



Continuous EEG Monitoring of Antiepileptic Therapy Using Unsupervised Deep Learning: A Proof-Of-Concept Study

Branislav S. Čeperković, Slobodan M. Janković

Fakultet medicinskih nauka, Univerzitet u Kragujevcu, Kragujevac, Srbija

SUMMARY

Introduction: Epilepsy affects more than 50 million people worldwide. In approximately one-third of patients, seizures cannot be adequately controlled with medication. A major challenge in clinical practice is that the effectiveness of a drug can only be assessed after weeks or months of observation - during which the patient may be receiving suboptimal treatment. The electroencephalogram (EEG) provides a direct, continuous window into brain electrical activity, yet in routine practice it is used almost exclusively for diagnosis rather than for ongoing therapy monitoring.

Aim: To develop and test a computational model that continuously quantifies how far the EEG signal is from a normal brain state - a measure we call latent distance - and to evaluate whether this measure could serve as a real-time, non-invasive biomarker of antiepileptic drug effect.

Material and Method: A deep learning model was trained on a publicly available EEG dataset (Bonn, Germany) using a self-supervised approach that requires no manual labelling of the data. The model was then tested on an entirely independent dataset (CHB-MIT, Boston, USA) containing EEG recordings from paediatric patients with drug-resistant epilepsy - a strict test of whether the model generalises beyond its training data. Performance was assessed using standard diagnostic accuracy metrics.

Results: On the external validation dataset, the model achieved an area under the ROC curve (AUC) of 0.80 (95% CI: 0.73-0.95), a precision-recall AUC of 0.97, and a balanced accuracy of 81%. Crucially, the latent distance signal followed a clinically meaningful trajectory: it rose before the documented seizure onset, peaked during the seizure, and gradually returned to baseline in the post-ictal period.

Conclusion: This proof-of-concept study demonstrates that a self-supervised EEG model can detect seizures and continuously track deviations from normal brain activity in patients it has never encountered before. Latent distance may represent a novel, continuous EEG biomarker with potential utility in real-time monitoring of antiepileptic therapy - particularly in identifying drug effects before changes in seizure frequency become apparent. Prospective clinical studies are needed to confirm this hypothesis.

Keywords: Epilepsy, EEG, Antiepileptic Drugs, Biomarker, Deep Learning, Seizure Detection

Corresponding author:

Full Professor, Primarius Slobodan M. Janković, MD, PhD
Specialist in Clinical Pharmacology

Department of Pharmacology and Toxicology, Faculty of Medical Sciences, University of Kragujevac, Kragujevac, Serbia

E-mail: sjankovic@fmn.kg.ac.rs

INTRODUCTION

Epilepsy and the challenge of treatment monitoring.

Epilepsy is one of the most common neurological disorders, affecting over 50 million people worldwide. It is characterised by recurrent, unprovoked seizures caused by abnormal, synchronised electrical discharges in brain networks. Although more than twenty antiepileptic drugs are currently available, approximately one-third of patients continues to experience seizures despite appropriate pharmacotherapy - a condition known as drug-resistant epilepsy [1,2].

A fundamental limitation of current clinical practice lies in the way treatment effectiveness is evaluated.

The standard approach involves observing the patient over a period of one to three months while maintaining an unchanged medication regimen and recording the number of seizures that occur. This method has an inherent delay: the clinician cannot determine whether a drug is effective until sufficient time has passed to observe whether seizures occur. During this waiting period, the patient may receive doses that are either too low to be effective or high enough to cause adverse effects, without any objective indication that treatment adjustment is needed. In addition, seizure reporting depends on patient recall and subjective experience, which further limits the reliability of this approach.

The underused potential of EEG.

Electroencephalography (EEG) records the brain's electrical activity in real time and has been the cornerstone of epilepsy diagnosis for decades. In routine clinical practice, however, EEG is used almost exclusively to confirm the diagnosis or to detect epileptiform discharges at a single point in time. Its potential as a continuous, quantitative monitor of the brain's physiological state - one that could track gradual medication-induced changes - has remained largely unexplored, primarily because the analytical tools required for such an application have not previously been available [3,4,5].

Artificial intelligence in EEG analysis.

In recent years, deep learning - a branch of artificial intelligence capable of learning complex patterns from large datasets - has demonstrated impressive performance in the automated analysis of EEG signals, particularly for seizure detection [19,20,21,22]. Most existing systems, however, are designed to answer a binary question: is a seizure occurring right now, or not? While useful, this approach provides only a snapshot in time and offers no information about the gradual, continuous changes in brain activity that occur between seizures and may reflect a pharmacological effect.

