Deep Neural Approaches for Detecting Seizures via Learned Feature Representations

Tony Livins

Solent University

United Kingdom

Contents

Abstract

This study investigates the effectiveness of deep learning in automatically identifying epileptic seizures from electroencephalogram (EEG) signals. The task is inherently challenging due to the significant variability in seizure patterns, both across different patients and within the same individual over time. To address this, we propose a hybrid deep neural architecture that integrates convolutional and recurrent components to jointly capture spatial, temporal, and spectral characteristics of EEG data.

The model learns a robust and spatially invariant representation of seizure activity, enabling consistent detection despite variations in signal patterns. Experimental evaluation demonstrates that the proposed approach substantially outperforms existing cross-patient seizure detection methods, achieving improved sensitivity while reducing false positive rates. In addition, the model exhibits strong resilience to practical limitations, including missing EEG channels and differences in electrode configurations, highlighting its applicability in real-world clinical settings.

1 Introduction

Epilepsy is a chronic neurological condition affecting over 50 million individuals globally, representing a significant burden on healthcare systems and patient quality of life (Megiddo et al., 2016). The primary clinical tool for diagnosing and monitoring epilepsy is the electroencephalogram (EEG), which records electrical activity in the brain. In standard clinical workflows, EEG recordings are manually reviewed by trained neurologists to identify pathological patterns such as seizures and pre-ictal abnormalities. However, this process is both time-consuming and resource-intensive, often requiring several hours to analyse a single day of recordings for one patient. The reliance on specialised expertise further limits scalability, particularly in resource-constrained environments and developing regions.

These challenges have driven substantial interest in the development of automated seizure detection systems. Traditional approaches have largely relied on manually engineered features derived from EEG signals. Many methods focus on spectral characteristics (Tzallas et al., 2012), while others attempt to capture temporal dynamics associated with seizure activity (Shoeb, 2009). Although such approaches have achieved moderate success, they are fundamentally constrained by the variability of epileptic seizures. Seizure patterns are highly non-stationary and exhibit significant variation both across patients and within the same patient over time (CP, 2010). This variability limits the generalisability of models based on fixed, handcrafted representations.

To address these limitations, this study explores the use of deep learning within a supervised framework to automatically learn robust feature representations directly from EEG data. Deep learning models have demonstrated superior performance across a range of domains, including computer vision and speech recognition, due to their ability to capture complex, hierarchical patterns in data (LeCun and Bengio, 1995). The proposed architecture combines convolutional and recurrent neural network components to jointly model spatial, temporal, and spectral information present in EEG signals. This design enables the learning of a spatially invariant representation of seizure activity, which is particularly important for cross-patient generalisation.

The proposed method is evaluated using the publicly available CHB-MIT dataset, one of the most widely used benchmarks for seizure detection research (Shoeb, 2009). Experimental results demonstrate that the model achieves performance comparable to state-of-the-art approaches in patient-specific settings, while significantly outperforming existing methods in cross-patient detection tasks. In addition, the model shows robustness to practical challenges such as missing EEG channels and variations in electrode configurations, supporting its applicability in real-world clinical environments.

2 Problem Definition

Epileptic seizures arise from abnormal, excessive, and synchronised neuronal activity within the brain, often resulting in a range of neurological symptoms, including impaired consciousness, involuntary movements, and sensory disturbances. Clinically, seizures are broadly categorised into focal seizures, which originate in a specific brain region, and generalised seizures, which involve widespread neural activity across the brain. The presentation of seizures is highly heterogeneous, varying significantly both between patients and within the same individual over time.

Effective seizure management relies heavily on accurate detection and classification. In acute scenarios, rapid identification allows for timely intervention to terminate ongoing seizures. In long-term care, continuous monitoring supports treatment evaluation and optimisation. For patients with focal epilepsy, precise localisation of seizure onset is essential for surgical intervention, where resection of epileptogenic tissue may provide a curative outcome. These clinical requirements highlight the importance of reliable and scalable seizure detection methods.

Electroencephalography (EEG) remains the primary diagnostic modality for monitoring seizure activity. EEG captures electrical signals from the brain via electrodes placed on the scalp or, in some cases, directly on cortical surfaces. These recordings typically span extended periods, ranging from hours to several days. Manual analysis of such data by trained neurologists is both time-intensive and resource-demanding, motivating the development of automated detection systems to reduce workload and improve accessibility.

