# Interpretable and Imbalance-Aware Cognitive State Decoding from Pediatric EEG Using Topological Persistence and Entropy

Waleed Bin Owais[1,*], Herbert F. Jelinek, *SMIEEE* [2], Ibrahim (Abe) M. Elfadel, *SMIEEE* [1]

***Abstract*— *Objective:* Classifying cognitive state from pediatric mobile electroencephalography (EEG) recorded during naturalistic, high-movement behavior is challenging due to motion artifacts, low channel count, and severe imbalance between rest and task conditions. Rather than optimizing accuracy under favorable or imbalance-agnostic evaluation settings, this study proposes an imbalance-aware and computationally feasible framework to distinguish Minecraft gameplay from eyes-open rest using low-density pediatric EEG. *Methods:* We introduce Enriched Topological Features, which combine Takens phase-space embeddings with persistent homology, persistence landscapes, and persistence entropy. Evaluation was performed on four-channel Muse EEG from 43 children using subject-disjoint five-fold GroupKFold cross-validation. Class imbalance was handled exclusively within training folds using synthetic minority oversampling or within-subject random undersampling, while test folds preserved natural class proportions. Robustness to persistence landscape resolution and truncation depth was assessed using linear mixed-effects modeling. *Results:* Under oversampling-based training, the proposed framework achieved test balanced accuracy up to 65.39% and macro-averaged F1 score up to 64.48%, with minority-class rest recall reaching 49.51% using temporal electrodes. Feature attribution analysis associated gameplay predictions with frontal gamma-band topological structure, whereas rest predictions were associated with alpha-band organization at temporal sites. End-to-end processing required 0.676 s per 4.5 s epoch (0.15× real time), supporting near–real-time feasibility. *Conclusion:* The proposed approach enables imbalance-aware gameplay–rest classification from low-density pediatric mobile EEG and is compatible with near–real-time analysis. *Significance:* This work demonstrates that topological representations support robust cognitive-state monitoring in pediatric mobile neuroengineering under realistic motion and hardware constraints.**

***Index Terms*— Cognitive state classification, Mobile brain–body imaging, Pediatric electroencephalography, Persistent homology, Topological data analysis**

## I. Introduction

Understanding the neural basis of children's behavior in real-world settings remains a significant challenge in developmental neuroscience. Young children acquire cognitive, language, and social skills through dynamic, natural interactions, yet most neurophysiological studies still occur under solitary and highly controlled laboratory conditions [1]. Decoding pediatric electroencephalography (EEG) during naturalistic activities such as free play or gameplay with unconstrained movement offers more ecologically valid insights into brain function that are directly relevant to children's everyday experiences [1].

Advances in wearable EEG technology have made it possible to record brain activity outside of traditional laboratory environments [2]. This includes contexts where conventional participation is difficult, such as children with autism who may find lab visits stressful or families in remote communities. These developments reflect a growing effort to bridge the gap between laboratory-based protocols and real-world pediatric neuroscience.

Conventional EEG paradigms, however, often fall short of this goal. Most developmental EEG studies restrict child movement and minimize environmental complexity, typically instructing children to remain seated and view repetitive screen-based stimuli [1]. Adding to the challenge, children, especially infants and toddlers, are naturally active. Studies show that even during "calm" screen-based tasks, children move during the majority of recording time [3]. Such movement introduces artifacts that are difficult to suppress without either discarding large portions of data or removing meaningful neural signals [3]. These limitations highlight the need for alternative methods capable of tolerating motion while preserving signal interpretability.

To address these challenges, Mobile Brain–Body Imaging (MoBI) and mobile EEG systems have emerged as effective tools for studying brain activity during real-world behavior [4]. MoBI combines EEG with motion capture or inertial sensors, enabling participants to move freely while cognitive processes are recorded in context [5]. This approach treats movement not as noise but as meaningful cognitive data. However, ecologically valid pediatric EEG datasets acquired from freely behaving children in dynamic, non-clinical environments remain rare, limiting systematic evaluation of decoding methods under realistic motion and noise conditions.

Rather than seeking to maximize classification accuracy under favorable or imbalance-agnostic conditions, this study evaluates which EEG representations remain informative when subject-wise separation, severe rest–task imbalance, and realistic motion constraints are explicitly enforced. EEG recorded from children engaged in naturalistic video game play is analyzed to determine whether task-related neural signatures can be recovered even under high-movement conditions using low-density mobile recordings. Together, these analyses demonstrate the feasibility of studying pediatric brain dynamics *in situ* and align with the broader transition toward naturalistic developmental neuroscience.

**Contributions.** This work makes four primary contributions. First, we present a rigorously imbalance-aware and subject-disjoint evaluation of cognitive state decoding from pediatric mobile EEG, demonstrating how accuracy-centric reporting can substantially overestimate performance under severe rest–task imbalance. Second, we introduce Enriched Topological Features (ETF), which augment persistence landscape representations with entropy-based descriptors and temporal aggregation, yielding a compact and interpretable characterization of EEG dynamics that remains informative under low-density, high-motion pediatric recordings. Third, we systematically assess the robustness of topological feature parameterization using linear mixed-effects modeling, showing that ETF performance is stable across persistence landscape resolution and truncation depth, thereby reducing sensitivity to heuristic tuning. Finally, we demonstrate that ETF supports interpretable, imbalance-aware gameplay–rest

*Corresponding author: W. B. Owais (100063776@ku.ac.ae).
Manuscript submitted December XX, 2026.
[1]Department of Computer and Information Engineering and the Center of Secure Cyber Physical Systems, Khalifa University, Abu Dhabi, UAE (e-mail: ibrahim.elfadel@ku.ac.ae).
[2]Department of Medical Sciences, Khalifa University, Abu Dhabi, UAE (e-mail: herbert.jelinek@ku.ac.ae).

discrimination with near-real-time computational efficiency using as few as two EEG channels, supporting practical deployment in naturalistic pediatric settings.

## II. Related Work

Conventional EEG analysis approaches for cognitive state classification commonly rely on spectral power measures, frequency decomposition, and functional connectivity metrics. While these techniques are well established, they typically assume linear and stationary signal behavior, assumptions that are frequently violated in naturalistic and task-evoked EEG recordings [6]. These limitations are particularly pronounced in pediatric EEG, where motion artifacts, high inter-subject variability, and limited durations of artifact-free data further degrade signal reliability and model generalization.

Neural network–based models, including convolutional architectures, have been applied to EEG to learn task-related representations directly from the data. However, their deployment in MoBI settings is constrained by practical factors such as low-channel EEG systems, small cohort sizes, and substantial recording noise [7]. Studies attempting to distinguish resting from task states under similar mobile and pediatric conditions have reported only moderate performance, reflecting the difficulty of extracting stable neural signatures under these constraints [7].

EEG microstate analysis has been proposed as an intermediate-scale representation of brain dynamics, modeling EEG as sequences of quasi-stable spatial patterns. Microstate temporal statistics and transition structures have been investigated as biomarkers across several neurological and psychiatric conditions, including schizophrenia, tinnitus, trigeminal neuralgia, and major depressive disorder [8]–[11]. Although microstate-based approaches capture temporal organization beyond static spectral summaries, they discretize neural activity into a small set of states and characterize dynamics primarily through state occupancy and transitions. This limits their ability to represent continuous, multiscale temporal structure, particularly in noisy, low-density mobile EEG recordings. However, reliable microstate estimation typically requires higher spatial sampling than is available in low-density mobile EEG systems, limiting applicability in the present setting.