A new approach: continuous quantification of brain state.

This study proposes a different framework. Rather than simply classifying an EEG segment as „seizure” or „no seizure”, the aim was to train a model capable of continuously measuring how far the current brain state is from a normal physiological state. This measure - termed *latent distance* - is a continuous numerical score that rises and falls with the dynamic changes of the brain's electrical activity over time. Unlike binary classification, it has the potential to reflect pharmacodynamic drug effects even before any change in seizure frequency is detectable.

Self-supervised learning: a label-free approach.

To train this model, a method called self-supervised contrastive learning approach was employed. Unlike conventional supervised learning, which requires large amounts of expert-labelled data, self-supervised learning enables the model to extract meaningful EEG features from unlabelled recordings by learning to recognise different versions of the same signal as similar while distinguishing signals originating from different brain states [23,24,25].

AIM

The primary aims were: (1) to develop and train a model that learns a continuous representation of EEG-based brain state, using only publicly available, anonymised data; (2) to evaluate the model's ability to detect seizures

on an entirely independent dataset from a different institution, different patient population, and different recording equipment - the most stringent test of generalisation; and (3) to assess whether the latent distance measure follows a clinically meaningful temporal trajectory around documented seizures, supporting its potential utility as a biomarker of antiepileptic drug effect.

We hypothesise that the latent distance of the EEG signal from a normal brain state can serve as a novel, continuous, and non-invasive biomarker for monitoring antiepileptic therapy, and that, in combination with plasma drug concentration measurements, it could enable earlier dose adjustment without waiting for a change in seizure frequency.

MATERIAL AND METHODS

Study design

This is an academic, non-commercial phase IV study. No commercial sponsor was involved, and the study was not part of any funded project. Ethics Committee approval was not required, as the study was conducted exclusively on publicly available, fully anonymised datasets (Bonn EEG Dataset and CHB-MIT Scalp EEG Database) with no direct involvement of human subjects.

Training dataset (Bonn EEG Dataset)

The model was trained on a publicly available, anonymised EEG dataset collected at the University of Bonn, Germany - one of the most widely used reference datasets in epilepsy research [28]. It contains 500 EEG segments organised into five groups of 100 recordings each. Each segment is approximately 23.6 seconds long, recorded as a single-channel signal under standardised conditions.

The recordings originate from three clinically distinct sources: healthy volunteers (eyes open or closed), patients with epilepsy

during the seizure-free interval (interictal phase), and patients during an active epileptic seizure (ictal phase). All ictal recordings were verified by expert visual inspection. For this study, recordings from healthy volunteers and interictal patients were grouped as „normal“, and ictal recordings as „seizure“, yielding 400 normal and 100 seizure segments.

It should be noted that the dataset does not contain information on whether patients were receiving antiepileptic medication at the time of recording, or on drug doses and plasma concentrations. This is an important limitation for the interpretation of results and is discussed below.

External validation dataset (CHB-MIT Scalp EEG Dataset)

To test whether the model could generalise beyond its training data, it was evaluated on the CHB-MIT Scalp EEG Dataset, collected at Children's Hospital Boston in the United States and freely available through the PhysioNet platform [26,27]. This dataset contains continuous, long-term EEG recordings from paediatric patients with drug-resistant epilepsy. Seizure onset and offset times were annotated by clinical neurophysiologists as part of the original hospital monitoring.

A single recording was selected for analysis, containing a documented seizure occurring between 2,996 and 3,036 seconds (duration 40 seconds). Analysis focused on a 600-second window surrounding this event, using one standard EEG channel (FP1-F7). The model had no access to any CHB-MIT data during training, ensuring a strict zero-shot evaluation.

The key differences between the two datasets - which make the external validation a meaningful test of the model's ability to generalise - are summarised in Table 1.

Before being entered into the model, all EEG signals underwent a uniform preprocessing procedure. Each segment was stan-

Table 1. Comparison of the two EEG datasets used in this study

Feature	Bonn Dataset	CHB-MIT Dataset
Origin	University of Bonn, Germany	Children's Hospital Boston, USA
Recording type	Isolated segments	Continuous multi-hour recording
Sampling frequency	173.6 Hz	256 Hz
Number of channels	1 (single-channel)	16 (multi-channel)
Patient population	Adults	Paediatric patients
Use in this study	Model training	External validation

dardised by subtracting its mean value and dividing by its standard deviation — a step that removes differences in signal amplitude between patients and between recording devices, ensuring that the model learns from the shape of the signal rather than its absolute magnitude. All signals were adjusted to a uniform length of 4,096 data points. For the CHB-MIT recordings, which were sampled at a different frequency than the training data, the signal was resampled to match the input format expected by the model.