Two primary detection scenarios are commonly considered. In patient-specific detection, models are trained using annotated data from an individual patient, enabling highly tailored performance. In contrast, cross-patient detection aims to generalise across individuals, identifying seizures in patients for whom labelled data may not be available. While patient-specific approaches often achieve higher accuracy, cross-patient models are more applicable in real-world settings, particularly in resource-limited environments.

This study evaluates both scenarios using the CHB-MIT dataset, a widely used benchmark in seizure detection research (Shoeb, 2009; Tzallas et al., 2012). The dataset comprises 969 hours of EEG recordings from 23 patients, including 173 annotated seizure events. It includes diverse seizure types and patient demographics, providing a realistic evaluation environment. The detection task is formulated as a binary classification problem, where each 30-second segment of EEG data is labelled as either containing a seizure or not.

3 Previous Work

Research in automated seizure analysis has traditionally focused on two distinct problems: seizure prediction and seizure detection (Gotman, 1999). Seizure prediction aims to anticipate the onset of seizures before they occur, enabling proactive intervention strategies (Mormann et al., 2007). In contrast, seizure detection focuses on analysing completed EEG recordings to identify seizure events retrospectively, supporting diagnosis, monitoring, and treatment planning (SJ, 2005). The present study focuses on offline seizure detection due to its direct relevance to clinical workflows.

Model performance in seizure detection is typically evaluated using sensitivity and false detection rate (Tzallas et al., 2012). Sensitivity measures the proportion of true seizures correctly identified, while the false detection rate quantifies the number of incorrect detections per hour of recording. Achieving a balance between these metrics is critical, as high sensitivity often comes at the cost of increased false positives. While some studies report specificity instead, this measure can be misleading in this context, as even high specificity can correspond to an unacceptably high false alarm rate when applied to long-duration EEG recordings.

3.1 Patient-Specific Detectors

Patient-specific approaches have shown strong performance by leveraging machine learning techniques combined with handcrafted EEG features. These features typically capture spectral, spatial, and temporal characteristics associated with seizure activity. For example, Shoeb (2009) demonstrated high detection performance using a support vector machine trained on engineered features, achieving a sensitivity of 96% with a low false detection rate. Similar results have been reported in subsequent studies, reinforcing the effectiveness of tailored models when sufficient patient-specific data is available (Fotiadis, 2016).

3.2 Cross-Patient Detectors

In contrast, cross-patient detection remains a significantly more challenging problem due to the variability of seizure patterns across individuals. Differences in seizure morphology, duration, and spatial origin complicate the development of generalised models. Large-scale studies, such as that by Furbass et al. (2014), provide benchmarks for this task, reporting moderate sensitivity and false detection rates across multi-centre datasets. Commercial systems have demonstrated limited performance in this setting, highlighting the need for more robust and generalisable approaches (Wilson et al., 2004).

4 Methods

To address the challenges associated with cross-patient variability, we propose a deep learning architecture designed to capture the complex characteristics of seizure activity. The model combines convolutional and recurrent neural network components to jointly model spatial, temporal, and spectral features within EEG data.

The detection pipeline consists of two main stages. First, multi-channel EEG signals are transformed into an image-based representation, enabling the integration of spatial and frequency-domain information. This transformation incorporates domain knowledge to preserve meaningful signal structure. Second, a recurrent convolutional neural network is trained on these representations to classify each segment as seizure or non-seizure. By learning hierarchical feature representations directly from data, the model aims to achieve improved generalisation across patients while maintaining sensitivity to clinically relevant patterns.

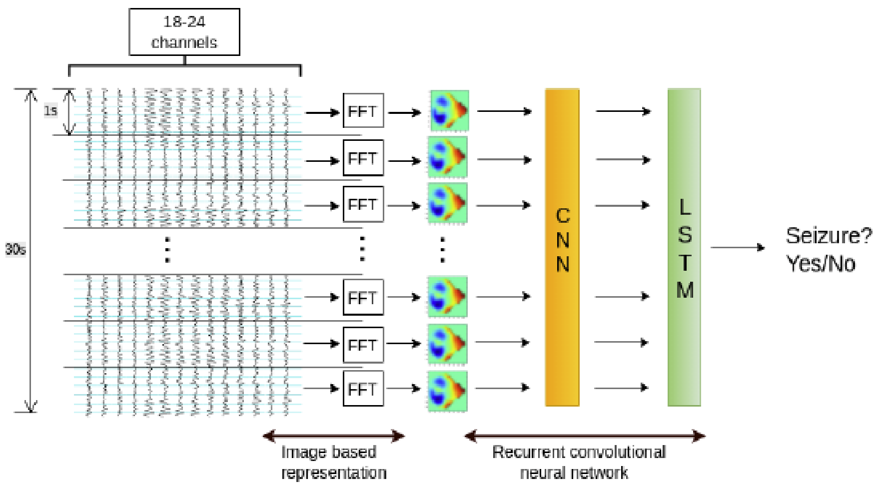

Figure 1 A Recurrent Convolutional Framework for EEG Analysis Using Image-Based Representations