Beyond discrete-state representations, nonlinear time-series analysis methods aim to characterize EEG dynamics through measures such as entropy, fractal scaling, and phase-space reconstruction. These approaches move beyond pairwise associations and offer insight into higher-order temporal organization, but often exhibit sensitivity to parameter selection and limited robustness across subjects and recording conditions [12].

Recent advances in topological signal processing (TSP) have introduced noise-tolerant frameworks for characterizing the geometric structure of neural dynamics [13]. By applying persistent homology to phase-space embeddings of EEG signals, TSP captures invariant topological features that reflect global dynamical organization rather than localized signal properties. This perspective has demonstrated utility in applications such as neurodegenerative disease classification and shows promise for decoding cognitive states in complex, naturalistic scenarios [14].

A notable application of this paradigm is the nonlinear EEG topological dynamics analysis (NETDA) framework [15], which applied topological data analysis to pediatric EEG recorded during real-world gameplay. NETDA reported high accuracy under accuracy-centric evaluation by extracting persistence-based features from reconstructed phase spaces. Its effectiveness stems from the noise tolerance of topological representations and their suitability for low-channel EEG. However, the framework remains sensitive to embedding parameter choices, provides limited interpretability of extracted features, and does not explicitly address class imbalance or deployment constraints. These limitations motivate further refinement of topological approaches for robust and interpretable cognitive state decoding in mobile pediatric EEG.

Despite these advances, several gaps remain in the application of topological methods to pediatric mobile EEG. Prior work has largely emphasized overall-accuracy–focused evaluation in the presence of severe class imbalance, often without reporting balanced accuracy or minority-class sensitivity. In addition, existing frameworks provide limited analysis of parameter robustness, offer minimal insight into the interpretability of extracted topological features, and less frequently report computational feasibility for online or near–real-time use.

Accordingly, there is a need for evaluation protocols that explicitly account for subject-wise separation, class imbalance, and deployment constraints. The goal of this work is to identify EEG representations that remain informative under subject-disjoint testing and severe rest–task imbalance using imbalance-aware metrics as primary outcomes. To this end, we introduce Enriched Topological Features (ETF), which extend persistence-based representations through entropy augmentation and systematic sensitivity analysis. ETF is evaluated using train-only balancing and subject-disjoint cross-validation, with results reported using class-balanced metrics alongside accuracy-centric metrics for baseline comparability.

## III. Topological Foundations of Feature Extraction

### A. Vietoris–Rips Complexes and Simplicial Filtration

Topological Data Analysis (TDA) offers a formalism for capturing the intrinsic structure of data through topological invariants. Central to this framework is the notion of a *simplicial complex*, a combinatorial structure that generalizes graphs to higher dimensions. A $k$-simplex consists of $k+1$ affinely independent points; examples include 0-simplices (vertices), 1-simplices (edges), 2-simplices (triangles), and so on [16].

For a finite point cloud $X \subset \mathbb{R}^d$, the Vietoris–Rips complex $\mathcal{R}(X, \epsilon)$ at scale $\epsilon > 0$ is defined as the set of simplices $\sigma \subset X$ for which all pairwise distances between points in $\sigma$ are at most $\epsilon$:

$$\mathcal{R}(X, \epsilon) = \{\sigma \subset X \mid \mathrm{diam}(\sigma) \leq \epsilon\}. \tag{1}$$

As $\epsilon$ varies, these complexes form a filtration that encodes the data's shape across multiple spatial resolutions.

### B. Persistent Homology and Multiscale Topology

The evolution of topological features across a filtration is quantified using persistent homology. Each topological structure, such as a connected component or loop, is tracked by its birth and death scales, forming a homology class $\mathbb{H}_k$, where $\mathbb{H}_0$ corresponds to connected components and $\mathbb{H}_1$ denotes one-dimensional cycles [17].

Each feature is represented as a birth–death pair $(b_i^k, d_i^k)$, with persistence (or lifetime) defined as

$$\ell_i^k = d_i^k - b_i^k. \tag{2}$$

The collection of all such pairs forms a persistence diagram, which summarizes the multiscale lifespans of topological features.

In this work, analysis is restricted to the first two homology dimensions, $\mathbb{H}_0$ and $\mathbb{H}_1$, which capture coarse connectivity and loop-like structure in reconstructed EEG phase-space trajectories.

### C. Topological Descriptors: Landscapes and Entropy

While persistence diagrams are extensive in information, they are not directly amenable to vector-based analysis. To overcome this, persistence landscapes transform diagrams into functions in a Hilbert space, enabling consistent and differentiable representations [18]. These are constructed by projecting birth–death intervals onto piecewise-linear functions, retaining the most significant layers for further analysis.

In addition, persistence entropy [19] serves as a compact summary of diagram complexity. For a persistence diagram $D = \{(b_i, d_i)\}$, the persistence entropy is given by:

$$\mathrm{PE}(D) = -\sum_i p_i \log p_i, \tag{3}$$

$$p_i = \frac{d_i - b_i}{\sum_j (d_j - b_j)}. \tag{4}$$

This measure reflects the spread and uniformity of topological lifespans within a given homology dimension and provides a scalar descriptor of structural regularity. Persistence landscapes preserve multiscale geometric structure, while persistence entropy provides a compact summary of diagram complexity; together, they offer complementary representations of topological organization.

## IV. Dataset Description

This study uses the publicly available Mobile Brain–Body Imaging (MoBI) dataset collected at the Children's Museum of Houston [20]. The dataset comprises electroencephalogram (EEG) and inertial sensor recordings from 232 children (ages 6 to 16 years; mean age $8.7 \pm 2.2$) engaged in naturalistic, museum-based gameplay. Each participant completed a 1-minute eyes-open resting baseline followed by up to 20 minutes of free-play Minecraft gameplay in a real-world environment.

EEG signals were recorded using a 4-channel Muse headband (TP9, TP10, AF7, AF8) at a sampling rate of 220 Hz with an online 60 Hz notch filter. Concurrently, 3-axis head motion was captured via onboard accelerometers at 50 Hz. The dataset includes raw EEG signals, head movement recordings, and annotated event markers enabling segmentation of rest and gameplay periods. After quality control and requiring valid epochs from both rest and gameplay conditions, 43 subjects were retained for downstream analysis.

All EEG and accelerometer recordings were obtained directly from the publicly released `.mat` files provided by the dataset authors. No proprietary Muse preprocessing software (e.g., Muse SDK or MuseDirect) was used. All preprocessing, epoch segmentation, and artifact rejection were performed using a custom MATLAB pipeline (Section V-A).

### A. Ethical Approval

This work constitutes a secondary analysis of a publicly available, fully de-identified pediatric EEG dataset collected as part of the Mobile Brain–Body Imaging (MoBI) project at the Children's Museum of Houston. The original data collection protocol was approved by the Institutional Review Board (IRB) at the University of Houston. Written informed consent was obtained from parents or legal guardians, and informed assent was obtained from all participating children, as reported in the original dataset publications [20]. The dataset is openly available through IEEE DataPort (DOI: 10.21227/H23W88) and contains no personally identifiable information. No new data were collected, and no interaction with human participants occurred in the present study.

## V. Methodology: Feature Extraction and Classification with Enriched Topological Features (ETF)

The Enriched Topological Features (ETF) framework is designed to support decoding of cognitive brain states from low-density pediatric EEG recorded in naturalistic environments. It combines topological signal processing with entropy-based temporal summarization to capture nonlinear patterns in brain activity while remaining stable under real-world noise and variability. To handle class imbalance and subject heterogeneity, ETF uses both random undersampling and synthetic oversampling methods. The pipeline is modular and interpretable, and was designed with computational efficiency in mind for mobile neuroimaging and time-sensitive applications. The method is summarized in Algorithm 1.

