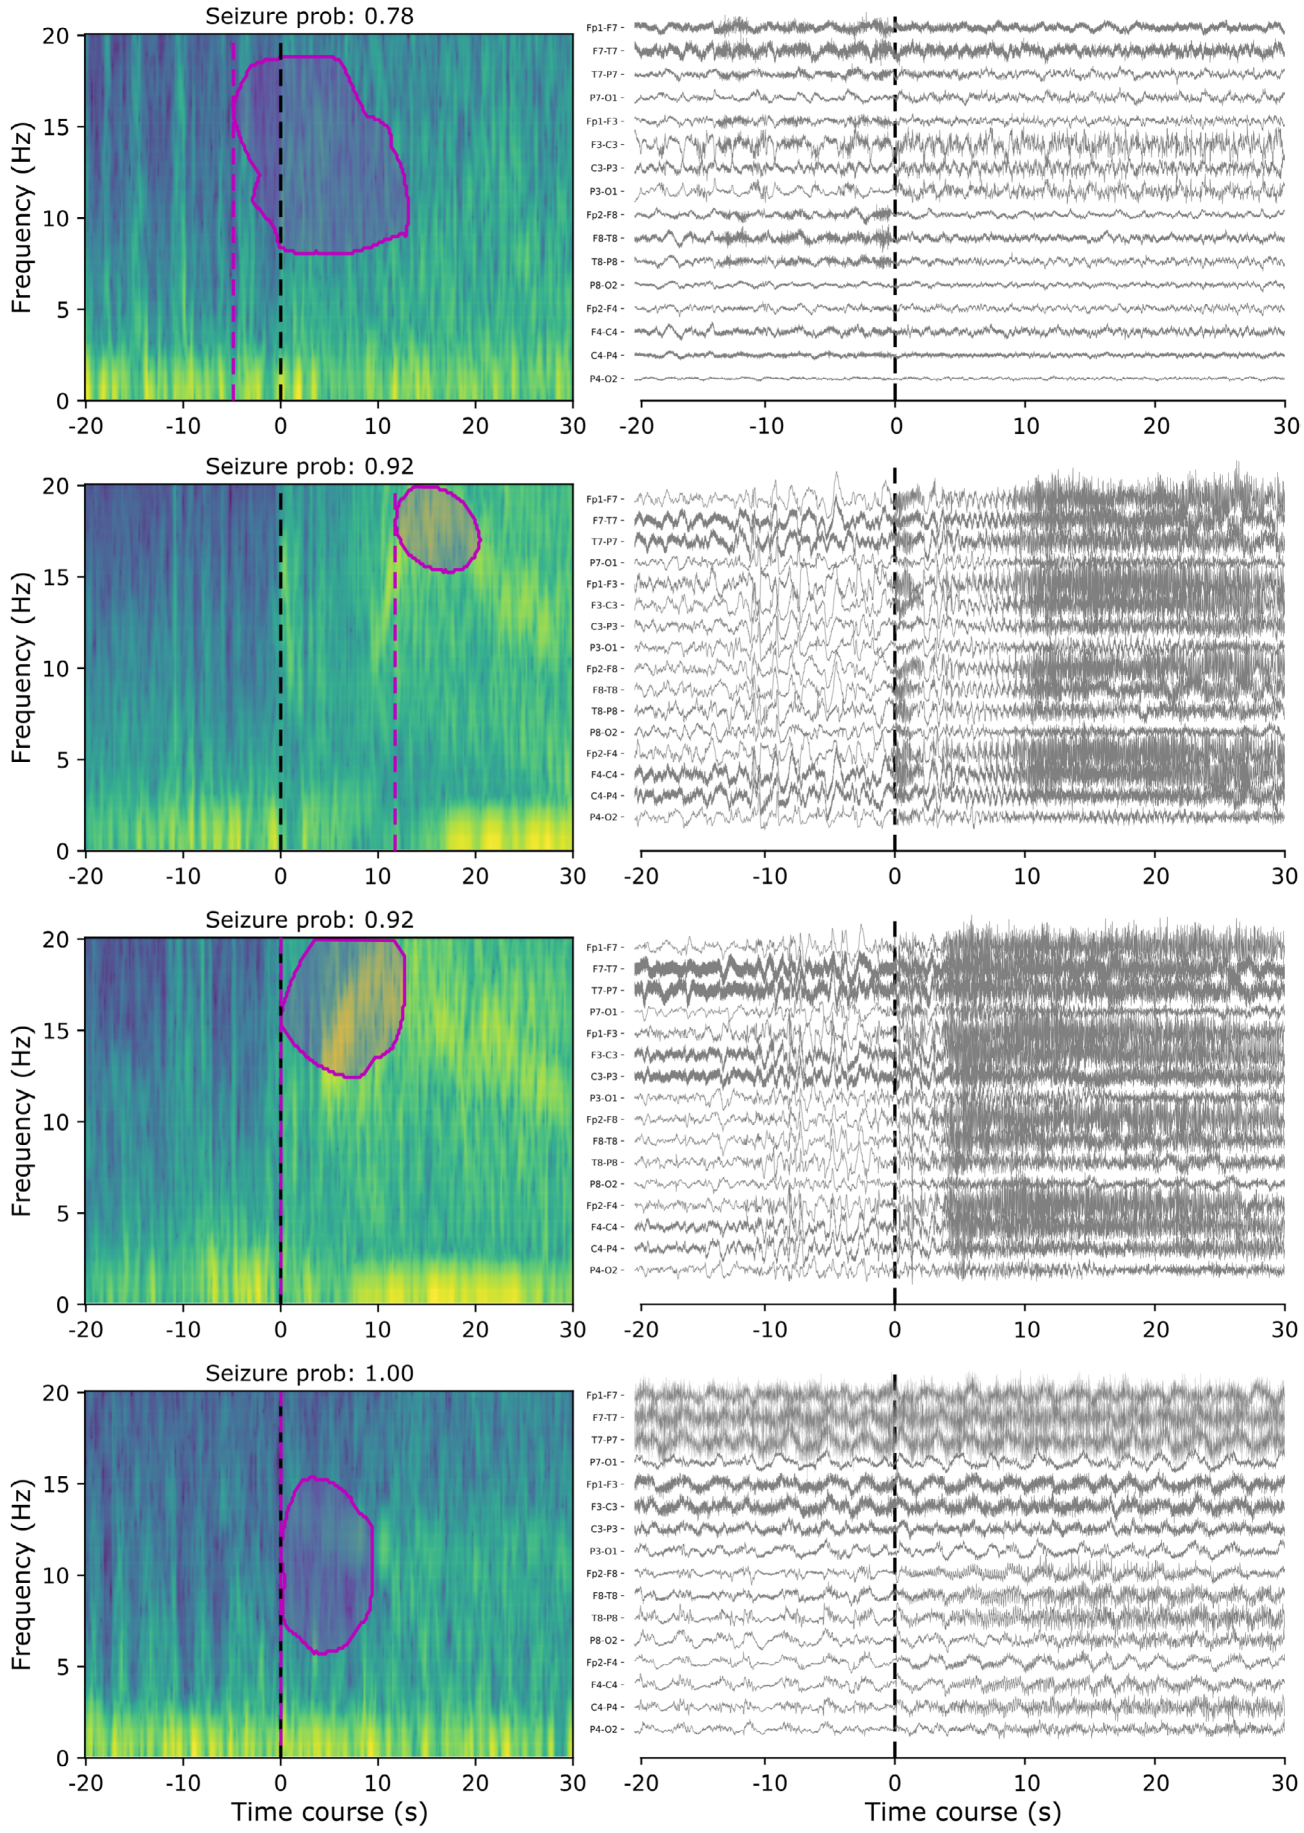

FIGURE 5 | Legend on next page.

FIGURE 5 | GRAD-CAM analysis of four different seizures across subjects. Left column shows the time–frequency plot of the EEG signal; right column shows the corresponding time courses of the double banana montage. EEG time courses that are displayed were first band passed filtered between 0.5 to 50 Hz. Purple circles represent the focus of the model for assigning the pattern to the Seizure class. Black dotted line is the annotated onset of the seizure. Purple dotted line is the estimated onset based on the temporal focus of the model.

clinical practice, a low precision for LRDA detection could lead to an excessive number of false alarms, increasing the risk of alarm fatigue and potentially undermining trust in the model's reliability during cEEG monitoring for RDA. As this limitation arises from an objective function that mirrors difficulties observed in expert clinician annotations [8], we underscore the need for further refinement in the encoding of spatial information within the network. Finally, following established calibration standards, the model achieved a good ECE score of 2.7% suggesting a balanced confidence in class predictions especially for clinical scenarios [22]. Importantly, this level of calibration can support the use of output probabilities in clinical decision-making, as well-calibrated confidence scores are essential for threshold-based alarms in EEG pattern detection.