4.1 Image-Based Representation of EEG Signals

To effectively capture the spatial relationships inherent in seizure activity, the proposed approach first transforms raw EEG signals into an image-based representation. This transformation incorporates domain-specific knowledge of electrode placement, enabling the model to exploit spatial locality within the data. The method follows a structured process inspired by prior work in EEG visualisation (Bashivan et al., 2016).

The initial step involves mapping the three-dimensional coordinates of EEG electrodes onto a two-dimensional plane. To preserve the relative spatial arrangement of electrodes, a polar projection technique is employed, ensuring that distances between electrodes in three-dimensional space are approximately maintained in the resulting two-dimensional representation (Snyder and Parr, 1987).

Once projected, each electrode location is assigned values across three channels, corresponding to distinct frequency bands extracted from the EEG signal. These bands capture low-frequency (0–7 Hz), mid-frequency (7–14 Hz), and high-frequency (14–49 Hz) components within a one-second segment of the signal. This multi-channel representation allows the model to encode spectral information alongside spatial structure.

To generate a continuous image, the discrete electrode values are interpolated across the two-dimensional grid using cubic interpolation. This step produces smooth spatial maps that approximate the underlying signal distribution. The final output is a set of images with dimensions 3×16×16, where the three channels correspond to frequency bands, and the spatial resolution reflects the interpolated electrode layout.

This representation enables the application of convolutional neural networks, allowing the model to learn spatial patterns associated with seizure activity in a manner analogous to image processing tasks.

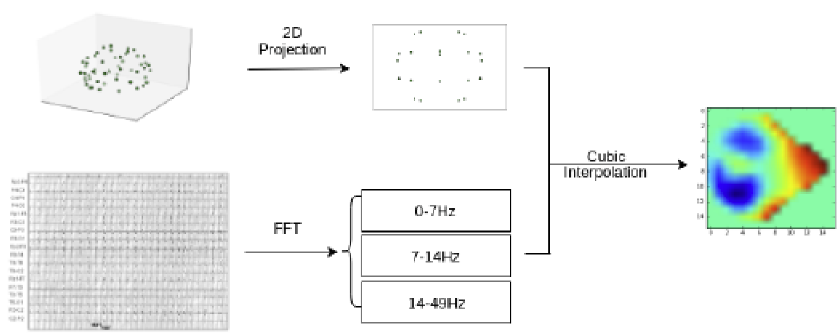

Figure 2 EEG Signal Representation in Image Form

4.2 Recurrent Convolutional Neural Network

Convolutional neural networks (CNNs) are a class of deep learning models inspired by the hierarchical processing mechanisms of the human visual system. They have demonstrated strong performance across a wide range of computer vision tasks due to their ability to learn spatially invariant feature representations (Krizhevsky *et al.*, 2012; LeCun and Bengio, 1995). This property is particularly advantageous for EEG analysis, as seizure activity may occur in different spatial regions of the brain. In this context, convolutional layers enable the model to learn generalisable patterns that are robust to spatial variations, effectively capturing common characteristics of seizures across different electrode locations.

A typical convolutional architecture consists of successive convolutional and pooling layers, followed by one or more fully connected layers. In this work, the CNN component is adapted from architectures that have achieved strong performance in large-scale image recognition tasks. Rather than relying on manually engineered features, the network is used as an automated feature extractor, learning hierarchical representations directly from the image-based EEG inputs described in Section 4.1.

To capture the temporal dynamics of EEG signals, the convolutional component is integrated with a recurrent neural network (RNN). RNNs are designed to process sequential data by maintaining an internal state that evolves over time, allowing the model to capture temporal dependencies. In this study, Long Short-Term Memory (LSTM) units are employed due to their ability to model long-range dependencies and mitigate issues such as vanishing gradients (Hochreiter and Schmidhuber, 1997).

