

Long-Term EEG Partitioning for Seizure Onset Detection

Zheng Chen¹, Yasuko Matsubara¹, Yasushi Sakurai¹, Jimeng Sun^{2,3}

¹SANKEN, Osaka University

²University of Illinois Urbana-Champaign

³Carle Illinois College of Medicine, University of Illinois Urbana-Champaign
{chenz, yasuko, yasushi}@sanken.osaka-u.ac.jp; jimeng@illinois.edu

Abstract

Deep learning models have recently shown great success in classifying epileptic patients using EEG recordings. Unfortunately, classification-based methods lack a sound mechanism to detect the onset of seizure events. In this work, we propose a two-stage framework, *SODor*, that explicitly models seizure onset through a novel task formulation of subsequence clustering. Given an EEG sequence, the framework first learns a set of second-level embeddings with label supervision. It then employs model-based clustering to explicitly capture long-term temporal dependencies in EEG sequences and identify meaningful subsequences. Epochs within a subsequence share a common cluster assignment (normal or seizure), with cluster or state transitions representing successful onset detections. Extensive experiments on three datasets demonstrate that our method can correct misclassifications, achieving 5%-11% classification improvements over other baselines and accurately detecting seizure onsets.

Introduction

Epilepsy affects 60 million of the population worldwide, and approximately 40% of patients have drug-resistant epilepsy with recurrent seizures that cannot be controlled by available medications (WHO 2024). This problem leads to an increased risk of sudden death. Deep learning models have demonstrated impressive success in automating seizure detection using electroencephalogram (EEG) data (Yang et al. 2023; Ho and Armanfard 2023; Chen et al. 2023). Successful methods typically form a classification task. They divide EEG recordings into a sequence of second-level epochs and aim to classify them accurately. While state-of-the-art (SOTA) performance has been demonstrated, *these methods cannot provide inherent information of seizure detection research, i.e., the seizure onset (SO)*. Clinically, many epileptic patients benefit from accurate SO detection as it helps localize and surgically remove the onset zone in the brain, which exhibits the earliest electrophysiological changes during a seizure event. Also, successful detection provides optimal timing to adjust abnormal electrical activities in the brain by neuromodulatory devices (Conrad et al. 2019). Despite great importance, the objective of classification-based methods is

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

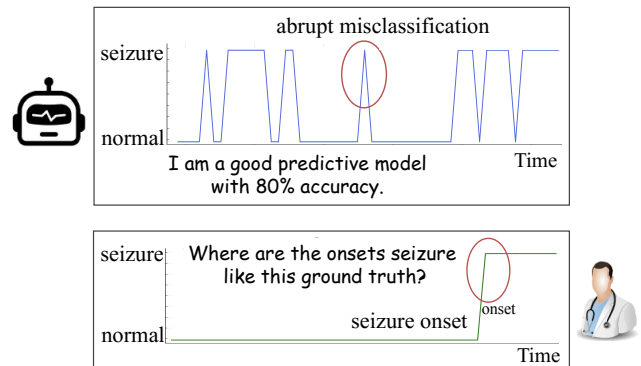


Figure 1: A visualization of 55 two-second epochs using a second-level classification method shows some unexpected abrupt misclassifications. This issue may lead to unexplainable outcomes for clinicians.

to determine whether a seizure exists within an EEG, without explicitly modeling and showing SO position over a long sequence. Several abrupt misclassifications randomly appear within a state-consistent sequence, as shown in Figure 1. These misclassifications inevitably increase false alarms of SO detection and lead to unexplainable outcomes.

A few studies propose a two-stage method (Burrello et al. 2020; Li et al. 2021b; Boo et al. 2021; Batista et al. 2024). They employ post-processing to smooth abrupt changes and reduce the false detection rate. Samples within a sliding window are re-assigned a consistent label based on a majority vote. However, several limitations remain unsolved.

- **Still lacking explicit SO detection modeling.** The objective of post-processing is to smooth misclassifications. No existing works effectively formulate SO detection as a learning task and directly output the SO information. They are required to manually set parameters, e.g., voting threshold and window size, which hinder scalability.
- **Insufficient feature utilization.** Existing works estimate a simple statistical observation of label assignments. They ignore the underlying features that characterize an EEG sample and its relation to a seizure event. Such empirical observations fail to offer any explanation for detection.
- **Lacking long-sequence dependencies modeling.** Either

classification-based methods or voting operations treat all second-level epochs equally and uniformly. They fail to take into context information or long-term dependencies within a sequence. This often yields suboptimal results.

To tackle these, we propose a new **seizure onset detector**, `SODOR`, with the following contributions:

- **Subsequence clustering formulation for explicit SO detection.** We propose a deep clustering method to explicitly model the SO detection task, which is formulated as a subsequence clustering problem to *find* and *segment* the subsequences that characterize a consistent state (normal vs. seizure) automatically. The segment points show state transitions that are identified as SO timestamps.
- **Channel logits representations.** We propose a channel correlation representation and successfully formulate it as the clustering objective. This approach benefits from the classification model by leveraging logits of label assignments from EEG channels, rather than smoothing label assignments. By learning the correlations between these logits, the method provides insights into *how multi-channel interactions relate to seizures*.
- **Temporal consistency modeling.** We propose modeling time-invariant interactions within epochs and long-sequence consistency through a clustering constraint. This approach helps mitigate abrupt changes and encourages neighboring epochs to be assigned to the same cluster.

To our knowledge, we are the first to formulate SO detection as a clustering task and explicitly output SO information.