In addition to the work from Jing and colleagues [17], we also evaluated the model's temporal performance in detecting seizures relative to their onset. Although limited by a small-to-medium sample size (135 seizures), the analysis demonstrated that the model achieved high sensitivity as early as 30 s after onset, especially if specifically tuned to detect seizures (84% of sensitivity and 22% of specificity). This result can be critical for ensuring prompt therapeutic interventions and is comparable with the human detection process. Moreover, given that this model was trained to classify six distinct patterns without leveraging temporal autocorrelations, we believe that the results achieved in seizure detection across time are particularly noteworthy. We want to emphasize that we avoided an explicit comparison with the existing literature on seizure detection and prediction (for a recent systematic review on the topic see [32]) as the cost function of the EfficientNet was not optimized for seizure detection only but was set to detect multiple annotations of RPPs.

We exploited the annotations of seizure onsets to explore whether the network activations can also serve as a tool for inspecting the spectro-temporal features of the patterns of interest (Figure 5). Using Grad-CAM analysis, we show that our model is often capable of segmenting seizure activity within a close temporal neighborhood to its onset. This is particularly significant given that no explicit onset information was used during training, yet the model demonstrates a qualitatively good performance on the test set. As the EEG pattern gets classified over time, clinicians can thus examine the specific regions within the time–frequency spectrum that influenced the network's conclusions. Previous research on neural network interpretability has demonstrated improvements in RPPs scoring simply by presenting the embeddings of the network without attention maps [18]. We believe this latter new feature can further validate clinical decisions and additionally highlight cortical biomarkers that might otherwise go unnoticed. This rapid detection, its tunability, and its interpretability thus highlight its potential for improving patient outcomes in time-sensitive situations [10]. Nevertheless, while beyond the scope of the current project, further quantitative analysis of the estimated onset should be addressed with the temporal performance score on a larger annotated dataset [33].

A major strength of our approach is its ability to detect EEG patterns with a good performance and minimal preprocessing, which is particularly valuable in critical care settings where timely interventions are essential. Integrating raw EEG data with spectral analysis allowed us to capture a more comprehensive and clinically meaningful representation of brain activity. Future model iterations may benefit from incorporating multimodal data to further enhance performance. For instance, while video integration may offer limited additional value in ICU settings with concomitant high resource costs, incorporating real-time physiological data, such as heart rate variability [34, 35], intracranial pressure, or partial pressure of oxygen in interstitial brain tissue [36], could refine EEG pattern interpretation and provide a sustainable and effective solution for continuous monitoring.

5 | Limitations

Despite these encouraging results, our study has several limitations. First, the model's generalizability remains untested on fully independent datasets, as testing was performed considering only seizures and a small cohort from a single institution. Moreover, the training of our model was affected by the limited amount of RDA activity, especially when lateralized, highlighting the need for a more balanced dataset. Multicenter validation studies are thus required to confirm its broader applicability. Second, the model's spatial resolution requires improvement to better detect lateralized EEG patterns. Third, future research should explore the model's ability to identify RPPs matching the ictal–interictal continuum [4], which indicates secondary neuronal damage and requires timely intervention.

6 | Conclusion

In conclusion, our lightweight CNN model provides an efficient and accurate tool for detecting seizures and other critical ICU-EEG pattern. Its minimal preprocessing requirements and robust performance make it well-suited for deployment in both resource-limited and advanced clinical settings, where it could reduce the workload of neurophysiologists by prescoring EEG data [18]. However, further research is necessary to validate its use in broader clinical contexts and to extend its capability to detect a wider range of clinically significant EEG patterns.

Author Contributions

Giulio Degano: conceptualization, data curation, formal analysis, investigation, methodology, writing - original draft. **Hervé Quintard:**

supervision; writing - review and editing. **Andreas Kleinschmidt:** supervision; writing - review and editing. **Nikita Francini and Oana E. Sarbu:** seizure annotation. **Pia De Stefano:** conceptualization, data curation, formal analysis; investigation, supervision, writing - original draft.

Acknowledgments

We gratefully acknowledge Jin Jing, Zhen Lin, Chaoqi Yang, Ashley Chow, Sohier Dane, Jimeng Sun, and M. Brandon Westover for their invaluable contribution to the open dataset, HBAC—Harmful Brain Activity Classification, which allowed the creation of our model. Open access publishing facilitated by Universite de Geneve, as part of the Wiley - Universite de Geneve agreement via the Consortium Of Swiss Academic Libraries.

Conflicts of Interest

G.D., H.Q., N.F., O.E.S., P.D.S. declares no competing interests. A.K. has received honoraria for consulting from Abbvie, Eli Lilly, Lundbeck, Mitsubishi Tanabe, Novartis, and TEVA that were paid to a teaching and research fund at the University Hospital Geneva.