The combination of convolutional and recurrent layers enables the model to jointly learn spatial and temporal representations of EEG data. This is particularly important for seizure detection, as seizures are not only characterised by spatial patterns across electrodes but also by temporal evolution over several seconds. By integrating these components, the proposed architecture provides a unified framework for modelling complex EEG dynamics, improving its ability to detect seizures across both patient-specific and cross-patient scenarios.

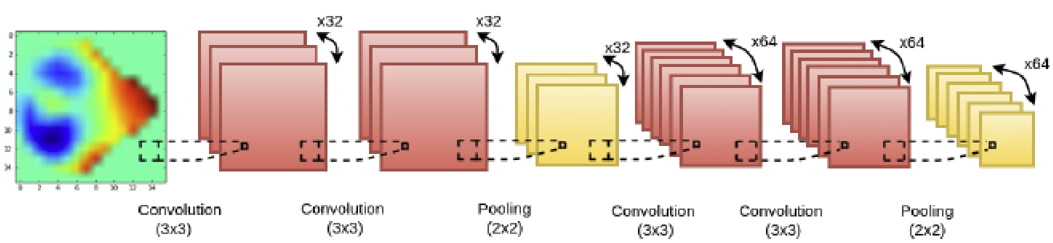

Figure 3 Structural Design of the Convolutional Component

Seizure activity typically unfolds over multiple consecutive one-second segments, making it essential to capture temporal dependencies across these intervals. To address this, the model incorporates bidirectional recurrent neural networks, which extend standard recurrent architectures by processing sequences in both forward and backward directions (Graves and Schmidhuber, 2005). This design allows the model to leverage information not only from preceding time steps but also from subsequent ones.

Such an approach closely reflects clinical practice, where neurologists often examine both prior and future signal segments when assessing a particular EEG window. In the proposed architecture, one set of LSTM units processes the sequence in chronological order, while a second set processes it in reverse order. By combining these two perspectives, the model is able to form a more comprehensive representation of temporal patterns associated with seizure activity, thereby improving detection performance.

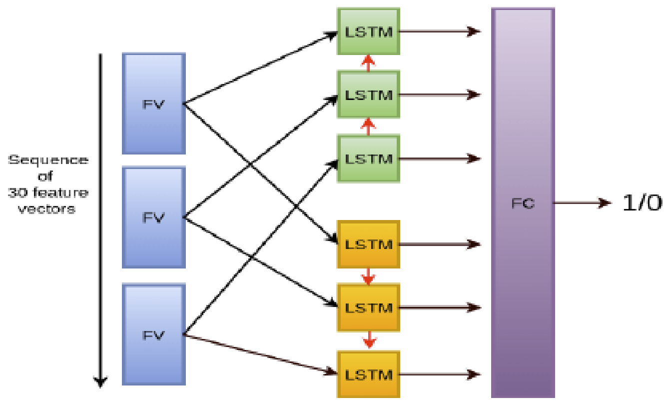

Figure 4 Recurrent neural network architecture. FV denotes a 64-dimensional feature vector, LSTM represents Long Short-Term Memory units with 128 hidden states, and FC indicates a fully connected layer with 512 hidden units.

Integration of Convolutional and Recurrent Components

To integrate spatial and temporal modelling, the output of the convolutional network is used as input to the recurrent component. Specifically, each EEG segment of one second is transformed into an image and processed by the convolutional network, producing a feature vector of dimension 64. These feature vectors are then sequentially fed into the recurrent architecture.

The recurrent network operates on sequences of 30 such feature vectors, corresponding to 30 seconds of EEG data. This design enables the model to capture temporal dependencies across multiple time windows. The combined recurrent–convolutional architecture is trained end-to-end using gradient-based optimisation, allowing both components to jointly learn complementary representations.

Model hyperparameters are selected through random search over a predefined parameter space, following established practices in deep learning (Goodfellow *et al.*, 2016). The final configuration includes a batch size of 128, the RMSProp optimiser, a learning rate of 0.001, and no dropout regularisation. To prevent overfitting, early stopping is employed, whereby training is halted once validation performance begins to degrade (Yao *et al.*, 2007).

Training Strategy

For patient-specific detection, the model is trained exclusively on data from a single patient. Due to the limited number of seizure events, a leave-one-out evaluation strategy is adopted. In this setup, the model is trained on all but one seizure instance and evaluated on the withheld instance. This process is repeated such that each seizure is used once for testing, ensuring robust performance estimation.