Related Works

Seizure Onset Detection

Existing SO detection methods can be categorized into end-to-end and two-stage approaches. End-to-end methods frame SO detection as a classification task, labeling second-level epochs as either normal or seizure. Various deep learning models such as CNNs (Eldele et al. 2021; Feizbakhsh and Omranpour 2023), Transformers (Chen et al. 2022a,b; Yang et al. 2023; Kotoge et al. 2024), and Graph models (Tang et al. 2022; Ho and Armanfard 2023; Cai et al. 2023) are employed to automate feature extraction and classification. While an accurate model may provide some SO information, these methods do not explicitly model and detect SOs. Some studies set up onset EEG fragments as a separate class for three-way classification (Rasheed et al. 2021; Dissanayake and Fookes 2021; Chen et al. 2023). Since they focus only on classification accuracy, end-to-end methods often result in abrupt misclassifications at random timestamps, increasing false SO detections. Two-stage methods involve post-processing the outputs of classification models. They re-assign consistent state labels to sequences of epochs within a window based on majority voting (Burrello et al. 2020; Li et al. 2021b; Batista et al. 2024) or thresholding (Khan et al. 2018; Boo et al. 2021; Feizbakhsh and Omranpour 2023). However, these works lack explicit learning for SO detection and oversimplify label smoothing, resulting in suboptimal performance that often requires data-specific manual adjustments. Different from existing methods, we

propose a novel formulation of SO detection as a time series subsequence clustering task.

Time Series Subsequence Clustering

Subsequence clustering is an important task in time series data mining. The objective is to distinct states in sequences of time series, without relying on known labels and segments (Matsubara et al. 2014; Obata et al. 2024). For example, in a dance routine, a multivariate time series captures transitions between motion states such as “walk,” “run,” “jump,” and “kick.” Subsequence clustering segments these time observations into concise segments and assigns each segment a motion label. Model-based methods are commonly used, where each cluster is represented as a model, and sequences are fitted to these models (Hallac et al. 2017; Kawabata et al. 2021; Nakamura et al. 2023; Koki et al. 2023). Recent researchers propose some deep clustering involving two stages (Nagano et al. 2019; Wang et al. 2023; Lai et al. 2024). The first stage learns a latent embedding by deep learning models. Then, they view a sequence of embeddings with multiple dimensions as a multivariate time series, and a clustering stage identifies the state of embeddings. `SODOR` follows this line but differs from these methods, which use unexplainable embeddings for clustering. Inspired by (Conrad et al. 2019; Li et al. 2021a), we propose an explainable method that treats the learned channel-independent logits as multivariate time series and models their state-specific structures.

Problem Formulation

This section presents the problem formulation for SO detection using EEG data. We first introduce preliminaries.

Definition 1 (Multi-channel EEG recordings). *EEG data captures neuronal activities from different brain regions over time. Let $\mathcal{X} := \{\mathcal{X}^{(n)}\}_{n=1}^N$ denote a longitudinal EEG set of N patients. For n -th patient, $\mathcal{X}^{(n)} \in \mathbb{R}^{C \times T}$ represents several recordings from C channels over a duration of T time points.*

To prepare for detecting the SO timestamps, each patient recording is segmented into a sequence of second-level epochs by a sliding window, denoted as $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_P]$, each $\mathbf{X}_p \in \mathbb{R}^{C \times L}$ and L is the window size. Each \mathbf{X}_p is associated with a label $\mathbf{y}_p \in \{0, 1\}$.

Definition 2 (Subsequence clustering). *Suppose $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_P]$ is a set of multivariate time series. Subsequence clustering wants to group these P time steps or epochs into K clusters by:*

- *finding a set of M non-overlapping subsequences of \mathbf{X} , i.e., $\tilde{\mathbf{X}} = [\tilde{\mathbf{X}}_1, \dots, \tilde{\mathbf{X}}_M]$, where $M \ll P$.*
- *estimate the model representations of K clusters, i.e., $\Theta = \{\Theta_1, \dots, \Theta_K\}$.*
- *assign each subsequence to one cluster, Θ_k i.e., $\tilde{\mathbf{y}} = \{\tilde{y}_1, \dots, \tilde{y}_K\}$ where $\tilde{y}_k \subset \{1, \dots, P\}$.*

Epochs in a subsequence $\tilde{\mathbf{X}}_m$ are consecutive and each subsequence is dependent on its neighbors. That is, $\tilde{\mathbf{X}}_m = \{t_s, t_e\}$ represents the starting and ending timestamps, serving SO detection. Moreover, model-based methods define

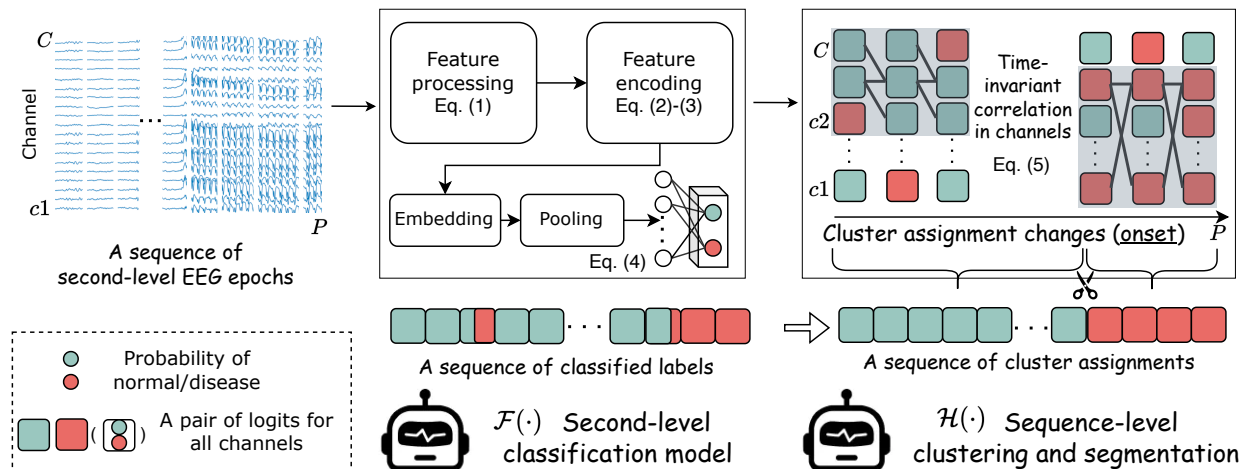


Figure 2: System overview: SODOOR is a two-stage framework designed to explicitly detect seizure onset, consisting of a classification model ($\mathcal{F}(\cdot)$) and a subsequence clustering model ($\mathcal{H}(\cdot)$). $\mathcal{F}(\cdot)$ learns second-level correlation representations for channel-wise logits in normal and seizure states through supervised learning. $\mathcal{H}(\cdot)$ then clusters a sequence of second-level epochs into subsequences. Epochs in a subsequence are consecutive, and each subsequence is dependent on its neighbors. A transition between subsequence/cluster assignments can be viewed as a seizure onset.

each cluster Θ by statistical models, such as Markov chains (Nakamura et al. 2023) or Gaussian (Matsubara et al. 2014). This method provides interpretability to each cluster’s attributes and the reasoning behind the assignments.