### A. Preprocessing and Signal Transformation

EEG data were acquired from four Muse headband channels (TP9, TP10, AF7, AF8) at a sampling rate of 220 Hz. Recordings were segmented into rest and gameplay intervals using manually validated event markers embedded in the recording metadata. To enhance signal fidelity and suppress non-neural artifacts, a structured preprocessing pipeline was applied.

Raw EEG signals were first high-pass filtered at 1 Hz using a zero-phase 4th-order Butterworth filter to remove DC offset and low-frequency drift. A zero-phase notch filter centered at 60 Hz was then applied to attenuate line noise contamination [21], followed by band-specific zero-phase FIR filtering corresponding to the analyzed frequency bands. Signals were demeaned per channel after filtering, and all subsequent phase-space reconstruction and topological analyses were performed on these band-limited, demeaned signals.

Each epoch was subjected to a multi-stage quality control procedure. Epochs exhibiting excessive amplitude fluctuations, defined as a standard deviation exceeding 75 $\mu$V on any channel, were rejected. Epochs with abnormally low variance (standard deviation below 0.05 $\mu$V), indicative of flatlined signals or sensor detachment, were also excluded. To mitigate electromyographic (EMG) contamination, power spectral density (PSD) was estimated using Welch's method [22] on signals after high-pass (1 Hz) and notch (60 Hz) filtering but prior to band-specific FIR filtering. Epochs in which high-frequency power in the 35 to 45 Hz range exceeded 25% of the total spectral power integrated over 1 to 45 Hz on any channel were discarded. The QC high-frequency band (35 to 45 Hz) was chosen above the analyzed gamma range (30 to 35 Hz) to detect high-frequency contamination without discarding physiologically relevant gamma-band content. Only epochs that satisfied all quality criteria across all channels were retained.

To ensure subject-level comparability, only subjects containing valid epochs from both rest and gameplay conditions were retained. In addition, to prevent disproportionate influence from subjects with extended gameplay recordings, the number of gameplay epochs retained per subject was limited relative to the corresponding rest duration. This constraint reduced extreme subject-level dominance while preserving the natural imbalance

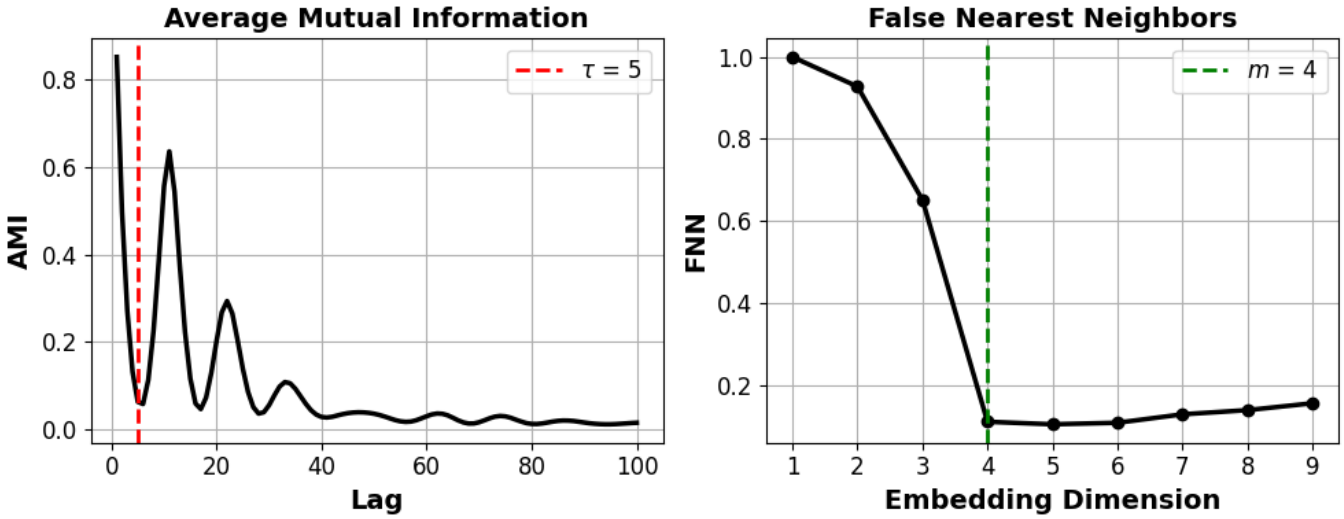


Fig. 1: Estimation of phase-space embedding parameters using Average Mutual Information (AMI) and False Nearest Neighbors (FNN) analysis for a representative subject (Subject 29). (Left) AMI as a function of time delay, with the first local minimum at $\tau = 5$. (Right) FNN proportion as a function of embedding dimension, showing a sharp decrease at $m = 4$. These trends are representative of those observed across subjects, channels, frequency bands, and behavioral conditions and motivated the fixed embedding parameters used in this study.

between conditions and was applied prior to any class balancing or classification procedures. Cleaned epochs were retained in both reduced (TP9 and TP10) and full (TP9, TP10, AF7, and AF8) channel configurations for downstream feature extraction and analysis.

After preprocessing and quality control, 43 subjects containing valid epochs from both rest and gameplay conditions were retained for downstream analysis. Subjects were excluded if required segmentation markers (baseline_start, baseline_end, task_start, task_end) were missing or if an insufficient number of valid 4.5-second epochs were available for either condition. From this final cohort, a total of 3,788 epochs were extracted, comprising 369 rest epochs and 3,419 gameplay epochs, corresponding to an approximate class imbalance of 9.3:1. This imbalance reflects the experimental design rather than sampling bias, as resting-state EEG was limited in duration whereas gameplay sessions were substantially longer.

Each band-limited EEG signal was transformed via phase-space reconstruction (PSR) using Takens' embedding theorem [23] to recover the underlying state-space dynamics. To determine suitable embedding parameters, Average Mutual Information (AMI) [24] and False Nearest Neighbors (FNN) [25] analyses were conducted across the full subject cohort using band-limited 4.5 s EEG epochs. AMI was used to estimate the time delay ($\tau$) by identifying the first local minimum of the mutual information curve, while FNN identified the embedding dimension ($m$) as the smallest dimension for which the fraction of false neighbors fell below 10%. These analyses yielded highly consistent results across subjects, channels, frequency bands, and behavioral conditions, with delay estimates tightly clustered around $\tau \approx 5$ to 6 samples and embedding dimensions consistently centered at $m = 4$. Based on these cohort-level trends, fixed values of $\tau = 5$ and $m = 4$ were adopted for all recordings, providing a robust and computationally efficient embedding choice while remaining within the dominant parameter optima observed across the dataset. This yielded embedded trajectories of the form:

$$\Phi(x_t) = [x_t, x_{t+\tau}, x_{t+2\tau}, x_{t+3\tau}] \in \mathbb{R}^4, \tag{5}$$

which preserves the geometric structure of temporal evolution and serves as input to the topological analysis pipeline described in the following section. Representative AMI and FNN curves for a single subject are shown in Fig. 1 to illustrate the typical trends observed across the dataset.

### B. Topological Feature Extraction

Topological data analysis (TDA) was employed to characterize the geometric structure of EEG dynamics. For each phase-space-reconstructed (PSR) trajectory, persistent homology was computed up to dimension one using Vietoris–Rips complexes, yielding persistence diagrams for connected components ($\mathbb{H}_0$) and one-dimensional cycles ($\mathbb{H}_1$).