Conceptual overview

The model consists of two components: an encoder and a projection head.

The encoder is the core of the system. It receives a 23.6-second EEG segment and compresses it into a compact numerical representation consisting of 64 values that capture the essential characteristics of the brain's electrical activity at that moment. This representation can be viewed as the model's internal representation, or „fingerprint”, of the brain state encoded in that EEG segment. Technically, the encoder uses a series of filtering and compression steps to progressively extract meaningful patterns from the raw signal.

The projection head is a smaller auxiliary component used exclusively during training. It functions as a temporary training scaffold that encourages the encoder to learn clearer, more robust representations. Once training is complete, the projection head is discarded; all subsequent analyses rely exclusively on the 64-number output of the encoder.

Training strategy: self-supervised contrastive learning

The model was trained using a method called contrastive learning, in which the model learns to recognise that two slightly altered versions of the same EEG segment should have similar internal representations, while segments from different recordings should differ [23,24,25]. To generate these modified versions, three physiologically motivated transformations were applied: addition of random background noise (simulating electrode impedance variation), random scaling of signal amplitude (simulating differences in recording gain), and random time-shifting of the signal (simulating brief fluctuations in patient alertness). These

transformations were chosen specifically because they mimic the natural variability seen in clinical EEG recordings without distorting the underlying physiological patterns.

The model was trained for 15 cycles (epochs) using a standard optimisation algorithm. No seizure labels were used during this phase - the model learned entirely from the structure of the signals themselves.

Seizure classifier

After the encoder was trained, its 64-number output was used to train a simple statistical classifier (logistic regression) to distinguish between normal EEG segments and seizure recordings, using the labelled Bonn dataset. This classifier was trained with a correction for the fact that seizure recordings were far fewer in number than normal recordings. Both the encoder and the classifier were then „frozen” - they received no further adjustment when applied to the CHB-MIT dataset.

Latent distance

For each EEG window analysed from the CHB-MIT recording, a latent distance score was computed. This score represents the Euclidean distance between the numerical fingerprint of that window and the average fingerprint of all normal EEG segments in the training dataset. A high latent distance indicates that the brain state captured in that window differs substantially from the model's learned representation of normal activity; a low latent distance indicates proximity to normality.

Crucially, this is a continuous, real-valued score, not a binary classification. It can therefore rise and fall smoothly over time, tracking the dynamic evolution of brain state - including gradual transitions between the interictal, ictal, and post-ictal phases.

Model performance was assessed using three standard diagnostic metrics: (1) the area under the receiver operating characteristic curve (AUC-ROC), which quantifies the model's overall ability to discriminate between seizure and non-seizure windows; (2) the area under the precision-recall curve (AUPRC), which is particularly informative when seizure events are rare relative to normal recordings; and (3) balanced accuracy, which accounts for class imbalance. An optimal classification threshold was identified using the Youden J

Table 2. Model performance on external validation (CHB-MIT dataset)

Metric	Value
AUC (ROC)	0.798 (95% CI: 0.730 - 0.945)
AUPRC (Precision-Recall AUC)	0.973
Balanced Accuracy	81.0%
Optimal classification threshold	0.072

statistic, which selects the cut-off that maximises the sum of sensitivity and specificity. The 95% confidence interval for the AUC was estimated using bootstrap resampling with 1,000 iterations.

RESULTS

The model was tested on 90 EEG windows from the CHB-MIT recording: 21 ictal (seizure) and 69 interictal (normal) windows. None of these data had been seen by the model during training. The results are summarised in Table 2.

AUC of 0.80: An AUC of 0.80 means that if the model were presented with one randomly selected seizure window and one randomly selected normal window, it would correctly rank the seizure as more abnormal in 80% of cases. This indicates good discriminative ability, particularly given that the model was tested on patients and recording equipment it had never encountered.

AUPRC of 0.97: In any EEG recording, seizures are rare events compared to the total duration of normal activity. Standard accuracy metrics can be misleading in this context, because a model that simply labels everything as „normal” would still appear highly


accurate.

The precision-recall AUC addresses this problem: a value of 0.97 indicates that when the model classifies a window as a seizure, that classification is correct in the vast majority of cases - that is, the rate of false alarms is very low. This property is directly relevant to clinical use, where frequent false alarms rapidly erode trust in automated monitoring systems.