Deep subsequence clustering introduces a feature learning stage before clustering. It projects \mathbf{X} into a set of feature vectors $\mathbf{Z} = \{z_1, \dots, z_P\}$. Each z_p , with d dimensions, is treated as an observation from d sensors or channels.

Problem (Seizure onset detection). *Given a recording of segmented EEG epochs \mathbf{X} , seizure detection aims to predict whether a seizure exists. In contrast, SO detection seeks to identify the interval of a seizure event, including the starting (s) and ending (e) positions, represented as $\mathcal{O} = \{(\mathbf{X}_s, \mathbf{X}_e) \mid 1 \leq s < e \leq P\}$, where $(\mathbf{X}_{s-1}, \mathbf{X}_s)$ marks the transition from a normal state (0) to a seizure state (1).*

Typically, this problem is solved by a classification model with post-processing. Our solution, however, is to develop an integrative framework with two distinct learning models. Conceptually, we formulate this problem as a deep subsequence clustering task:

1. $\mathcal{F}(\cdot) : \mathbf{X} \rightarrow \{0, 1\}^{(y)}$, which classifies each epoch, resulting in a set of embeddings \mathbf{Z} that encapsulates the second-level label information $\mathbf{y} = \{y_1, \dots, y_P\}$.
2. $\mathcal{H}(\cdot) : \mathbf{Z} \rightarrow \tilde{\mathbf{y}}$, which partitions a set of subsequences for a sequence and assigns each subsequence a state label, $\tilde{\mathbf{z}} = \{\tilde{z}_1, \dots, \tilde{z}_M\}$ and again $M \ll P$.

For a seizure interval (i.e., a subsequence), we aim to achieve $\mathcal{O} = \{\tilde{z}_s, \tilde{z}_e\}$ where $\forall \tilde{y} \in \mathcal{O}$, $\tilde{y} = 1$ and \tilde{z}_s explicitly provides SO timestamp. However, existing methods use unexplainable feature embeddings as the cluster targets (as described in Definition 2). In contrast, we aim for the discovered subsequence clusters to be interpretable because, in

epileptic research, it is often more important to learn discriminative and interpretable patterns that reflect seizures.

Proposed Method

This section presents our SODOOR framework, as shown in Figure 2. Specifically, it comprises two phases: second-level representation learning and sequence-level clustering aimed at explicitly detecting SO in EEG recordings.

Second-Level Representation Learning

Seizure is fundamentally a network disease (Burns et al. 2014). The goal of this graph model is to characterize this network and show how it is relevant to a seizure event. A network constructed from EEG data $\mathbf{X} = [\mathbf{X}_1, \dots, \mathbf{X}_P]$ can be represented as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ spanning all P epochs. \mathcal{V} denotes the set of vertices, corresponding to the EEG channels, and \mathcal{E} represents the set of edges, which capture correlations between channels. An edge is defined as $e_{ij} = (v_i, v_j)$, where $i, j \in C$.

Channel Correlation for Graph Construction To construct a graph, we extract frequency features in each channel as the node embeddings. We apply the Fast Fourier Transform (FFT) to each normalized epoch \mathbf{X}_p to decompose the signal into its frequency spectrum. The absolute values, $\mathbf{x}_p = \{x_p^{(1)}, \dots, x_p^{(|V|)}\}$, are then extracted to represent the node set \mathcal{V} . For edge initialization, we employ a dynamic connectivity modeling method (Tang et al. 2022). We compute the absolute value of the normalized cross-correlation between embeddings $(x_p^{(i)}, x_p^{(j)})$ of each node pair (v_i, v_j) .

Specifically, we calculate each edge weight e_{ij} as follows:

$$e_{ij} := \begin{cases} \mathbf{x}_p^{(i)} \circ \mathbf{x}_p^{(j)} & \text{if } v_j \in \mathcal{V}_{\text{top}}(v_i) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where \circ denotes the normalized cross-correlation, and \mathcal{V}_{top} represents the set of neighbor nodes with the highest correlation scores. This results in a sparse graph representing each second-level EEG epoch for subsequent feature encoding.

Channel Logits Representation Learning While SODOR follows the deep clustering paradigm, unlike existing works, it generates explainable feature embeddings as the target of subsequent modeling. Specifically, we propose learning a set of channel-wise logits. These logits encode the information from the original data and represent the probabilities of seizures in various channels. That is, we shift the observation from a Boolean value (in post-process methods) and unexplainable embeddings to probabilities, offering insights into “how multi-channel interactions in an EEG epoch relate to a seizure event.” To learn the logit representations, $\mathcal{F}(\cdot)$ is designed to capture the spatial and temporal dependencies within each epoch. Given \mathbf{X}_p with a processed graph structure, it first learns the spatial information by:

$$\mathbf{Z}_p^{\text{spatial}} = \text{ReLU} \left(f_{\text{conv}} \left(\sum_{k=0}^K \theta_k \cdot (D^{-1}V_p)^k \mathbf{X}_p \right) \right), \quad (2)$$