Data Availability Statement

The HBAC datasets analyzed during the current study are available in the Kaggle repository (https://www.kaggle.com/competitions/hms-harmful-brain-activity-classification). The HUG data generated and analyzed during the current study are not publicly available due to ethical and legal considerations, including patient confidentiality and privacy protection. Access to the data and code can be made available from the corresponding author on reasonable request and review from the institutional ethics committee.

References

1. P. Vespa, D. Menon, P. Le Roux, and the Participants in the International Multi-Disciplinary Consensus Conference on Multimodality Monitoring, "The International Multi-Disciplinary Consensus Conference on Multimodality Monitoring: Future Directions and Emerging Technologies," *Neurocritical Care* 21, no. S2 (2014): 270–281.

2. L. R. Lara and H. A. Püttgen, "Multimodality Monitoring in the Neurocritical Care Unit," *Continuum* 24, no. 6 (2018): 1776–1788.

3. R. Sutter, P. Fuhr, L. Grize, S. Marsch, and S. Rüegg, "Continuous Video-EEG Monitoring Increases Detection Rate of Nonconvulsive Status Epilepticus in the ICU: CVEM NCSE Detection in ICU," *Epilepsia* 52, no. 3 (2011): 453–457.

4. L. J. Hirsch, M. W. K. Fong, M. Leitinger, et al., "American Clinical Neurophysiology Society's Standardized Critical Care EEG Terminology: 2021 Version," *Journal of Clinical Neurophysiology* 38, no. 1 (2021): 1–29.

5. G. B. Young and J. Mantia, "Continuous EEG Monitoring in the Intensive Care Unit," in *Handbook of Clinical Neurology* (Elsevier, 2017), 107–116, https://linkinghub.elsevier.com/retrieve/pii/B9780444636003000076.

6. J. Claassen, S. A. Mayer, R. G. Kowalski, R. G. Emerson, and L. J. Hirsch, "Detection of Electrographic Seizures With Continuous EEG Monitoring in Critically Ill Patients," *Neurology* 62, no. 10 (2004): 1743–1748.

7. A. O. Rossetti, K. Schindler, R. Sutter, et al., "Continuous vs Routine Electroencephalogram in Critically Ill Adults With Altered Consciousness and No Recent Seizure: A Multicenter Randomized Clinical Trial," *JAMA Neurology* 77, no. 10 (2020): 1225–1232.

8. J. Jing, W. Ge, A. F. Struck, et al., "Interrater Reliability of Expert Electroencephalographers Identifying Seizures and Rhythmic and Periodic Patterns in EEGs," *Neurology* 100, no. 17 (2023): e1737–e1749, https://doi.org/10.1212/WNL.0000000000201670.

9. C. E. Hill, L. J. Blank, D. Thibault, et al., "Continuous EEG Is Associated With Favorable Hospitalization Outcomes for Critically Ill Patients," *Neurology* 92, no. 1 (2019): e9–e18, https://doi.org/10.1212/WNL.0000000000006689.

10. F. Misirocchi, H. Quintard, A. Kleinschmidt, et al., "ICU-Electroencephalogram Unit Improves Outcome in Status Epilepticus Patients: A Retrospective Before-After Study," *Critical Care Medicine* 52, no. 11 (2024): e545–e556, https://doi.org/10.1097/CCM.0000000000006393.

11. U. R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan, and H. Adeli, "Deep Convolutional Neural Network for the Automated Detection and Diagnosis of Seizure Using EEG Signals," *Computers in Biology and Medicine* 100 (2018): 270–278.

12. C. Baumgartner and J. P. Koren, "Seizure Detection Using ScalpEEG," *Epilepsia* 59, no. S1 (2018): 14–22.

13. C. da Silva Lourenço, M. C. Tjepkema-Cloostermans, and M. J. A. M. van Putten, "Machine Learning for Detection of Interictal Epileptiform Discharges," *Clinical Neurophysiology* 132, no. 7 (2021): 1433–1443.

14. K. G. van Leeuwen, H. Sun, M. Tabaeizadeh, et al., "Detecting Abnormal Electroencephalograms Using Deep Convolutional Networks," *Clinical Neurophysiology* 130, no. 1 (2019): 77–84.

15. S. Saminu, G. Xu, Z. Shuai, et al., "A Recent Investigation on Detection and Classification of Epileptic Seizure Techniques Using EEG Signal," *Brain Sciences* 11, no. 5 (2021): 668.