Two complementary topological descriptors were extracted from each persistence diagram:

- **Persistence Landscapes (PL):** Birth–death intervals from $\mathbb{H}_0$ and $\mathbb{H}_1$ were transformed into persistence landscape functions, as described in Section III-C. Each landscape was sampled at PL = 30 uniformly spaced evaluation points to obtain fixed-length feature vectors.
- **Persistence Entropy (PE):** Each persistence diagram was summarized using persistence entropy, as defined in Section III-C (Eq. (3)–(4)), yielding a single scalar that quantifies the distribution of topological lifespans within each homology dimension.

Each EEG epoch was represented by a fixed-length feature vector composed of persistence landscapes (30 values) and persistence entropy (1 value) computed across four frequency bands (theta, alpha, beta, gamma), two homology dimensions ($\mathbb{H}_0$, $\mathbb{H}_1$), and either two channels (TP9, TP10) or all four Muse channels (TP9, TP10, AF7, AF8). This resulted in feature dimensionalities of 496 and 992 for the two-channel and four-channel configurations, respectively. Features were computed within sliding

**Algorithm 1** *Feature extraction and classification with Enriched Topological Features (ETF)*

**Require:** Raw EEG $\{x_i^{(c)}(t)\}$ for subjects $i = 1..N$, channels $c \in \{TP9, TP10, AF7, AF8\}$, sampling rate $f_s = 220$ Hz, task markers (baseline/task start/end), embedding ($m = 4, \tau = 5$), epoch length $T_e = 4.5$ s, window length $w = 330$, step $s = 110$, bands $\mathcal{B} = \{\theta, \alpha, \beta, \gamma\}$ with $\theta : 4-8$, $\alpha : 8-13$, $\beta : 13-30$, $\gamma : 30-35$ Hz, max homology $k_{\max} = 1$, persistence landscape resolution PL = 30, truncation depth $K = 1$.
**Ensure:** Feature matrix $\mathbf{U}$, labels $\mathbf{y}$, subject groups $\mathcal{G}$

**(A) Subject segmentation and artifact-validated epoch selection**
**for** each subject $i = 1$ to $N$ **do**
  Extract rest/gameplay intervals using annotated markers; set group $g_i \leftarrow i$
  Segment each interval into non-overlapping epochs of length $T_e$
  **for** each epoch $e$ **do**
    `fail` ← `False`
    **for** each channel $c$ **do**
      Preprocess $x_{i,e}^{(c)}$: HPF 1 Hz, notch 60 Hz, then band-specific FIR filtering
      Flag epoch $e$ if any QC criterion fails on channel $c$ (amplitude SD threshold, flatline detection, line-noise check, high-frequency contamination proxy; see Section V-A)
      **if** QC fails on channel $c$ **then**
        `fail` ← `True`
      **end if**
    **end for**
    **if** `fail` = `True` **then**
      Discard epoch $e$
    **else**
      Retain epoch $e$
    **end if**
  **end for**
**end for**
**(B) Topological feature extraction (per retained epoch)**
**for** each retained epoch $e$ from subject $i$ **do**
  Initialize epoch feature vector $\mathbf{u}_{i,e} \leftarrow [\ ]$
  **for** each channel $c$ (or subset, e.g., TP9/TP10) **do**
    **for** each band $b \in \mathcal{B}$ **do**
      Bandpass filter epoch to obtain $x_{i,e}^{(c,b)}$
      **for** each sliding window $\mathbf{w}$ in $x_{i,e}^{(c,b)}$ (length $w$, step $s$) **do**
        Takens embedding $\Phi(\mathbf{w}) = \{[x_t, x_{t+\tau}, x_{t+2\tau}, x_{t+3\tau}]\}_t \subset \mathbb{R}^m$
        Compute persistence diagrams $D_0, D_1$ (Vietoris–Rips, $k_{\max} = 1$)
        Compute persistence landscapes $\lambda_0^{(c,b)}(\mathbf{w}), \lambda_1^{(c,b)}(\mathbf{w})$ sampled at PL points (retain first layer, $K = 1$)
        Compute persistence entropies $E_0^{(c,b)}(\mathbf{w}), E_1^{(c,b)}(\mathbf{w})$
      **end for**
      Temporal aggregation within epoch: $\bar{\lambda}_0^{(c,b)} = \text{mean}_{\mathbf{w}} \lambda_0^{(c,b)}(\mathbf{w})$, $\bar{\lambda}_1^{(c,b)} = \text{mean}_{\mathbf{w}} \lambda_1^{(c,b)}(\mathbf{w})$, $\bar{E}_0^{(c,b)} = \text{mean}_{\mathbf{w}} E_0^{(c,b)}(\mathbf{w})$, $\bar{E}_1^{(c,b)} = \text{mean}_{\mathbf{w}} E_1^{(c,b)}(\mathbf{w})$
      Append $[\bar{\lambda}_0^{(c,b)}, \bar{\lambda}_1^{(c,b)}, \bar{E}_0^{(c,b)}, \bar{E}_1^{(c,b)}]$ to $\mathbf{u}_{i,e}$
    **end for**
  **end for**
  Store $\mathbf{U} \leftarrow \mathbf{U} \cup \{\mathbf{u}_{i,e}\}$, label $y_{i,e} \in \{0 \text{ (rest)}, 1 \text{ (game)}\}$, and group $\mathcal{G} \leftarrow \mathcal{G} \cup \{g_i\}$
**end for**
Resulting epoch feature dimensionality: $d = 496$ (2-channel) or $d = 992$ (4-channel)
**(C) Subject-wise evaluation (train-only balancing)**
Create 5 folds using GroupKFold over groups $\mathcal{G}$
**for** fold $k = 1$ to 5 **do**
  Split into $\mathcal{D}_{\text{train}}^{(k)}$ and $\mathcal{D}_{\text{test}}^{(k)}$ (subject-disjoint)
  Apply class balancing to $\mathcal{D}_{\text{train}}^{(k)}$ only:
    Option 1: SMOTE (minority → 50–70% of majority) with class weights
    Option 2: Within-subject random undersampling of gameplay epochs in $\mathcal{D}_{\text{train}}^{(k)}$
  **for** each classifier $f \in \{\text{RF}, \text{XGB}, \text{SVM}, \text{DT}\}$ **do**
    Train $f$ on balanced $\mathcal{D}_{\text{train}}^{(k)}$
    If undersampling is used, tune a classifier-specific decision threshold $\theta_k$ on $\mathcal{D}_{\text{train}}^{(k)}$ to maximize balanced accuracy; otherwise use the default decision threshold
    Apply $\theta_k$ to obtain predictions on untouched $\mathcal{D}_{\text{test}}^{(k)}$; record metrics
  **end for**
**end for**

windows and averaged over time to produce a single vector per epoch, capturing both geometric structure and topological complexity across multiple spatiotemporal scales.

To account for nonstationary signal behavior, a sliding-window approach was applied within each EEG epoch using a window length of 330 samples and a step size of 110 samples. Topological features were extracted per window and temporally averaged across all windows. Representative topological signatures derived from Subject 27 are shown in Figure 2, illustrating differences in barcodes, persistence diagrams, and landscapes between gameplay and rest states.

### C. Class Balancing and Feature Integration

To mitigate the severe class imbalance between rest and gameplay epochs, two class-balancing strategies were implemented during training: (i) oversampling with SMOTE [26] combined with cost-sensitive learning, and (ii) within-subject random undersampling [27].