where V_p denotes the adjacency matrix defined by Eq. (1) and D is the diagonal degree matrix. $f_{\text{conv}}(\cdot)$ denotes a diffusion convolution with a ReLU activation function, referring to (Tang et al. 2022). Hence $(D^{-1}W)^k$ represents the diffusion process across k steps. θ_k are trainable parameters. $\mathbf{Z}_p^{\text{spatial}} \in \mathbb{R}^{C \times S}$ where $S \ll L$, is a set of channel-wise feature embeddings. For each channel, we aggregate frequency messages $\mathbf{x}_p^{(j)}$ across edges by $\frac{1}{|\mathcal{V}_{\text{top}}(i)|} \sum_{j \in \mathcal{V}_{\text{top}}(i)} \mathbf{x}_p^{(j)}$ where $\mathcal{V}_{\text{top}}(i)$ represents the set of neighbors of the node i defined by normalized cross-correlations. Afterward, we use a recurrent neural network to model temporal dependencies along with S time steps:

$$\{\mathbf{z}_p^{(1)}, \dots, \mathbf{z}_p^{(C)}\} = \sigma(\text{GRU}(\mathbf{Z}_p^{\text{spatial}})), \quad (3)$$

The c -th channel $\mathbf{z}_p^{(c)}$ contains a pair of logits for normal or seizure, obtained by using a softmax $\sigma(\cdot)$ to the hidden state of the last time step. A max-pooling layer selects the channel with the maximum logit for loss calculation:

$$\mathcal{L}_{\text{BCE}} := - [y \log(\mathbf{z}_p^{\max}) + (1 - y) \log(1 - \mathbf{z}_p^{\max})] \quad (4)$$

where \mathbf{z}_p^{\max} represents the maximum logit value across all channels. Pooling operations are commonly used to aggregate node features, with each method making different assumptions about the graph structure. We assume that max-pooling retains only the most prominent signals, which may lead to better performance, particularly in tasks where certain strong features are indicative of a seizure event. This assumption aligns with findings suggesting that a few unique, abnormal connections across EEG channels can serve as SO markers (Li et al. 2021a; Boo et al. 2021). A comprehensive evaluation is provided in the Experiment section.

Sequence-Level Clustering for SO Detection

Formulating SO Detection as A Subsequence Clustering

Given second-level representations $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_P\}$ with C multivariate sequences, our goal is to cluster and segment them into subsequences $\tilde{\mathbf{Z}} = \{\tilde{\mathbf{Z}}_1, \dots, \tilde{\mathbf{Z}}_M\}$. As described in Definition 2, each subsequence contains several EEG epochs with consistent clustering, and a pair $\{\tilde{\mathbf{Z}}_m, \tilde{\mathbf{Z}}_{m+1}\}$ represents a cluster transition, which facilitates SO detection. In this work, $\mathcal{H}(\cdot)$ is designed on top of a novel Toeplitz Inverse Covariance-based Clustering (Hallac et al. 2017). This method employs a graphical lasso to estimate sparse Gaussian inverse covariance matrices, also known as precision matrices (Obata et al. 2024), to represent cluster models $\{\Theta_k | k = \text{normal, seizure}\}$. Each matrix provides insights into pairwise conditional independencies among EEG channels, determining which correlations contribute most significantly to cluster assignments. The formulation is:

$$\underset{\Theta, \tilde{\mathbf{y}}}{\text{argmin}} \sum_{k=1}^K \left[\sum_{y=1}^{|Y|} \sum_{\mathbf{z}_p \in \tilde{\mathbf{y}}_k} \underbrace{(-\ell\ell(\mathbf{z}_p, \Theta_k))}_{\text{log likelihood}} \right. \quad (5)$$

$$\left. + \underbrace{\beta \mathbb{1}\{\mathbf{z}_{p-1} \notin \tilde{\mathbf{y}}_k\}}_{\text{temporal consistency}} + \underbrace{\|\lambda \odot \Theta_k\|_1}_{\text{time-invariant representation}} \right].$$

where the log-likelihood term measures the probability that the p -th epoch belongs to cluster k by observing its representation \mathbf{z}_p . The temporal consistency term models long-term dependencies, encouraging neighboring epochs to be assigned to the same cluster. $\|\lambda \odot \Theta_k\|_1$ is an ℓ_1 -norm penalty to control the sparseness of Θ_k . This enforces the preservation of cluster-specific correlations between EEG channels, enabling the learning of time-invariant representations.

- Limitation. Notably, when $|Y| = 1$, Eq. (5) denotes the original clustering algorithm (Hallac et al. 2017). However, while the original clustering algorithm operates on multivariate time series as input, our pairwise logit representation is structured as a two-dimensional tensor time series. Next, we address how to define and infer the model Θ_k .

Logits Toeplitz Matrices Instead of clustering each epoch independently, we assume neighboring epochs should be consecutive, so we redefine the “epoch” by a sliding window $\omega \ll P$, represented as $\mathbf{Z}_p := \{\mathbf{Z}_{p-\omega+1}, \dots, \mathbf{Z}_p\}$. Thus, we cluster these short-duration matrices and then fit all variables into Θ , characterized by block Toeplitz inverse covariance matrices. These block-wise constraints are designed to capture time-invariant structural patterns within \mathbf{Z}_p , helping to smooth abrupt changes. A matrix can be expressed as:

$$\Theta_k := \begin{bmatrix} A^{(0)} & (A^{(1)})^P & \dots & \dots & (A^{(w-1)})^P \\ A^{(1)} & A^{(0)} & \ddots & & \vdots \\ A^{(2)} & A^{(1)} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & (A^{(1)})^P & (A^{(2)})^P \\ \vdots & & & A^{(1)} & A^{(0)} & (A^{(1)})^P \\ A^{(w-1)} & \dots & A^{(2)} & A^{(1)} & A^{(0)} \end{bmatrix},$$

where each A represents logit correlations among C channels within ω time observations. An element $a_{i,j} \in A$ refers to the relationship between the i -th and j -th channels at the same ω -th epoch or between the $(\omega - 1)$ -th and ω -th epochs.