Confidence interval: The upper bound of the 95% confidence interval for the AUC reached 0.945, confirming that the observed performance is unlikely to be due to chance.

Classification threshold: The optimal threshold of 0.072 differs markedly from the intuitively expected value of 0.5. This shift is explained by systematic differences between the training and test datasets - a phenomenon known as domain shift - arising from differences in patient age, recording equipment, and clinical setting. Adjusting the threshold to compensate for these differences is a standard and expected step in applying models across different clinical environments.

When the model's output probabilities were plotted for seizure and normal windows separately, a clear - though not complete - separation was visible. Normal windows clustered around low probability values, while ictal windows showed substantially higher values. The partial overlap between the two distributions is expected in real clinical data, reflecting the continuous nature of seizure evolution: a seizure does not begin abruptly at a single moment but develops gradually, with some windows capturing transitional states that are

Figure 1. Clinical evaluation of zero-shot transfer performance: discrimination between ictal and interictal EEG windows in an independent external dataset (CHB-MIT)

neither fully normal nor fully ictal.

The diagnostic performance of the model is summarised in Figure 1, which presents the ROC curve, precision–recall curve, confusion matrix, and output probability distributions for the external validation dataset.

Figure 2 shows the most clinically relevant finding of this study. For the 600-second window surrounding the documented seizure, the latent distance was computed for each successive EEG window and plotted as a continuous time series. The resulting trajectory exhibits a pattern that is both statistically clear and physiologically meaningful:

- Interictal period (before second 2,996): latent distance was stable and low, indicating that the brain's electrical activity closely resembled the model's learned representation of normal brain activity.

- Pre-ictal phase: latent distance began to increase before the clinically documented seizure onset, potentially reflecting a gradual increase in cortical excitability during the pre-ictal period.

- Ictal period (seconds 2,996–3,036): latent distance reached its maximum, potentially reflecting with the time window of the clinically verified seizure.

- Post-ictal period: latent distance gradually declined, reflecting the brain's return towards its baseline state - consistent with the known post-ictal period of reduced neuronal excitability.

This finding demonstrates that the model does not merely classify EEG segments as seizure or no seizure, but continuously tracks the dynamic evolution of brain state over time - a property essential for the proposed application in therapy monitoring.

DISCUSSIONS

The principal finding of this study is that a self-supervised deep learning model, trained without any labelled seizure data, spontaneously learns an internal representation of EEG signals in which the continuous distance from a normal brain state follows a clinically interpretable and physiologically coherent trajectory. This behavior was observed in a patient whose data were never used during training, and whose recordings were obtained using different equipment, at a different sampling frequency, and in a different patient population than the training dataset.

The significance of this observation extends beyond the detection performance metrics. A binary seizure detector, regardless of its accuracy, can only indicate whether a seizure is occurring at a given moment. Latent distance, in contrast, provides a continuous, quantitative signal that rises before the seizure, peaks during it, and declines afterwards. Such temporal resolution has the potential to capture information about the dynamics of the epileptic process - information that is lost when brain activity is reduced to a binary label.

The high precision-recall AUC (0.97) further supports the clinical relevance of this approach. In any EEG monitoring context, false-positive alarms - events flagged as seizures that are in fact normal activity - are a major practical obstacle. Excessive false alarm rates are one of the primary reasons that automated EEG monitoring systems have not been widely adopted in clinical practice. The very low false alarm rate suggested by the AU-

Clinical Evaluation of Zero-Shot Transfer (Bonn → CHB-MIT)




Figure 2. Continuous quantification of brain-state deviation from normality (latent distance) over a 600-second EEG recording interval encompassing a documented epileptic seizure

PRC value of this model represents a meaningful advantage over many existing approaches [14,15,16,17,18].

Although preliminary, the present findings raise an important clinical question: could latent distance serve as an objective, real-time indicator of antiepileptic drug effect?

The current standard for assessing treatment response in epilepsy - counting seizures over weeks or months - has a fundamental structural limitation.

A drug that reduces cortical excitability may be exerting a measurable physiological effect on the EEG long before that effect is reflected in a change in seizure frequency [9,10,11,12,13]. Conversely, temporary reduction in seizure frequency due to natural variability may lead to an incorrect impression of treatment response. Both scenarios have the potential to result in suboptimal therapeutic decisions.