In the *oversampling configuration*, SMOTE was applied exclusively within each training fold to synthetically generate minority-class (rest) samples by interpolating between nearest neighbors in the feature space. The sampling ratio was selected to increase minority-class representation without fully equalizing class sizes. In addition, cost-sensitive learning was employed by assigning higher misclassification penalties to rest epochs via explicit class weights during classifier training. This discouraged trivial majority-class predictions and improved sensitivity to underrepresented rest states. SMOTE was performed in the ETF feature space rather than on raw EEG signals and was restricted to training folds only. Because synthetically generated samples may not correspond to physiologically realizable EEG time series, SMOTE results are interpreted as facilitating minority-class decision boundary learning rather than as evidence of additional recorded rest-state data. For SMOTE-based experiments, classifiers used the default decision threshold; no threshold optimization was performed to avoid introducing an additional degree of freedom alongside synthetic oversampling and cost-sensitive learning.

In the undersampling configuration, gameplay epochs were randomly subsampled without replacement within each subject and within each training fold to match the number of rest epochs for that subject. This strategy preserves within-subject balance and data authenticity, albeit at the cost of reduced gameplay sample diversity and statistical power.

In both configurations, each EEG epoch was represented by a unified topological feature vector composed of persistence landscapes and entropy values derived from $\mathbb{H}_0$ and $\mathbb{H}_1$ diagrams across four canonical EEG bands and selected EEG channels (TP9 and TP10 in the reduced configuration, and TP9, TP10, AF7, and AF8 in the full configuration). Persistence landscapes were vectorized and concatenated across bands, channels, and homology dimensions, while persistence entropy contributed one scalar per homology dimension. The final feature representation is defined as:

$$\mathbf{u} = [\boldsymbol{\lambda}_0, \boldsymbol{\lambda}_1, E_0, E_1] \in \mathbb{R}^d, \tag{6}$$

where $\boldsymbol{\lambda}_k, k \in \{0, 1\}$, denotes the vectorized persistence landscape functions and $E_k, k \in \{0, 1\}$, denotes the corresponding persistence entropy. Here, $k = 0$ and $k = 1$ correspond to connected components ($\mathbb{H}_0$) and one-dimensional cycles ($\mathbb{H}_1$), respectively.

### D. Classification and Evaluation

Four supervised classifiers were selected to evaluate the proposed topological features across a spectrum of model complexities and interpretability. XGBoost (XGB) was chosen due to its strong performance in imbalanced, high-dimensional data settings and proven utility in EEG classification [28]. Random Forest (RF)

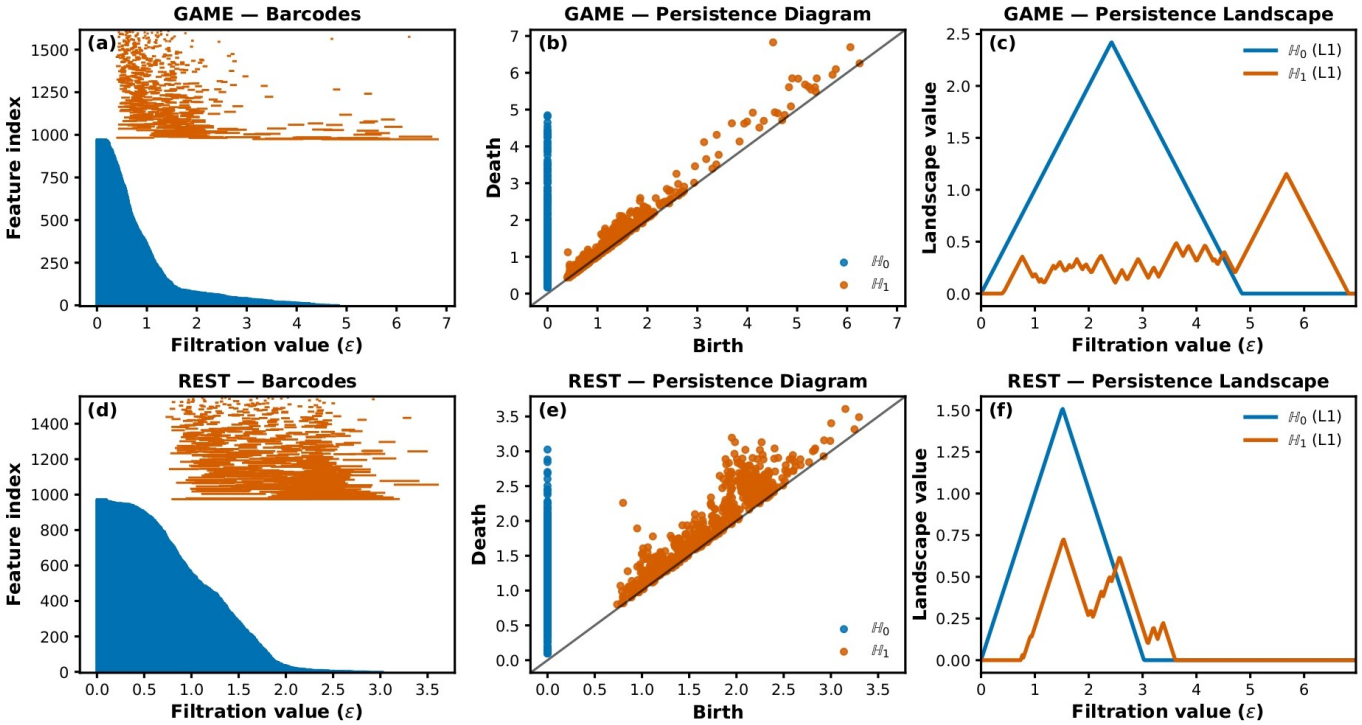


Fig. 2: Topological representations of EEG dynamics from **Subject 27** (alpha band, channel AF8) during **gameplay** (top row) and **rest** (bottom row). Panels (a,d) show persistence barcodes, (b,e) persistence diagrams, and (c,f) first-layer persistence landscapes. Barcodes and landscapes are plotted as functions of the filtration value ($\epsilon$). Persistence diagrams depict birth–death pairs for connected components ($\mathbb{H}_0$, blue) and loops ($\mathbb{H}_1$, orange), with the diagonal indicating equal birth and death.

offers resilience to noise and feature redundancy, making it well suited for topological EEG features [15]. Support Vector Machine (SVM) with an RBF kernel enables nonlinear decision boundaries, while Decision Tree (DT) provides an interpretable baseline with clear rule-based structure [29], [30].

To assess classification performance, 5-fold cross-validation was conducted using the GroupKFold strategy which ensures subject-disjoint evaluation by assigning all EEG segments from a given participant to the same fold. For each fold, class balancing was applied exclusively to the training split using one of two strategies: (i) SMOTE, or (ii) random undersampling of the majority class (gameplay epochs). Test folds were left untouched. All preprocessing and feature-extraction settings were fixed *a priori* and held constant across folds. For undersampling experiments, a classifier-specific decision threshold $\theta_k$ was optimized on the training fold to maximize balanced accuracy and then applied unchanged to the corresponding test fold.

In the SMOTE-based configuration, class weights were specified in each classifier to penalize misclassification of the minority class. Each EEG segment was represented by a feature vector of 496 or 992 dimensions, corresponding to the 2-channel and 4-channel configurations, respectively.