However, the original graphical lasso estimates Θ_k based on single observations T . Instead, we formulate a pair of logits as a two-dimensional tensor, treating it as a single observation. Such estimations can become too high-dimensional. (Koki et al. 2023) proposes separating tensor Θ_k into multi-mode, where $\mathbf{a}_{i,j}^{(n)} \in A^{(n)}$ refers to the relationship between the i -th and j -th variables in mode- n . This may lead to over-representation, since the logit for "normal" already implies the probability of "seizure". Since Θ are covariance matrices, we solve this by the linearity property of covariance (Wackery, Mendenhall, and Scheaffer 2008).

Proposition 1. *Given a pair of logits $\{z_{nor}, z_{sei}\}$, denoted as $\{A, \tilde{A}\}$, under the constraint $z_{nor} + z_{sei} = 1$, computing $Cov(A, \tilde{A})$, as it fully captures the covariance relationship between A and \tilde{A} , due to the linearity property and $Cov(A, 1) = 0$, as denoted by $Cov(A, A) \equiv Cov\{A, \tilde{A}\}$.*

The proof is provided in the appendix. Thus, we focus on analyzing the logit representation of the seizure state, which captures pairwise logit correlations. This proof allows us to estimate Θ . Furthermore, this single logit representation can seamlessly serve as the optimization objective for the BCE loss in Eq. (4).

Clustering for SO detection After estimating the cluster representation Θ , we identify the optimal Θ_k for each \mathbf{Z}_p . More importantly, we incorporate temporal consistency by ensuring that consecutive EEG epochs are aligned to a consistent representation, thereby further modeling long-term dependencies in a state-consistent sequence, as denoted by:

$$\text{minimize} \sum_{k=1}^K \sum_{\mathbf{z}_p \in \tilde{\mathbf{y}}_k} -\ell(\mathbf{Z}_p, \Theta_k) + \beta \mathbb{1}\{\mathbf{z}_{p-1} \notin \tilde{\mathbf{y}}_k\}. \quad (6)$$

This formulation jointly maximizes the log-likelihood and maintains temporal consistency. The balance between these objectives is controlled by β . This is an indicator function: when the same cluster $\tilde{\mathbf{y}}_k$ is made between neighboring epochs, there is no penalty, but $\mathbb{1}\{t - 1 \notin P_k\}$ is 1, if \mathbf{Z}_{p-1} does not belong to the same cluster as \mathbf{Z}_p .

Minimizing this constraint ensures that neighboring EEG epochs are assigned to the same clusters. Any assignment that deviates from this constraint indicates a state transition. The corresponding p -index marks the onset of the next state, identifying SO when transitioning from a normal state to a seizure state in the EEG sequence.

Optimization The subsequence clustering is optimized by the expectation-maximization (EM) algorithm to iteratively learn the cluster assignments $\tilde{\mathbf{y}}$ and the structural patterns Θ until convergence. Specifically, in Eq. (5), the log-likelihood term and the sparsity term, which can be considered as a typical graphical lasso problem, have a solution guaranteed to converge to the global optimum using the alternating direction method of multipliers (ADMM) (Boyd et al. 2011). The

clustering is formed by a dynamic programming optimization that finds the minimum cost Viterbi path for a sequence (Hallac et al. 2017). More detailed methods, implementations, and optimizations can be found in the appendix.

Experiments and Results

We evaluate SODOR to determine if it addresses the following questions:

- Can SODOR filter out the false detections, e.g., abrupt misclassifications, and can automatically detect the SO.
- Are the channel logit correlation representations robust and beneficial for SO detection?
- Does SODOR potentially provide explainable clusters?

Datasets

We evaluated SODOR for the seizure onset (SO) detection task on three real-world datasets. In addition, we assessed a more challenging task: seizure prediction, which aims to identify the preictal state preceding seizures. This task is critical in clinical settings, and we provided onset information for it, referred to as PSO (Preictal Seizure Onset).

CHB-MIT comprises 844 hours of continuous scalp EEG data from 22 patients, recorded across 22 channels, with a total of 163 seizure episodes. For the PSO detection task, we preprocessed this dataset by defining the pre-seizure state as the 5 minutes preceding seizure onset.

HUH is collected from University of Helsinki, Finland. It consists of scalp 21-channel EEG data of 79 patients, serving the seizure detection task.

TUSZ dataset is part of the Temple University Hospital EEG Seizure Corpus. It comprises 5,612 EEG recordings with 3,050 clinically annotated seizures. We utilized 19 EEG channels, following the standard 10-20 system.

We divided each dataset into 70%/20%/10% for training, testing, and validation. We stored the IDs of all epochs in the patient recordings, enabling recall in long-term recordings to verify detection accuracy.

Baselines

We compared SODOR with three post-process (PP) SO detection baselines, two classification-based (CIs) seizure detection baselines, and two deep clustering (Clu) methods.

1. (Burrello et al. 2020) proposed a post-processing that uses a sliding 5-second window to re-assign labels based on a patient-specific voting threshold.
2. (Boo et al. 2021) proposed a two-step method involving a deep model to classify second-level epochs first and a weighting phase to score the probabilities of SO.
3. (Li et al. 2021b) proposed an ensemble learning post-process based on four machine learning models and used majority voting to determine the SO threshold.
4. (Tang et al. 2022) proposed a GNN-based method for seizure detection and classification tasks.
5. (Ho and Armanfard 2023) proposed a GNN model incorporating contrastive learning for seizure classification.