If latent distance tracks the brain's physiological state on a continuous timescale, a drug that reduces excitability might be expected to lower the baseline latent distance - to shift the EEG signature of the resting brain closer to a normal physiological state - even in the absence of seizures. In principle, this change could be detectable within days of starting or adjusting medication, rather than after the weeks or months currently required for reliable assessment of seizure frequency. This would represent a fundamentally different paradigm for treatment monitoring.

Antiepileptic drug therapy is guided in part by pharmacokinetic principles - the relationship between drug dose, plasma concentration, and clinical effect. Therapeutic drug monitoring, which involves measuring plasma drug levels to ensure they fall within an effective range, is standard practice for many antiepileptic medications [29,30]. However, plasma concentration reflects drug availability, not the drug's actual effect on the brain's electrical activity. The relationship between a given plasma concentration and a given clinical or neurophysiological effect varies substantially between patients.

We hypothesise that latent distance could serve as a bridge between pharmacokinetics (what the drug does to the body) and pharmacodynamics (what the drug does to the brain). If a consistent relationship exists between plasma drug levels and the latent distance of the EEG at steady state, latent distance

could provide a non-invasive, continuous proxy for pharmacodynamic effect - information that plasma levels alone cannot provide.

The combination of latent distance monitoring with routine therapeutic drug monitoring could therefore enable more precise, individualised dose optimisation.

It must be emphasised that this hypothesis has not been tested in the present study, as neither dataset contains information on drug administration or plasma concentrations. The hypothesis is derived from first principles and from existing evidence that EEG spectral characteristics are sensitive to antiepileptic drug levels [9,10,11]. It requires prospective experimental verification.

This study has several important limitations that must be considered when interpreting the findings.

Limited scope of external validation

External validation was performed on a single EEG recording from one patient, containing one documented seizure. While this is sufficient as a proof of concept, it does not provide the statistical power needed for definitive clinical conclusions. Performance may vary considerably across patients, seizure types, and recording conditions.

Absence of pharmacological data

Neither dataset contains information on antiepileptic medication, dosing, or plasma drug concentrations. The central clinical hypothesis of this paper - that latent distance reflects pharmacodynamic drug effect - therefore remains untested. Future studies must incorporate this information prospectively.

Domain shift

Systematic differences between the training and validation datasets (patient age, recording equipment, sampling frequency, electrode configuration) required adjustment of the classification threshold and may limit the model's out-of-the-box applicability in new clinical settings. Some degree of calibration will likely be needed when deploying this approach in different hospitals or with different EEG systems.

Single-channel analysis

The model was developed and validated using a single EEG channel. Clinical EEG recordings typically involve multiple channels, and different seizure types originate from different

brain regions. Extending the model to multi-channel input could substantially improve both detection performance and physiological interpretability.

Absence of clinical trial data

This work presents a computational proof of concept, not a clinical study. No conclusions can be drawn about patient outcomes, clinical utility, or safety. Prospective clinical studies - ideally randomised controlled trials comparing standard treatment monitoring with EEG latent distance-guided monitoring - are required before any clinical recommendations can be made.

The present findings support the design of a prospective clinical study in which patients starting or changing antiepileptic therapy undergo simultaneous EEG monitoring and serial plasma drug level measurements. The primary research question would be whether latent distance at steady state correlates with plasma drug concentration, and whether early changes in latent distance predict subsequent seizure frequency response. Such a study would require collaboration between neurologists, clinical neurophysiologists, and clinical pharmacologists, and would need to address the practical challenges of ambulatory EEG monitoring in a real-world outpatient setting.

In parallel, the model should be extended to multi-channel EEG, validated across a wider range of seizure types and antiepileptic drugs, and evaluated in paediatric as well as adult populations. The development of a calibration procedure that allows the model to be adapted to a new recording system without retraining would be an important step towards clinical deployment.

CONCLUSION

This proof-of-concept study demonstrates that a deep learning model trained without manually labelled data can learn a continuous representation of EEG-based brain state that generalises to patients and recording equipment not seen during training. When applied to an independent external dataset, the model achieved good seizure detection performance (AUC 0.80, AUPRC 0.97, balanced accuracy 81%), with a very low false alarm rate.

More importantly, the latent distance

measure - a continuous score quantifying how far the EEG signal is from a normal brain state - followed a physiologically coherent trajectory around the documented seizure: rising in the pre-ictal phase, peaking during the seizure, and declining post-ictally. This behaviour was observed without any adaptation of the model to the new data.