Model performance was quantified using accuracy, balanced accuracy, precision, recall, and F1 score. Metrics were computed independently for each class, with particular emphasis on balanced accuracy, macro-averaged F1, and minority-class (rest) recall, which are more informative under severe class imbalance. All metrics were reported separately for training and test sets, and mean values were computed across the five cross-validation folds to ensure reliable performance estimates. All experiments were repeated under two spatial configurations:

1) A reduced 2-channel setting using TP9 and TP10 (temporal electrodes), and
2) A full 4-channel setting using all Muse electrodes (TP9, TP10, AF7, AF8).

### E. Sensitivity Analysis of Topological Parameters

To assess the robustness of the proposed framework to topological feature parameterization, a linear mixed-effects (LME) analysis was conducted under the SMOTE-based training regime. The analysis evaluated the influence of persistence landscape resolution (PL) and the number of retained landscape layers ($K$) on classification performance, while accounting for repeated measurements across subject-wise cross-validation folds.

Persistence landscape resolution (PL) and truncation depth ($K$) were treated as robustness factors rather than tunable hyperparameters. This analysis was designed to evaluate sensitivity to parameterization and to inform the selection of fixed values for subsequent experiments.

### F. Implementation Details, Computational Complexity, and Runtime

EEG recordings were segmented into non-overlapping 4.5 s epochs (990 samples at 220 Hz). Within each epoch, sliding windows of 330 samples with a step size of 110 samples were extracted, yielding seven windows per epoch. Phase-space reconstruction (PSR) was performed using fixed parameters $(m, \tau) = (4, 5)$, producing embedded point clouds of 315 points in $\mathbb{R}^4$ per window. Persistent homology was computed up to dimension

$k = 1$ for each window, frequency band, and channel. Persistence landscapes and entropy values derived from $\mathbb{H}_0$ and $\mathbb{H}_1$ were then averaged across windows to obtain a single feature vector per epoch, resulting in feature dimensionalities of 496 and 992 for the two- and four-channel configurations, respectively.

The dominant computational cost arises from persistent homology computation on the embedded point clouds. For each retained epoch, persistent homology was evaluated 56 times in the two-channel configuration and 112 times in the four-channel configuration (four frequency bands × seven windows × number of channels, with homology computed up to dimension one). Although Vietoris–Rips persistence has exponential worst-case complexity with respect to the number of points, practical runtime remained tractable for point clouds of size 315 using optimized implementations.

All feature extraction and classification procedures were implemented in Python 3.12 using standard scientific computing libraries, including NumPy, SciPy, scikit-learn, and matplotlib. Topological features were computed using `ripser.py` and `giotto-tda`, while data balancing was handled using `imblearn`. Experiments were executed on a workstation equipped with an Intel Core i7-11800H CPU (8 cores, 2.30 GHz) and 32 GB DDR4 RAM.

Empirical runtime was evaluated by timing end-to-end feature extraction on randomly selected EEG epochs using the same configuration employed in the main experiments (four channels, PL = 30, $K = 1$), with parallel execution across available CPU cores. Averaged over 200 epochs, the ETF pipeline required 676 ms to process a single 4.5 s EEG epoch, corresponding to a real-time factor of 0.15×. This indicates that the computational throughput of the proposed method substantially exceeds real-time requirements, enabling processing to complete well within the duration of the analyzed signal. While evaluated in an offline per-epoch setting, these results suggest that ETF is computationally efficient and compatible with online or near–real-time mobile EEG analysis without GPU acceleration.

## VI. Results

### A. Sensitivity Analysis of Topological Parameters

Prior to reporting classification performance, sensitivity of the proposed Enriched Topological Feature (ETF) representation to persistence landscape resolution (PL) and truncation depth ($K$) was evaluated using linear mixed-effects (LME) models under SMOTE-based training. Across both two-channel and four-channel configurations, no significant main effect of PL was observed for test balanced accuracy, macro-F1, or overall accuracy, nor was any PL×$K$ interaction detected ($p > 0.1$ in all cases). Performance variation across the evaluated PL range (10 to 60) remained limited (mean spreads < 4 percentage points).

In contrast, $K$ exhibited a modest but statistically significant effect on imbalance-aware metrics, driven primarily by performance degradation at higher-order landscapes rather than consistent improvement. No monotonic gains were observed beyond $K = 1$. Accordingly, PL was fixed to 30 and $K$ to 1 for all subsequent experiments, providing sufficient numerical resolution with negligible additional computational overhead.

### B. Imbalance-Aware Evaluation of ETF

Given the severe class imbalance (rest:game ≈ 1:9.3) and subject-disjoint evaluation, overall accuracy can be dominated by majority-class behavior. Accordingly, balanced accuracy, macro-F1, and minority-class (rest) recall are treated as primary outcomes, while accuracy-centric metrics are reported separately for contextual comparison. All imbalance-aware results are reported as mean and 95% confidence intervals (CI) across five subject-wise GroupKFold folds.

Under subject-disjoint evaluation with severe class imbalance and unconstrained motion, these performance levels reflect recovery of task-relevant structure under conditions where chance-level balanced accuracy is 50%.

*1) SMOTE-Based Training:* Under SMOTE-based training, ETF yields balanced accuracies consistently above chance across classifiers and channel configurations. Tree-based models (Random Forest and XGBoost) outperform margin-based methods, reflecting robustness to synthetic sample variability and heterogeneous feature distributions. In the four-channel setting, XGBoost achieves a test balanced accuracy of 62.34% [95% CI: 55.31 to 69.37] and macro-F1 of 63.47% [55.58 to 71.36], with rest recall of 34.88% [24.09 to 45.68]. In contrast, the two-channel configuration attains comparable or higher balanced accuracy, with Random Forest achieving 65.39% [58.39 to 72.40] and substantially higher rest recall of 49.51% [43.76 to 55.26], indicating that temporally located electrodes capture sufficient task-relevant topological structure under imbalance-aware evaluation.

*2) Within-Subject Random Undersampling:* Relative to SMOTE, undersampling generally reduces macro-F1 and yields wider confidence intervals, consistent with the loss of majority-class diversity after truncation. While minority-class recall can increase for certain classifier and configuration combinations, this increase reflects classifier bias toward the minority class (e.g., in decision trees) at the expense of majority-class precision, resulting in reduced macro-F1 and balanced accuracy. Overall, these trade-offs support SMOTE as the more stable strategy for this dataset.

### C. Interpretation Relative to Prior Work

Accuracy-centric baselines report high test accuracy and strong majority-class (game) performance, but do not evaluate minority-class (rest) sensitivity and do not report imbalance-aware metrics. Consequently, lower overall accuracy under imbalance-aware evaluation should not be interpreted as inferior discrimination, but rather as a direct consequence of penalizing majority-class dominance to recover minority-class sensitivity. In contrast, ETF is evaluated using balanced accuracy, macro-F1, and rest recall (Tables I–II), which provide a more informative assessment of class-balanced discrimination under severe imbalance. Notably, ETF recovers rest-state sensitivity using as few as two temporal electrodes, indicating that task-relevant topological structure remains informative even under low-density mobile EEG constraints.

### D. Accuracy-Centric Metrics Reported in Prior Work

Because rest epochs constitute approximately 9.3% of the dataset, overall accuracy and majority-class F1 can exceed 85 to 90% even when rest recall is near zero. Table III reports performance using the same accuracy-centric metrics adopted in prior work, namely overall test accuracy and majority-class precision, recall, and F1 score. Because prior work reports majority-class metrics and accuracy without evaluating minority-class sensitivity, this table is included only for protocol comparability and should not be interpreted as an imbalance-aware benchmark.