Baseline	CHB-MIT			HUH			TUH		
	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC
Burrello, TBME, 2020 (PP)	0.897	0.875	0.882	0.574	0.728	0.661	0.683	0.656	0.717
Boony, TSNRE, 2021 (PP)	0.842	0.744	0.868	0.732	0.715	0.749	0.670	0.692	0.713
Li et al., TSNRE, 2021 (PP)	0.840	0.859	0.870	0.643	0.625	0.639	0.617	0.604	0.699
Siyi et al., ICLR, 2022 (Cls)	0.865	0.867	0.939	0.778	0.744	0.810	0.764	0.743	0.816
Thi et al., AAAI, 2023 (Cls)	0.815	0.817	0.889	0.728	0.694	0.762	0.716	0.695	0.768
Time2State, WWW, 2023 (Clu)	0.770	0.730	0.765	0.632	0.654	0.733	0.664	0.673	0.756
E2Usd, WWW, 2024 (Clu)	0.834	0.815	0.852	0.669	0.686	0.704	0.668	0.677	0.760
SODor	0.979	0.964	0.981	0.873	0.823	0.845	0.842	0.863	0.879

Table 1: Performance comparison on the CHB-MIT, HUH, and TUH datasets. We retain the initial classification model stage and utilize the post-processing and clustering methods from the baseline approaches. **Bold**: best; PP: post-process; Cls: one-step classification-based method; Clu: deep clustering method.

- Time2State (Wang et al. 2023) is a deep clustering method tailored for multivariate time series.
- E2Usd (Lai et al. 2024) formulates each dimension of feature embeddings as a multivariate time series and conducts a subsequence clustering on them.

Experiment: Overall SO Detection

Setup. We compared selected baselines to SODor. For a fair comparison, with (Burrello et al. 2020), (Boo et al. 2021), and (Li et al. 2021b), we maintained our first-stage second-level learning and compared their post-processing methods. For one-step classification methods (Tang et al. 2022) and (Ho and Armanfard 2023), we used their learning models as the backbone and added our subsequence clustering. For Time2State and E2Usd, since they are tailored for time series, we used our classification model and incorporated it into their model-based clustering, the Dirichlet Process Gaussian Mixture Model.

Metrics. We evaluated performance using two clustering metrics: normalized mutual information (NMI) and adjusted Rand index (ARI), along with accuracy (ACC) for assignment analysis.

Results (main). Table 1 presents the main SO detection results across three datasets. SODor outperforms all baseline methods. Specifically, for CHB-MIT, SODor achieved an NMI of 0.979, an ARI of 0.964, and an ACC of 0.981, significantly surpassing the performance of other approaches. The post-processing method (Burrello et al. 2020) performed well on this dataset, achieving an NMI of 0.897, an ARI of 0.875, and an ACC of 0.882; however, SODor clearly demonstrated superior performance. One-step classification methods outperformed the two deep subsequence clustering methods (Time2State and E2Usd), as these clustering methods lacked a temporal consistency term. Post-processing and deep clustering methods performed worse on the TUH dataset, likely due to the presence of diverse seizure types and rapid transitions (Tang et al. 2022). These transitions or preictal phases often contain numerous misclassifications that dominate sequences (Daoud and Bayoumi 2019), presenting significant challenges for majority voting.

Results (visualization). Figure 3 shows two case study visualizations for the TUH and CHB-MIT datasets. Figure 3(a) shows a sequence with 55 epochs. The second row shows the

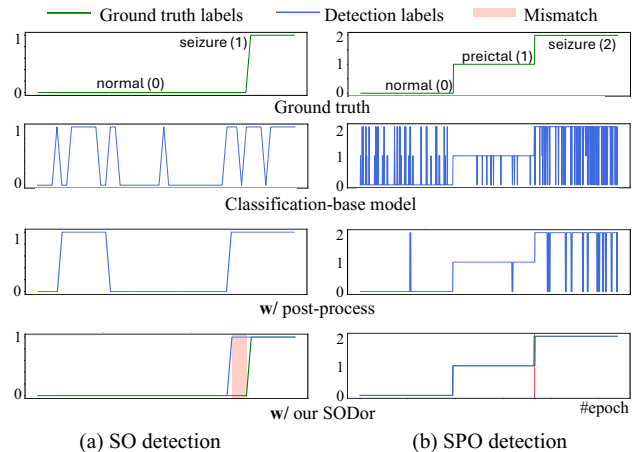


Figure 3: Comparison of (a) SO detection and (b) SPO (Seizure Preictal Onset) detection. The figure illustrates the performance of the classification model ((Tang et al. 2022)), w/ post-processing methods ((Burrello et al. 2020)), and w/ our proposed subsequence clustering approach. The mismatch between SODor and ground truth is marked.

results of the SOTA method (Tang et al. 2022). We applied the post-processing method from (Burrello et al. 2020) to the results of the second classification model using a window $\omega = \{2, 3, 5, 7, 10\}$ and a search of voting threshold within $\{0.5, 0.65, 0.75\}$. The final row shows the results with our subsequence clustering. As shown in Figure 3(a), although SODor exhibits some mismatches between the ground truth and detection labels, it avoids abrupt changes and operates in an automated manner. In contrast, while post-processing reduces false detections, it relies on data-specific manual parameter tuning. Figure 3(b) extends the analysis to SPO detection. SODor also shows significantly fewer mismatches, with minor discrepancies highlighted in the red box.

Experiment: Representation Analysis

Setup. To maintain consistency in the analysis, we further visualized the TUH case used in Figure 3(a), focusing on the logits from different channels after training. As discussed in

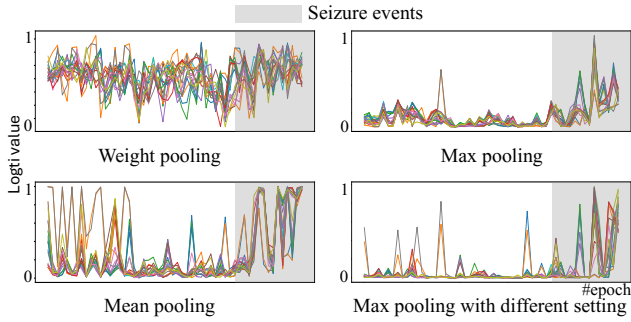


Figure 4: Ablation of pooling methods (i.e., max, mean, and weighted) on channel logits across different training settings. The results show the robustness of max pooling. We visualize logits in max pooling under different parameter settings, and it maintains consistent performance.

the “Proposed Method” section, pooling operations play a critical role in our framework, particularly in summarizing and extracting significant features from channels. To evaluate their impact, we replaced the proposed max pooling with weighted pooling (using an MLP) and mean pooling, keeping the same parameter settings. Interestingly, the final classification accuracy remained unchanged at 0.816.