16. J. Tveit, H. Aurlien, S. Plis, et al., "Automated Interpretation of Clinical Electroencephalograms Using Artificial Intelligence," *JAMA Neurology* 80, no. 8 (2023): 805–812.

17. J. Jing, W. Ge, S. Hong, et al., "Development of Expert-Level Classification of Seizures and Rhythmic and Periodic Patterns During EEG Interpretation," *Neurology* 100, no. 17 (2023): e1750–e1762, doi/10.1212/WNL.0000000000207127.

18. A. J. Barnett, Z. Guo, J. Jing, et al., "Improving Clinician Performance in Classifying EEG Patterns on the Ictal–Interictal Injury Continuum Using Interpretable Machine Learning," *NEJM Ai* 1, no. 6 (2024), https://doi.org/10.1056/AIoa2300331.

19. M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," (2019), arXiv, http://arxiv.org/abs/1905.11946.

20. K. Han, C. Liu, and D. Friedman, "Artificial Intelligence/Machine Learning for Epilepsy and Seizure Diagnosis," *Epilepsy & Behavior* 155 (2024): 109736.

21. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization" (2017), http://arxiv.org/abs/1412.6980.

22. C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On Calibration of Modern Neural Networks" (2017), http://arxiv.org/abs/1706.04599.

23. R. V. Lenth, "emmeans: Estimated Marginal Means, aka Least-Squares Means" (2017), 1.10.5, https://CRAN.R-project.org/package=emmeans.

24. R. R. Selvaraju, M. Cogswell, A. Das, et al., "Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization" (2016), https://arxiv.org/abs/1610.02391.

25. M. C. Cloostermans, C. C. de Vos, and M. J. A. M. van Putten, "A Novel Approach for Computer Assisted EEG Monitoring in the Adult ICU," *Clinical Neurophysiology* 122, no. 10 (2011): 2100–2109.

26. D. Mansilla, J. Tveit, H. Aurlien, et al., "Generalizability of Electroencephalographic Interpretation Using Artificial Intelligence: An External Validation Study," *Epilepsia* 65, no. 10 (2024): 3028–3037.

27. C. J. Stam and W. S. Pritchard, "Dynamics Underlying Rhythmic and Non-Rhythmic Variants of Abnormal, Waking Delta Activity," *International Journal of Psychophysiology* 34, no. 1 (1999): 5–20.

28. R. Sutter, R. D. Stevens, and P. W. Kaplan, "Clinical and Imaging Correlates of EEG Patterns in Hospitalized Patients With Encephalopathy," *Journal of Neurology* 260, no. 4 (2013): 1087–1098.

29. P. De Stefano, M. Carboni, R. Marquis, et al., "Increased Delta Power as a Scalp Marker of Epileptic Activity: A Simultaneous Scalp and Intracranial Electroencephalography Study," *European Journal of Neurology* 29, no. 1 (2022): 26–35.

30. E. A. Accolla, P. W. Kaplan, M. Maeder-Ingvar, S. Jukopila, and A. O. Rossetti, "Clinical Correlates of Frontal Intermittent Rhythmic Delta Activity (FIRDA)," *Clinical Neurophysiology* 122, no. 1 (2011): 27–31.

31. F. Brigo, "Intermittent Rhythmic Delta Activity Patterns," *Epilepsy & Behavior* 20, no. 2 (2011): 254–256.

32. P. Handa, Lavanya, N. Goel, and N. Garg, "Software Advancements in Automatic Epilepsy Diagnosis and Seizure Detection: 10-Year Review," *Artificial Intelligence Review* 57, no. 7 (2024): 181.

33. J. Dan, U. Pale, A. Amirshahi, et al., "SzCORE: Seizure Community Open-Source Research Evaluation Framework for the Validation of Electroencephalography Based Automated Seizure Detection Algorithms," *Epilepsia* (2024): 1–11.

34. B. Olmi, C. Manfredi, L. Frassineti, et al., "Heart Rate Variability Analysis for Seizure Detection in Neonatal Intensive Care Units," *Bioengineering* 9, no. 4 (2022): 165.

35. F. Mason, A. Scarabello, L. Taruffi, et al., "Heart Rate Variability as a Tool for Seizure Prediction: A Scoping Review," *Journal of Clinical Medicine* 13, no. 3 (2024): 747.

36. J. Witsch, H.-P. Frey, J. M. Schmidt, et al., "Electroencephalographic Periodic Discharges and Frequency-Dependent Brain Tissue Hypoxia in Acute Brain Injury," *JAMA Neurology* 74, no. 3 (2017): 301–309.