We propose that latent distance may represent a novel EEG biomarker with potential utility in the monitoring of antiepileptic therapy. If future prospective studies confirm a relationship between latent distance and pharmacodynamic drug effect, this approach could provide clinicians with an objective, continuous, and non-invasive tool for earlier assessment of treatment response - without waiting for changes in seizure frequency.

Such an approach may be particularly valuable for patients with drug-resistant epilepsy, in whom timely identification of ineffective therapies is critical, and could contribute to more individualised treatment optimisation based on objective neurophysiological measurements.

DISCLOSURE OF AI ASSISTANCE

The Python code used for EEG signal analysis and model training was developed with the assistance of an artificial intelligence tool (Claude, Anthropic, 2024–2025), used solely for code generation and optimisation. All scientific hypotheses, interpretation of results, and clinical conclusions represent the original intellectual contribution of the authors and are based on expert knowledge and independent analysis of the results.

FIGURE LEGENDS

Figure 1

The figure comprises four panels illustrating different aspects of model performance on data not seen during training. (A) Receiver operating characteristic (ROC) curve. The horizontal axis shows the false positive rate (1-specificity); the vertical axis shows the true positive rate (sensitivity). The area under the ROC curve (AUC) was 0.80 (95% CI: 0.73–0.95), indicating good overall discriminative ability. The black point marks the optimal classification threshold (0.072), selected using the Youden J statistic to maximise the sum of sensitivity and specificity. The position of this

threshold substantially below 0.5 reflects systematic differences between the training and test datasets (domain shift), a recognised and expected consequence of cross-institutional model deployment. (B) Precision–recall curve. The horizontal axis shows recall (sensitivity); the vertical axis shows precision (positive predictive value). The area under the precision–recall curve (AUPRC) was 0.97, indicating that the vast majority of windows classified as ictal represented true seizure activity. This metric is particularly informative in settings where seizure events are rare relative to the total recording duration. (C) Confusion matrix. The matrix summarises classification outcomes at the optimal threshold: 53 true positive (ictal correctly classified), 10 false negative (ictal windows missed), 2 false positive (normal windows misclassified as ictal), and 5 true negative (normal correctly classified). The very low false positive rate has direct clinical relevance, as frequent false alarms are a primary obstacle to the clinical adoption of automated EEG monitoring. (D) Probability score distribution. The horizontal axis shows the model-assigned probability of ictal activity (0–1); the vertical axis shows the kernel density estimate. The green distribution corresponds to interictal (normal) windows and the red distribution to ictal (seizure) windows; the vertical dashed line indicates the classification threshold. The partial overlap between distributions reflects the gradual, continuous nature of seizure evolution and the domain shift inherent in zero-shot cross-dataset evaluation.

Figure 2

The horizontal axis represents time (seconds) from the beginning of the analysed EEG segment. The vertical axis represents the latent distance score – a continuous, real-valued measure of how far the model’s internal representation of the current brain state deviates from the mean representation of normal EEG activity established during training. The solid dark line depicts the latent distance trajectory computed for successive 23.6-second EEG windows. The orange dashed line indicates the optimal classification threshold (0.072), above which the model assigns an increased probability of seizure activity. The red-shaded region corresponds to the clinically verified seizure interval (2,996–3,036 seconds). In the interictal period, latent distance remains low and stable, indicating a brain state consistent with

normality. A gradual rise in latent distance is observed in the pre-ictal phase, suggesting an increase in cortical excitability prior to electrographic seizure onset. The latent distance reaches its maximum during the ictal period and then progressively returns towards baseline in the post-ictal phase, consistent with the well-characterised period of reduced neuronal excitability following a seizure. This continuous trajectory illustrates the potential of latent distance as a pharmacodynamic biomarker: effective antiepileptic therapy would be expected to reduce the baseline latent distance and delay or prevent the transition into the high-excitability zone above the threshold.

CONFLICT OF INTEREST

All authors declare no conflict of interest.