TABLE I: ETF performance under SMOTE-based training (PL = 30, $K = 1$)

| Ch. | C. | Test BAcc | Test Macro-F1 | Rest Rec | Rest F1 | Game F1 |
|---|---|---|---|---|---|---|
| A4 | RF | 57.39 [45.98, 68.81] | 57.97 [45.80, 70.15] | 32.32 [20.74, 43.90] | 33.99 [16.79, 51.20] | 81.95 [74.22, 89.68] |
| | XGB | 62.34 [55.31, 69.37] | 63.47 [55.58, 71.36] | 34.88 [24.09, 45.68] | 40.37 [26.99, 53.74] | 86.58 [83.72, 89.44] |
| | SVM | 63.53 [52.48, 74.58] | 61.98 [47.82, 76.14] | 44.17 [26.89, 61.44] | 42.09 [15.41, 68.77] | 79.80 [68.16, 91.43] |
| | DT | 55.97 [51.04, 60.90] | 53.52 [48.79, 58.24] | 33.79 [25.66, 41.93] | 26.99 [19.33, 34.65] | 73.24 [71.03, 75.45] |
| T2 | RF | 65.39 [58.39, 72.40] | 64.48 [56.09, 72.86] | 49.51 [43.76, 55.26] | 45.71 [34.25, 57.18] | 83.24 [77.32, 89.17] |
| | XGB | 61.68 [57.54, 65.82] | 62.01 [56.46, 67.56] | 39.97 [32.76, 47.19] | 38.61 [34.51, 42.71] | 84.05 [78.87, 89.23] |
| | SVM | 67.53 [57.46, 77.61] | 64.20 [53.81, 74.59] | 48.19 [34.03, 62.35] | 40.24 [26.06, 54.42] | 80.21 [71.98, 88.43] |
| | DT | 57.28 [50.26, 64.29] | 53.94 [47.39, 60.48] | 35.68 [25.76, 45.61] | 27.81 [19.13, 36.48] | 72.19 [67.90, 76.48] |

Values are mean [95% confidence interval] across five subject-wise folds. BAcc denotes balanced accuracy. Ch. denotes EEG channel configuration (A4: TP9, TP10, AF7, AF8; T2: TP9, TP10). C. denotes classifier. Rest and Game metrics correspond to the minority and majority classes, respectively (rest is the minority class).

TABLE II: ETF performance under within-subject random undersampling applied to the training fold only (PL = 30, $K = 1$).

| Ch. | C. | Test BAcc | Test Macro-F1 | Rest Rec | Rest F1 | Game F1 |
|---|---|---|---|---|---|---|
| A4 | RF | 58.58 [43.54, 73.61] | 58.10 [42.27, 73.93] | 41.32 [25.18, 57.47] | 37.87 [18.94, 56.80] | 78.33 [64.51, 92.15] |
| | XGB | 60.04 [53.51, 66.57] | 60.77 [53.36, 68.18] | 31.81 [19.40, 44.21] | 36.25 [23.52, 48.97] | 85.30 [80.90, 89.70] |
| | SVM | 65.72 [54.47, 76.96] | 59.29 [48.21, 70.36] | 45.08 [29.96, 60.20] | 34.08 [20.13, 48.04] | 73.49 [65.95, 81.03] |
| | DT | 57.89 [51.07, 64.71] | 52.60 [45.16, 60.03] | 58.15 [47.43, 68.87] | 27.53 [18.46, 36.60] | 68.08 [61.78, 74.39] |
| T2 | RF | 62.71 [54.79, 70.63] | 62.24 [51.92, 72.57] | 44.36 [33.84, 54.87] | 42.22 [28.74, 55.70] | 82.26 [73.18, 91.34] |
| | XGB | 60.82 [56.97, 64.67] | 61.52 [56.80, 66.24] | 33.09 [22.21, 43.97] | 37.49 [29.18, 45.80] | 85.56 [81.17, 89.95] |
| | SVM | 64.54 [52.91, 76.17] | 60.92 [48.40, 73.43] | 59.00 [43.76, 74.24] | 37.06 [21.45, 52.68] | 77.04 [66.22, 87.87] |
| | DT | 57.94 [53.01, 62.88] | 52.37 [46.89, 57.85] | 58.43 [51.16, 65.70] | 27.29 [22.13, 32.44] | 67.82 [60.33, 75.30] |

Values are mean [95% confidence interval] across five subject-wise folds. BAcc denotes balanced accuracy. Ch. denotes EEG channel configuration (A4: TP9, TP10, AF7, AF8; T2: TP9, TP10). C. denotes classifier. Rest and Game metrics correspond to the minority and majority classes, respectively (rest is the minority class).

TABLE III: Accuracy-centric majority-class metrics reported in prior work (not imbalance-aware; not comparable to Tables I–II). A trivial majority-class baseline is included to illustrate the inflation of accuracy-centric metrics under severe class imbalance.

| Method | Channels | Classifier | Test Acc (%) | Prec (%) | Rec (%) | F1 (%) |
|---|---|---|---|---|---|---|
| Trivial majority-class baseline (always predict gameplay) | – | – | 90.7 | 90.7 | 100.0 | 95.1 |
| DL [7] | TP9,TP10 | – | 67.40 | 67.99 | 65.55 | 66.80 |
| | All 4 | – | 73.68 | 72.91 | 71.40 | 72.15 |
| NETDA [15] | TP9,TP10 | RF | 89.71 | 90.90 | 88.46 | 89.71 |
| | | SVM | 87.96 | 91.50 | 86.71 | 87.96 |
| | | DT | 83.08 | 84.27 | 81.33 | 83.08 |
| | All 4 | RF | 91.96 | 92.72 | 91.21 | 91.96 |
| | | SVM | 89.92 | 90.68 | 89.17 | 89.92 |
| | | DT | 88.22 | 88.98 | 87.47 | 88.22 |
| ETF (this work) | TP9,TP10 | RF | 74.48 | 85.52 | 81.28 | 83.24 |
| | | XGB | 74.95 | 83.56 | 84.75 | 84.05 |
| | | SVM | 71.61 | 87.58 | 74.20 | 80.21 |
| | | DT | 61.25 | 82.85 | 64.10 | 72.19 |
| | All 4 | RF | 71.77 | 81.84 | 82.47 | 81.95 |
| | | XGB | 78.12 | 83.66 | 89.80 | 86.58 |
| | | SVM | 70.61 | 84.97 | 76.17 | 79.80 |
| | | DT | 61.93 | 81.95 | 66.26 | 73.24 |

Only metrics reported in prior work are shown. All precision, recall, and F1 values correspond exclusively to the majority (gameplay) class. Minority-class (rest) and imbalance-aware metrics are intentionally excluded and reported separately in Table I.

## VII. Feature Analysis

To interpret the contribution of individual topological features to model decisions, SHAP (SHapley Additive exPlanations)

analysis [31] was performed using TreeExplainer on the trained XGBoost classifier. Features were derived from the ETF pipeline applied to each EEG epoch. SHAP values reflect feature contributions to the classifier's decision function and should not be interpreted as direct neurobiological measurements, source-localized effects, or clinical indicators.

SHAP values were computed on the resulting feature matrix to quantify the marginal contribution of each feature to the classifier output. Mean absolute SHAP values were used to rank feature importance across the dataset. For interpretability, features were labeled according to their frequency band, electrode location, homology dimension, and feature type (persistence landscape component or entropy).

### A. Key Topological Features

Figures 3(a) and 3(b) summarize the most influential features identified by SHAP analysis. Overall, gameplay classification was primarily attributed by the model to gamma-band persistence landscape features at frontal electrodes, whereas resting-state predictions were associated with alpha-band topological structure at temporal sites.