In the cross-patient setting, the model is trained on data from all but one patient and evaluated on the excluded individual. This process is repeated across all patients, providing a comprehensive assessment of generalisation capability.

4.3 Efficient Training Under Data Constraints

Training deep neural networks in the context of seizure detection presents several challenges, including class imbalance and limited availability of labelled data. To address class imbalance, the majority class (non-seizure data) is randomly subsampled during training, adjusting the class ratio to approximately 80:20. This rebalancing improves the learning process, as deep models often struggle with highly skewed datasets (Hensman and Masko, 2015). During evaluation, however, the full dataset is used to ensure realistic performance measurement. Prediction probabilities are appropriately rescaled to account for the altered training distribution.

Another challenge arises from the scarcity of seizure examples. To mitigate this, a pre-training strategy is employed. The convolutional component is first trained independently to classify individual one-second EEG segments. The learned weights are then used to initialise the full recurrent–convolutional model, which is subsequently trained on longer temporal sequences. This approach facilitates more stable optimisation and improves convergence (Erhan *et al.*, 2010).

For patient-specific models, where data is even more limited, transfer learning is applied. A general seizure representation is first learned using data from multiple patients. The model is then fine-tuned on data from a specific patient, leveraging previously learned weights to improve performance in low-data settings.

To further enhance model robustness and reduce prediction variance, an ensemble approach is adopted. Predictions from three independently initialised models with identical architectures are averaged, leading to more stable and reliable outputs (Zhou *et al.*, 2002).

5 Results on the CHB-MIT Dataset

The performance of the proposed model is evaluated on the CHB-MIT dataset. For patient-specific detection, results are compared against a benchmark model developed by Shoeb (2009), which represents a strong baseline using handcrafted features and classical machine learning techniques.

In the cross-patient setting, the proposed model is compared with the commercial seizure detection system REVEAL (Wilson *et al.*, 2004), using published results from prior studies. These comparisons provide a meaningful benchmark for assessing the effectiveness of the proposed approach in both personalised and generalised detection scenarios.

5.1 Patient-Specific Detection

Benchmarking in seizure detection remains challenging due to variations in experimental setups, evaluation protocols, and definitions of seizure events across studies. Despite these challenges, patient-specific detection provides a useful setting for comparing model performance under controlled conditions.

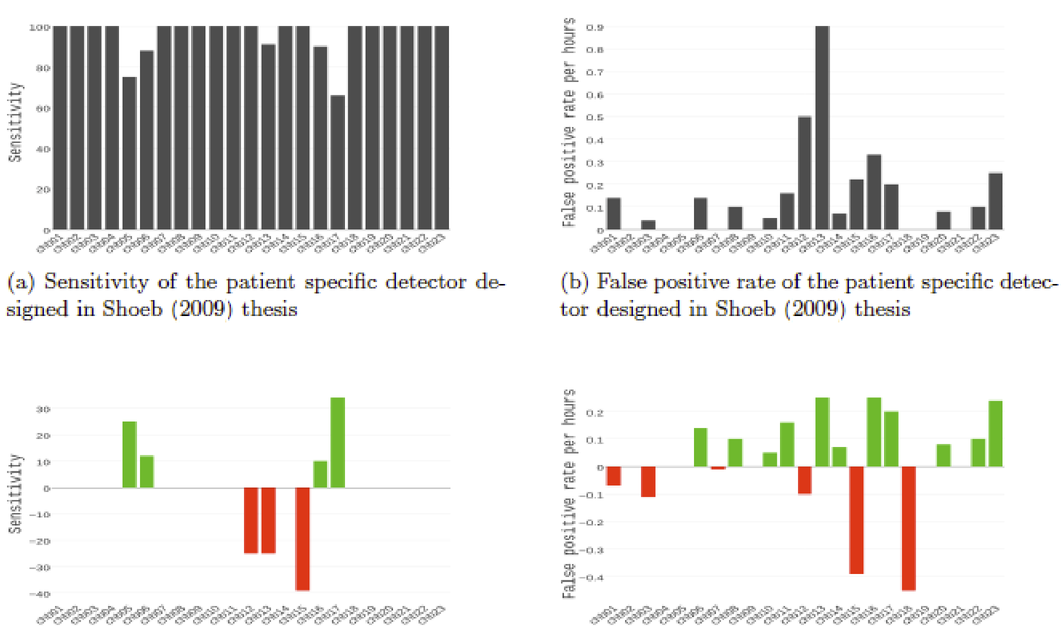

(a) Sensitivity of the patient specific detector designed in Shoeb (2009) thesis

(b) False positive rate of the patient specific detector designed in Shoeb (2009) thesis

(c) Sensitivity difference (recurrent convolutional neural network minus Shoeb detector)

(d) False positive rate difference(Shoeb detector minus recurrent convolutional neural network)

Figure 5 Performance comparison of the proposed patient-specific model with the approach introduced by Shoeb (2009), based on expert-engineered features and SVM classification.