Results. Figure 4 illustrates the impact of different pooling methods on channel-wise logits. It is clear that the representation generated by SOD_{or} has distinctive characteristics between normal and seizure states. Max pooling exhibits superior robustness compared to other methods. Even though the classification accuracy is the same, weighted pooling distorts the representations, and mean pooling averages the features of graph nodes, which leads to some channels having high probabilities of seizure in normal samples. Moreover, we empirically observe an unstable representation issue during different training settings, but max pooling remains robust across different settings, as shown in the last subfigure.

Experiment: Channel Correlation Visualization

Setup. Since Θ represents inverse covariance matrices and the sparsity term in Eq. (5) controls the sparsity of each cluster model, we extracted the sparse matrices as the adjacency matrix after training. Because our representation preserves the channel index across all learning and clustering stages, we mapped the channels back to the EEG electrodes and graph node positions, following the method outlined in (Tang et al. 2022), to visualize channel correlations in normal and seizure cluster models. For CHB-MIT, which contains more than 19 EEG channels, we used the standard 19-channel system for visualization, consistent with the settings for TUH. Each connection in the visualization corresponds to a “1” that persists in Θ .

Results. Figure 5 compares the visualizations of normal and seizure states in two different cases. The results show a consistent pattern: from sparse to dense connections, indicating that more channels are connected and the brain becomes more active during seizures. The central connection of “C3-CZ-C4” in the normal state still appears during seizures, as




Figure 5: Visualization of learned Θ . The connections become denser from normal to seizure. The green box highlights consistent connections, while the red circle indicates differences between the two clusters.

	SOD_{or}	MDL	DPGMM
NMI	0.979 ± 0.04	0.720 ± 0.13	0.613 ± 0.09
ARI	0.964 ± 0.04	0.694 ± 0.13	0.658 ± 0.09
ACC	0.981 ± 0.04	0.649 ± 0.14	0.626 ± 0.09

Table 2: Ablation study of different subsequence clustering methods: MDL (Matsubara et al. 2014; Koki et al. 2023) and DPGMM (Lai et al. 2024), in CHB-MIT dataset.

shown in Figure 5(a), but more channels are activated. Figure 5(b) shows that the Occipital lobe (O1 and O2) and Parietal lobe (PZ) remain stable across the transitions, suggesting that the seizure may be related to other areas.

Clustering Method Ablation

Tale 2 presents that our method can significantly outperform both MDL and DPGMM, achieving the highest values with the lowest standard deviation. The reason may be that the MDL-based method (Matsubara et al. 2014) focuses on model compression loss without explicitly considering temporal consistency. DPGMM (Lai et al. 2024) focuses more on estimating the number of clusters. Our method can also employ Bayesian information criterion (BIC) to estimate the number of clusters. Since the seizure states are predetermined, this work specifies it to two clusters for SO detection and three for SPO detection.

Conclusion

Several powerful seizure detection models have been proposed, yet no existing works explicitly model the seizure onset, often resulting in unexplainable labeling in long-sequence EEG recordings. We aimed to provide a robust SO detection framework that successfully formulates this task as subsequence clustering, identifying the state (normal or seizure) transition as the SO timestamp. One advantage of this framework is its two-stage learning process, allowing us to fully leverage the high capabilities of existing deep learning methods. Experimental results confirmed that our framework has a strong capacity for SO detection and may have potential for clinical applications.

Acknowledgments

This work was supported by JSPS KAKENHI Grant-in-Aid for Scientific Research Number JP24K20778, JST CREST JPMJCR23M3, and NSF award SCH-2205289, SCH-2014438, and IIS-2034479. We thank Rikuto Kotoge, Xihao Piao, and Kohei Obata for their valuable discussions.