REFERENCES

1. Engel J. *Epilepsy: A Comprehensive Textbook*. Philadelphia: Lippincott Williams & Wilkins; 2008.
2. Shorvon S, Perucca E, Engel J. *The Treatment of Epilepsy*. 3rd ed. Oxford: Wiley-Blackwell; 2015.
3. Niedermeyer E, Lopes da Silva F. *Electroencephalography: Basic Principles, Clinical Applications, and Related Fields*. 6th ed. Philadelphia: Lippincott Williams & Wilkins; 2011.
4. Sanei S, Chambers JA. *EEG Signal Processing*. Chichester: Wiley; 2007.
5. Schomer DL, Lopes da Silva F. *Niedermeyer’s Electroencephalography: Basic Principles, Clinical Applications, and Related Fields*. 7th ed. Oxford: Oxford University Press; 2018.
6. Devinsky O, Vezzani A, O’Brien T. *Epilepsy*. Oxford: Oxford University Press; 2019.
7. Stafstrom CE, Carmant L. *Seizures and Epilepsy in Childhood: A Clinician’s Guide*. Cambridge: Cambridge University Press; 2015.
8. Fisher RS. *Epileptic Seizures and Epilepsy*. Oxford: Oxford University Press; 2005.
9. Pitkänen A, Löscher W, Vezzani A, Becker AJ, Simonato M, Lukasiuk K, Gröhn O, Bankstahl JP, Friedman A, Aronica E, Gorter JA, Ravizza T, Sisodiya SM, Kokaia M, Beck H. Advances in the development of biomarkers for epilepsy. *Lancet Neurol*. 2016;15(8):843–56.
10. Engel J Jr, Bragin A, Staba R. Brain-state monitoring and epileptic biomarkers. *Epilepsia*. 2018;59(Suppl 2):160–71.
11. Iasemidis LD. Epileptic seizure prediction and control. *IEEE Trans Biomed Eng*. 2003;50(5):549–58.

12. Poh MZ, Loddenkemper T, Reinsberger C, Swenson NC, Goyal S, Sabtala MC, Madsen JR, Picard RW. Convulsive seizure detection using a wrist-worn electrodermal activity and accelerometry biosensor. *Epilepsia*. 2012;53(5):e93-7.
13. Henshall DC, Hamer HM, Pasterkamp RJ, Goldstein DB, Kjems J, Prehn JHM, Schorge S, Lamottke K, Rosenow F. MicroRNAs in epilepsy: pathophysiology and clinical utility. *Lancet Neurol*. 2016;15(13):1368-76.
14. Andrzejak RG, Lehnertz K, Mormann F, Rieke C, David P, Elger CE. Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state. *Phys Rev E*. 2001;64(6):061907.
15. Xun G, Jha K, He L, Goyal P, Zhang A. Epileptic seizure detection with EEG textual features and graph-based learning. *Proceedings of the 2016 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*; 2016 Dec 15-18; Shenzhen, China. IEEE; 2016.
16. Bao FS, Liu X, Zhang C. PyEDFlib: A Python library for EDF+/BDF+ files. *Proceedings of the 2009 International Joint Conference on Neural Networks*; 2009 Jun 14-19; Atlanta, GA, USA. IEEE; 2009. p. 1-8.
17. Lu D, Triesch J. Residual deep convolutional neural network for EEG signal classification in epilepsy. *arXiv*. 2019:1903.08100.
18. Shoji T, Itahashi T, Ikeda T, Hasegawa C, Yoshimura Y, Ito T. Compact convolutional neural network model for EEG-based abnormality detection. *Front Neurosci*. 2021;15:604385.
19. Zhang X, Yao L, Wang X, Monaghan J, Mcalpine D, Zhang Y. A survey on deep learning-based non-invasive brain signals: recent advances and new frontiers. *J Neural Eng*. 2024;21(1):011002.
20. Wu Y. Deep learning in epilepsy: a bibliometric and visualisation analysis. *Front Neurol*. 2025;16:1542873.
21. Wang Z, Yang J, Sawan M. A review of EEG-based epileptic seizure prediction methods. *Neurophysiologie Clinique*. 2024;54(1):102908.
22. Wong S, Simmons A, Rivera-Villicana J, Blunt S, Muthukumaraswamy S. EEG datasets for seizure detection and prediction – a review. *Epilepsia Open*. 2023;8(2):252-67.
23. Chen T, Kornblith S, Norouzi M, Hinton G. A simple framework for contrastive learning of visual representations. *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*; 2020 Jul 13-18; Vienna, Austria. PMLR; 2020. p. 1597-607.
24. He K, Fan H, Wu Y, Xie S, Girshick R. Momentum contrast for unsupervised visual representation learning. *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*; 2020 Jun 14-19; Seattle, WA, USA. IEEE; 2020. p. 9729-38.
25. Qi L, Wang Z, Chen X, Li J, Zhang Y. Contrastive learning for cross-patient seizure prediction from EEG. *IEEE Trans Neural Syst Rehabil Eng*. 2025;33:984-94.
26. Goldberger AL, Amaral LAN, Glass L, Hausdorff JM, Ivanov PCh, Mark RG, Mietus JE, Moody GB, Peng CK, Stanley HE. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation*. 2000;101(23):e215-20.
27. Shoeb AH. Application of machine learning to epileptic seizure onset detection and treatment [dissertation]. Cambridge (MA): Massachusetts Institute of Technology; 2009.
28. Andrzejak RG, Lehnertz K, Mormann F, Rieke C, David P, Elger CE. Bonn EEG dataset [internet]. Bonn: Department of Epileptology, University of Bonn; 2001 [cited 2025]. Available from: <https://www.ukbonn.de/epileptologie/arbeitsgruppen/ag-lehnertz-neurophysik/downloads/>
29. Patsalos PN, Berry DJ, Bourgeois BFD, Cloyd JC, Glauser TA, Johannessen SI, Leppik IE, Tomson T, Perucca E. Antiepileptic drugs – best practice guidelines for therapeutic drug monitoring: a position paper by the subcommission on therapeutic drug monitoring, ILAE Commission on Therapeutic Strategies. *Epilepsia*. 2008;49(7):1239-76. Updated: Patsalos PN. Antiepileptic drug pharmacokinetics: specific considerations for therapeutic drug monitoring. *Ther Drug Monit*. 2013;35(4):387-404.
30. Perucca E. Clinical pharmacology and therapeutic use of the new antiepileptic drugs. *Fundam Clin Pharmacol*. 2001;15(6):405-17. Updated: Perucca E, Tomson T. The pharmacological treatment of epilepsy in adults. *Lancet Neurol*. 2011;10(5):446-56.