Benchmarking seizure detection models remains inherently challenging due to variability in evaluation protocols and differences in how seizure events are defined across studies (Ronnera *et al.*, 2009). This challenge is particularly evident in patient-specific detection, where many approaches already achieve sensitivity levels within the 95–100% range, making fine-grained comparison difficult.

Figure 5 presents a comparison between the proposed neural architecture and the SVM-based detector introduced by Shoeb (2009). Both approaches demonstrate comparable performance in terms of sensitivity and false positive rate, indicating that they are similarly effective in patient-specific settings. The achieved performance levels suggest that both models are suitable for clinical application in personalised seizure detection scenarios (KM *et al.*, 2012). However, a key advantage of the proposed deep learning model lies in its robustness. As illustrated in Figure 6, the neural architecture maintains stable performance even in the presence of missing EEG channels, a condition that commonly arises in practical clinical environments.

5.2 Cross-Patient Detection

In contrast to patient-specific models, traditional approaches often struggle to generalise effectively across patients due to the high variability in seizure characteristics. This limitation is reflected in the performance of existing systems, where sensitivity can decline substantially when applied to unseen individuals. For example, the REVEAL system exhibits a notable reduction in detection accuracy compared to patient-specific models.

The proposed deep learning architecture addresses this challenge by learning generalisable representations of seizure activity. As shown in Figure 7, the model achieves an average sensitivity of 85%, significantly outperforming the REVEAL system, which achieves an average sensitivity of 67%. In addition to improved detection accuracy, the model also reduces the false positive rate from approximately 1.7 detections per hour to 0.8 detections per hour.

These results highlight the effectiveness of the proposed approach in cross-patient scenarios, demonstrating its ability to capture invariant features of seizure activity while maintaining robustness across diverse patient populations.

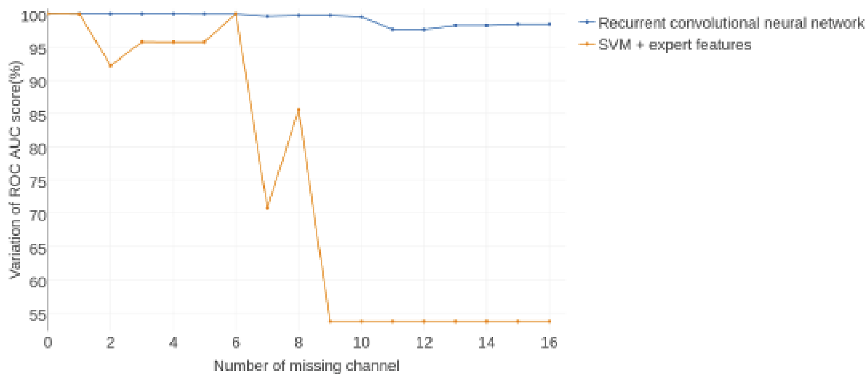

Figure 6 Comparison of robustness to missing EEG channels between the proposed recurrent convolutional neural network and an SVM model using handcrafted features. The SVM implementation follows the method described by Shoeb (2009).

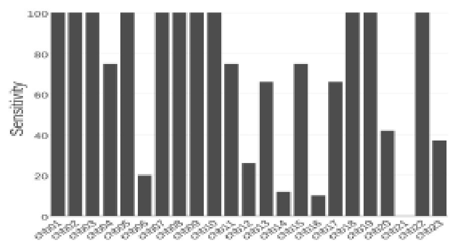

(a) Sensitivity of the cross patient algorithm REVEAL (Wilson et al., 2004)

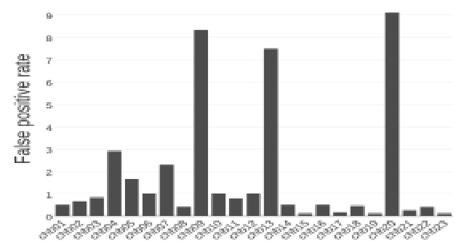

(b) False positive rate of the cross patient algorithm REVEAL (Wilson et al., 2004)