References

- Batista, J.; Pinto, M.; Tavares, M.; Lopes, F.; Oliveira, A.; and Teixeira, C. 2024. EEG epilepsy seizure prediction: the post-processing stage as a chronology. *Scientific Reports*.
- Boo, P.; et al. 2021. ScoreNet: A Neural Network-Based Post-Processing Model for Identifying Epileptic Seizure Onset and Offset in EEGs. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2474–2483.
- Boyd, S. P.; Parikh, N.; Chu, E.; Peleato, B.; and Eckstein, J. 2011. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trends Mach. Learn.*, 1–122.
- Burns, S. P.; Santaniello, S.; Yaffe, R. B.; Jouny, C. C.; Crone, N. E.; Bergey, G. K.; Anderson, W. S.; and Sarma, S. V. 2014. Network dynamics of the brain and influence of the epileptic seizure onset zone. *Proceedings of the National Academy of Sciences*, E5321–E5330.
- Burrello, A.; Schindler, K.; Benini, L.; and Rahimi, A. 2020. Hyperdimensional Computing With Local Binary Patterns: One-Shot Learning of Seizure Onset and Identification of Ictogenic Brain Regions Using Short-Time iEEG Recordings. *IEEE Transactions on Biomedical Engineering*, 601–613.
- Cai, D.; Chen, J.; Yang, Y.; Liu, T.; and Li, Y. 2023. MBrain: A Multi-Channel Self-Supervised Learning Framework for Brain Signals. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 130–141.
- Chen, J.; Yang, Y.; Yu, T.; Fan, Y.; Mo, X.; and Yang, C. 2022a. BrainNet: Epileptic Wave Detection from SEEG with Hierarchical Graph Diffusion Learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, 2741–2751.
- Chen, Z.; Zhu, L.; Jia, H.; and Matsubara, T. 2023. A Two-View EEG Representation for Brain Cognition by Composite Temporal-Spatial Contrastive Learning. In *SDM*, 334–342.
- Chen, Z.; Zhu, L.; Yang, Z.; and Zhang, R. 2022b. Multi-Tier Platform for Cognizing Massive Electroencephalogram. In *IJCAI-22*, 2464–2470.
- Conrad, E. C.; Tomlinson, S. B.; Wong, J. N.; Oechsel, K. F.; Shinohara, R. T.; Litt, B.; Davis, K. A.; and Marsh, E. D. 2019. Spatial distribution of interictal spikes fluctuates over time and localizes seizure onset. *Brain*, 554–569.
- Daoud, H.; and Bayoumi, M. A. 2019. Efficient Epileptic Seizure Prediction Based on Deep Learning. *IEEE Transactions on Biomedical Circuits and Systems*, 804–813.
- Dissanayake, T.; and Fookes, C. 2021. Deep Learning for Patient-Independent Epileptic Seizure Prediction Using Scalp EEG Signals. *IEEE Sensors*, 9377–9388.
- Eldele, E.; Ragab, M.; Chen, Z.; Wu, M.; Kwoh, C. K.; Li, X.; and Guan, C. 2021. Time-Series Representation Learning via Temporal and Contextual Contrasting. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, 2352–2359.
- Feizbakhsh, B.; and Omranpour, H. 2023. Cluster-based phase space density feature in multichannel scalp EEG for seizure prediction by deep learning. *Biomedical Signal Processing and Control*, 105276.
- Hallac, D.; Vare, S.; Boyd, S.; and Leskovec, J. 2017. Toeplitz Inverse Covariance-Based Clustering of Multivariate Time Series Data. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17*, 215–223.
- Ho, T. K. K.; and Armanfard, N. 2023. Self-Supervised Learning for Anomalous Channel Detection in EEG Graphs: Application to Seizure Analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7866–7874.
- Kawabata, K.; Bhatia, S.; Liu, R.; Wadhwa, M.; and Hooi, B. 2021. SSMF: Shifting Seasonal Matrix Factorization. In *Advances in Neural Information Processing Systems*, 3863–3873.
- Khan, H.; Marcuse, L.; Fields, M.; Swann, K.; and Yener, B. 2018. Focal Onset Seizure Prediction Using Convolutional Networks. *IEEE Transactions on Biomedical Engineering*, 2109–2118.
- Koki, K.; et al. 2023. Modeling Dynamic Interactions over Tensor Streams. In *Proceedings of the ACM Web Conference 2023, WWW '23*, 1793–1803.
- Kotoge, R.; Chen, Z.; Kimura, T.; Matsubara, Y.; Yanagisawa, T.; Kishima, H.; and Sakurai, Y. 2024. SplitSEE: A Splittable Self-supervised Framework for Single-Channel EEG Representation Learning. In *IEEE International Conference on Data Mining (ICDM)*.
- Lai, Z.; Li, H.; Zhang, D.; Zhao, Y.; Qian, W.; and Jensen, C. S. 2024. E2Usd: Efficient-yet-effective Unsupervised State Detection for Multivariate Time Series. In *Proceedings of the ACM on Web Conference 2024, WWW '24*, 3010–3021.
- Li, A.; Huynh, C.; Fitzgerald, Z.; Cajigas, I.; Brusko, D.; Jagid, J.; Claudio, A.; Kanner, A.; Hopp, J.; Chen, S.; Haagen, J.; Johnson, E.; Anderson, W.; Crone, N.; Inati, S.; Zaghoul, K.; Bulacio, J.; Gonzalez-Martinez, J.; and Sarma, S. 2021a. Neural fragility as an EEG marker of the seizure onset zone. *Nature Neuroscience*, 24: 1–10.
- Li, C.; Zhou, W.; Liu, G.; Zhang, Y.; Geng, M.; Liu, Z.; Wang, S.; and Shang, W. 2021b. Seizure Onset Detection Using Empirical Mode Decomposition and Common Spatial Pattern. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 458–467.
- Matsubara, Y.; et al. 2014. AutoPlait: automatic mining of co-evolving time sequences. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data, SIGMOD '14*, 193–204.
- Nagano, M.; Nakamura, T.; Nagai, T.; Mochihashi, D.; Kobayashi, I.; and Takano, W. 2019. HVGH: Unsuper-

vised Segmentation for High-Dimensional Time Series Using Deep Neural Compression and Statistical Generative Model. *Frontiers in Robotics and AI*.

Nakamura, K.; Matsubara, Y.; Kawabata, K.; Umeda, Y.; Wada, Y.; and Sakurai, Y. 2023. Fast and Multi-aspect Mining of Complex Time-stamped Event Streams. In *WWW*, 1638–1649.

Obata, K.; Kawabata, K.; Matsubara, Y.; and Sakurai, Y. 2024. Dynamic Multi-Network Mining of Tensor Time Series. In *Proceedings of the ACM on Web Conference 2024*, WWW '24, 4117–4127.

Rasheed, K.; Qadir, J.; O'Brien, T. J.; Kuhlmann, L.; and Razi, A. 2021. A Generative Model to Synthesize EEG Data for Epileptic Seizure Prediction. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2322–2332.

Tang, S.; Dunnmon, J.; Saab, K. K.; Zhang, X.; Huang, Q.; Dubost, F.; Rubin, D.; and Lee-Messer, C. 2022. Self-Supervised Graph Neural Networks for Improved Electroencephalographic Seizure Analysis. In *International Conference on Learning Representations*.

Wackerly, D. D.; Mendenhall, W.; and Scheaffer, R. L. 2008. *Mathematical statistics with applications*, volume 7. Thomson Brooks/Cole Belmont, CA.

Wang, C.; Wu, K.; Zhou, T.; and Cai, Z. 2023. Time2State: An Unsupervised Framework for Inferring the Latent States in Time Series Data. *Proc. ACM Manag. Data*.

WHO. 2024. Cancer - World Health Organization. <https://www.who.int/news-room/fact-sheets/detail/epilepsy>. Accessed: 7 February 2024.

Yang, C.; et al. 2023. BIOT: Biosignal Transformer for Cross-data Learning in the Wild. In *Thirty-seventh Conference on Neural Information Processing Systems*.