Kontinuirani EEG monitoring antiepileptične terapije primenom samonadgledanog dubokog učenja: studija dokaza koncepta

Branislav S. Čeperković, Slobodan M. Janković

Fakultet medicinskih nauka, Univerzitet u Kragujevcu, Kragujevac, Srbija

KRATAK SADRŽAJ

Uvod: Epilepsija pogađa više od 50 miliona ljudi u svetu. Kod oko jedne trećine bolesnika, napadi se ne mogu adekvatno kontrolisati lekovima. Ključni izazov u kliničkoj praksi je što se efikasnost leka može proceniti tek nakon nedelja ili meseci posmatranja - tokom kojih bolesnik možda prima suboptimalnu terapiju. Elektroencefalogram (EEG) pruža direktan, kontinuiran uvid u električnu aktivnost mozga, ali se u rutinskoj praksi gotovo isključivo koristi za dijagnostiku, a ne za praćenje terapije.

Cilj: Razviti i testirati računarski model koji kontinuirano kvantifikuje koliko daleko je EEG signal od normalnog stanja mozga - mera koju nazivamo latentnom distancom - i proceniti da li ova mera može služiti kao biomarker dejstva antiepileptika u realnom vremenu.

Materijal i metode: Model dubokog učenja treniran je na javno dostupnom skupu EEG podataka (Bon, Nemačka) primenom samonadziranog pristupa koji ne zahteva ručno obeležavanje podataka. Model je zatim testiran na potpuno nezavisnom skupu podataka (CHB-MIT, Boston, SAD) koji sadrži EEG snimke pedijatrijskih bolesnika s epilepsijom otpornom na lekove. Performanse su procenjene standardnim metrikama dijagnostičke tačnosti.

Rezultati: Na spoljnom validacionom skupu podataka, model je postigao površinu ispod ROC krive (AUC) od 0,80 (95% CI: 0,73-0,95), AUC za preciznost i osetljivost od 0,97 i uravnoteženu tačnost od 81%. Latentna distanca je pratila klinički smislenu putanju: rasla je pre zabeleženog početka napada, dostižući maksimum tokom napada, i postepeno se vraćala na bazalni nivo u postictalnom periodu.

Zaključak: Ova studija dokaza koncepta pokazuje da samonadgledani EEG model može da otkrije napade i kontinuirano prati odstupanja od normalne aktivnosti mozga kod bolesnika koje nikada nije video tokom obuke. Latentna distanca bi mogla predstavljati novi, kontinuirani EEG biomarker s potencijalnom upotrebom u praćenju antiepileptičke terapije u realnom vremenu - posebno za identifikaciju dejstva leka pre nego što promene u učestalosti napada postanu očigledne. Neophodne su prospektivne kliničke studije za potvrdu ove hipoteze.

Ključne reči: epilepsija, EEG, antiepileptici, biomarker, duboko učenje, detekcija napada

Received: May 4, 2026

Accepted: May 26, 2026