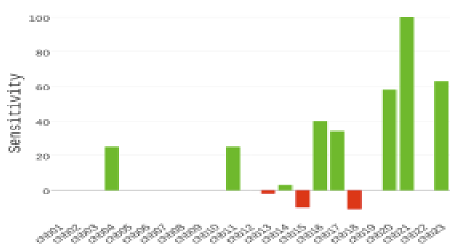

(c) Sensitivity difference (recurrent convolutional neural network minus REVEAL)

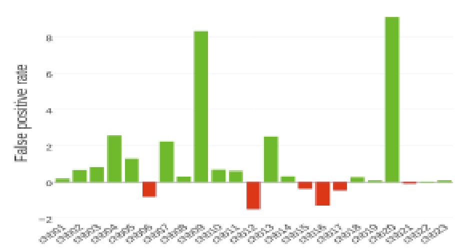

(d) False positive rate difference(REVEAL minus recurrent convolutional neural network)

Figure 7 Comparison between the proposed cross-patient model and the REVEAL algorithm (Wilson et al., 2004), evaluated on the CHB-MIT dataset as reported by Shoeb (2009).

6 Discussion

This study introduced a deep learning-based approach for automated seizure detection, designed to learn robust feature representations directly from EEG signals. By integrating spatial, temporal, and frequency-domain information, the proposed architecture is able to capture complex patterns associated with seizure activity. The results demonstrate that the model achieves performance comparable to state-of-the-art methods in patient-specific detection, while significantly improving performance in cross-patient scenarios. This highlights the model's ability to learn generalisable representations, addressing one of the key limitations of traditional approaches.

The practical implications of this work are substantial. Automating seizure detection has the potential to significantly reduce the burden on clinical workflows by minimising the need for manual EEG analysis. This can improve the efficiency of diagnosis, enable continuous patient monitoring, and support more informed treatment planning. The impact is particularly relevant in low-resource settings, where access to trained neurologists is limited, and automated tools could play a critical role in expanding healthcare accessibility.

An additional strength of the proposed approach lies in the use of an image-based representation of EEG data. By incorporating spatial relationships between electrodes and applying interpolation techniques, the model is able to operate effectively across different electrode configurations. This flexibility enhances its applicability in real-world clinical environments, where electrode montages may vary between patients or recording setups.

Furthermore, the architecture provides a degree of interpretability, which is often lacking in deep learning models. By systematically occluding regions of the input representation and evaluating the resulting impact on classification performance, it is possible to identify spatial areas that are critical for seizure detection. This occlusion-based analysis allows for approximate localisation of seizure activity within the brain. As illustrated in Figure 8, regions corresponding to the left parietal, frontal, and temporal lobes appear to play a significant role in detection. Such insights are clinically valuable, as localisation of seizure origin is an essential component of diagnosis and treatment planning, particularly in cases where surgical intervention is considered.

Overall, the proposed model not only improves detection performance but also offers enhanced robustness, adaptability, and interpretability, making it a promising candidate for deployment in practical clinical settings.

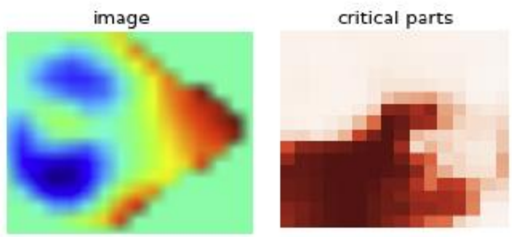

Figure 8 Occlusion-based visualisation highlighting the spatial regions of the EEG representation that are most influential in seizure detection.

Although the proposed model demonstrates strong performance in cross-patient evaluation, its effectiveness is not uniform across all individuals. In particular, sensitivity remains relatively low for a subset of patients. This discrepancy is likely attributable to variations in seizure characteristics that are not sufficiently represented in the training data, highlighting the importance of larger and more diverse datasets for improving generalisation.

In addition, the false positive rate observed in cross-patient detection remains higher than that achieved in patient-specific settings. This reflects the inherent difficulty of generalising across heterogeneous seizure patterns, where distinguishing between pathological and non-pathological activity becomes more challenging.

The training strategies outlined in Section 4.3 contribute significantly to improving model performance under limited data conditions. Techniques such as class rebalancing, pre-training, and transfer learning enable more efficient use of available data. However, the variability in prediction outputs remains a concern. Specifically, training on a reduced subset of negative samples while evaluating on a much larger and more diverse set introduces sensitivity to small changes in model parameters. This can lead to substantial fluctuations in the false positive rate.

This issue reflects a broader limitation of deep learning approaches when applied to small or imbalanced datasets. Future work could address this challenge by incorporating unsupervised or self-supervised pre-training strategies, allowing the model to learn more stable and generalisable representations prior to supervised fine-tuning.

References

Bashivan, P., Rish, I., Yeasin, M. and Codella, N., 2016. Learning representations from EEG with deep recurrent convolutional neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Chollet, F., 2015. *Keras*. Available at: https://github.com/fchollet/keras [Accessed 23 April 2026].

Erhan, D., Bengio, Y., Courville, A., Manzagol, P.A., Vincent, P. and Bengio, S., 2010. Why does unsupervised pre-training help deep learning?

Fotiadis, D.I., 2016. *Handbook of Research on Trends in the Diagnosis and Treatment of Chronic Conditions*. IGI Global.

Furbass, F., Ossenblok, P., Hartmann, M., Perko, H., Skupch, A.M., Lindinger, G., Elezi, L., Patarai, E., Colon, A.J., Baumgartner, C. and Kluge, T., 2014. Prospective multi-centre study of an automatic online seizure detection system for epilepsy monitoring units. *Clinical Neurophysiology*.

Goodfellow, I., Bengio, Y. and Courville, A., 2016. *Deep Learning*. Cambridge, MA: MIT Press. Available at: http://www.deeplearningbook.org [Accessed 23 April 2026].

Gotman, J., 1999. Automatic detection of seizures and spikes. *Journal of Clinical Neurophysiology*, 16, pp.130–140.

Graves, A. and Schmidhuber, J., 2005. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks*, 18(5), pp.602–610.

Hensman, P. and Masko, D., 2015. The impact of imbalanced training data for convolutional neural networks.

Hochreiter, S. and Schmidhuber, J., 1997. Long short-term memory. *Neural Computation*, 9(8), pp.1735–1780.

Kelly, K.M., Shiau, D.S., Kern, R.T., Chien, J.H., Yang, M.C.K., Yandora, K.A., Valeriano, J.P., Halford, J.J. and Sackellares, J.C., 2012. Assessment of a scalp EEG-based automated seizure detection system. *Clinical Neurophysiology*, 121.

Krizhevsky, A., Sutskever, I. and Hinton, G.E., 2012. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp.1097–1105.

LeCun, Y. and Bengio, Y., 1995. Convolutional networks for images, speech, and time series. In: M.A. Arbib, ed. *The Handbook of Brain Theory and Neural Networks*. MIT Press.

Megiddo, I., Colson, A., Chisholm, D., Dua, T., Nandi, A. and Laxminarayan, R., 2016. Health and economic benefits of public financing of epilepsy treatment in India: An agent-based simulation model. *Epilepsia*.

Mormann, F., Andrzejak, R.G., Elger, C.E. and Lehnertz, K., 2007. Seizure prediction: the long and winding road. *Brain*, 130, pp.314–333.

Panayiotopoulos, C.P., 2010. *A Clinical Guide to Epileptic Syndromes and Their Treatment*. Springer.

Shoeb, A., 2009. *Application of Machine Learning to Epileptic Seizure Onset Detection and Treatment*. PhD thesis. Massachusetts Institute of Technology.

Smith, S.J., 2005. EEG in the diagnosis, classification, and management of patients with epilepsy. *Journal of Neurology, Neurosurgery and Psychiatry*.

Snyder, J.P. and Parr, J., 1987. *Map Projections: A Working Manual*.

Tzallas, A.T., Tsipouras, M.G., Tsalikakis, D.G., Karvounis, E.C., Astrakas, L., Konitsiotis, S. and Tzaphlidou, M., 2012. Automated epileptic seizure detection methods: A review study. *Epilepsy - Histological, Electroencephalographic and Psychological Aspects*.

Wilson, S.B., Scheuer, M.L., Emerson, R.G. and Gabor, A.J., 2004. Seizure detection: evaluation of the REVEAL algorithm. *Clinical Neurophysiology*, 115, pp.2280–2291.

Yao, Y., Rosasco, L. and Caponnetto, A., 2007. On early stopping in gradient descent learning. *Constructive Approximation*, 26.

Zhou, Z.H., Wu, J. and Tang, W., 2002. Ensembling neural networks: many could be better than all. *Artificial Intelligence*, 137(1), pp.239–263.