- **Gamma (AF8) $\mathbb{H}_1$ PL 2** exhibited the highest mean absolute SHAP value ($\approx$0.16), indicating a strong association with gameplay classification. This feature captures persistent loop-like structure in gamma-band phase-space dynamics at AF8.
- **Alpha (TP9) $\mathbb{H}_1$ PL 15 and PL 19**, along with **Gamma (TP9) $\mathbb{H}_1$ PL 6**, showed large negative SHAP values (mean range: –0.4 to –0.6), indicating stronger association with resting-state epochs characterized by sustained low-frequency topological organization.
- **Entropy-based features** displayed mixed SHAP polarity across bands and channels, suggesting that changes in signal complexity contributed to both gameplay and rest predictions in a state-dependent manner rather than exhibiting unidirectional effects.

### B. Neurocognitive Interpretation

The prominence of frontal gamma-band persistence landscapes, particularly at AF8, is consistent with patterns commonly reported during increased cognitive engagement in goal-directed tasks. Gamma activity has been associated in prior work with rapid cortical processing, working memory maintenance, and executive control mechanisms that are heavily recruited during visually complex and goal-directed tasks such as *Minecraft* [32], [33]. Importantly, the ETF framework captures not gamma power per se, but the persistence of loop-like structures in the reconstructed phase space, reflecting stable dynamical patterns rather than transient amplitude fluctuations.

In contrast, temporal alpha-band persistence landscapes, most notably at TP9, were strongly linked to resting-state predictions. Alpha rhythms are known to index sensory inhibition and internally oriented processing, and their sustained topological structure is consistent with stable default-mode network dynamics during cognitive disengagement [34].

Lateralized effects were also observed in the beta band. **Beta (AF7) $\mathbb{H}_1$ PL 10** contributed positively to gameplay classification, whereas homologous right-temporal features (e.g., TP10) showed weaker influence. This asymmetry is compatible with left-hemisphere dominance in visuomotor planning and top-down control, and aligns with prior evidence of beta–gamma interactions during motor coordination and task execution [35].

### C. Topological Signatures of Cognitive Load

Beyond individual features, SHAP analysis revealed systematic differences in topological organization between behavioral states. Reduced alpha-band entropy during gameplay suggests a relative suppression of spontaneous cortical variability, consistent with constrained neural dynamics under elevated executive demand [36]. Concurrently, increased persistence of gamma-band $\mathbb{H}_1$ landscapes at frontal electrodes reflects more stable task-associated dynamical organization in the reconstructed phase space.

While high frequency EEG components in mobile settings may be susceptible to myogenic contamination, the dominance of persistence based features rather than raw spectral power suggests that classification relies on the stability of geometric structure in phase space. Nonetheless, gamma related findings are interpreted conservatively as markers of task associated dynamical organization rather than exclusively neurogenic oscillatory activity.

### D. Clinical and Developmental Relevance

The observed gamma–alpha dissociation and hemispheric asymmetries are particularly relevant in the context of pediatric neurodevelopment [37]. Frontal gamma persistence during gameplay may reflect compensatory or immature control strategies in developing cortico-cortical networks engaged during top-down regulation. Beta-band lateralization further implicates frontoparietal and premotor pathways involved in visuomotor integration and early-stage motor learning. Together, these results indicate that ETF captures dynamical features that vary with developmental factors and are not fully described by conventional spectral descriptors. These findings are intended to support clinician-facing monitoring and decision support in naturalistic settings rather than diagnostic classification or clinical state labeling.

Interpretability claims in this work refer to model-level attribution over engineered topological features using SHAP-based analysis rather than direct mechanistic inference from raw EEG signals. Consequently, attribution patterns are presented as hypothesis-generating descriptions of state-dependent dynamical structure, not as definitive neurophysiological biomarkers.

## VIII. Discussion

### A. Summary of Contributions

This study introduces *Enriched Topological Features* (ETF) as a statistically validated and physiologically interpretable framework for decoding cognitive brain states from low-density pediatric EEG acquired in naturalistic settings. By integrating persistent homology, time-aggregated entropy, and phase-space embeddings, ETF supports reliable discrimination between gameplay and rest under severe class imbalance, subject-wise evaluation, and low spatial resolution.

Rather than optimizing accuracy alone, ETF explicitly targets imbalance-aware performance, demonstrating consistent recovery of minority-class (rest) sensitivity while maintaining stable majority-class discrimination across classifiers and channel configurations. These findings indicate that topological data analysis provides a principled means of extracting neural features that remain informative under motion, environmental noise, participant movement, and limited spatial sampling.

Beyond performance considerations, ETF supports feature-level interpretability by linking classification outcomes to structured dynamical properties of EEG signals rather than purely statistical feature abstractions. Importantly, the contribution of this work lies not in surpassing prior accuracy-centric benchmarks, but in

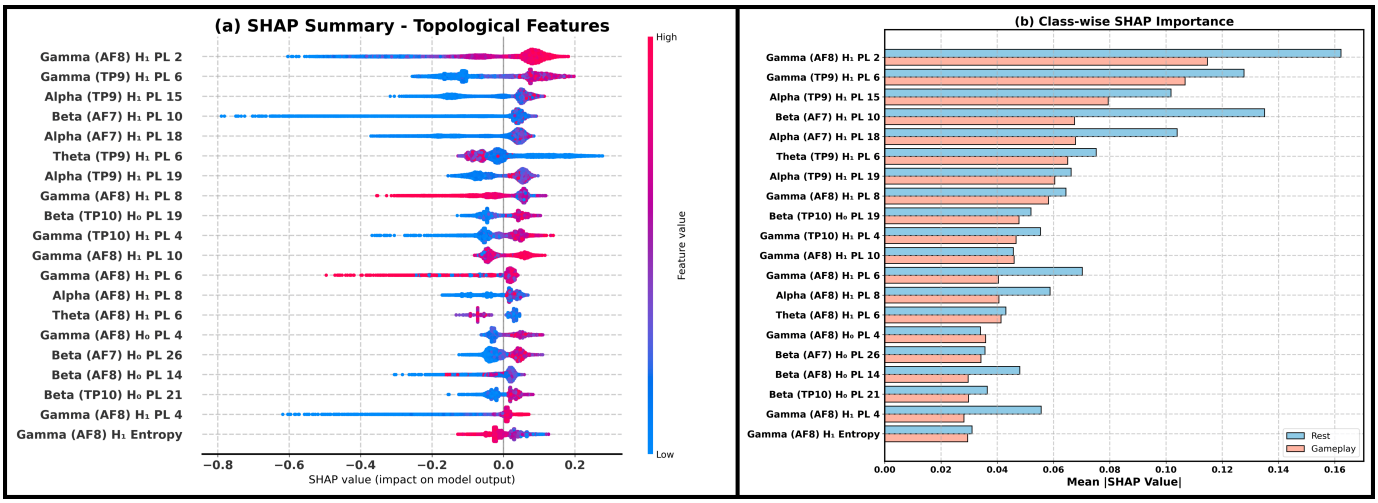


Fig. 3: Model interpretability via SHAP-based analysis of topological EEG features. (a) SHAP summary plot showing the top 20 features ranked by impact on classifier output. Each point represents a SHAP value for an instance, colored by the corresponding feature magnitude. (b) Class-wise SHAP importance plot showing mean absolute SHAP values for rest and gameplay conditions separately. Feature labels follow the format: *Band (Channel) Homology Feature Index*.

demonstrating which neural representations remain informative when subject-wise separation, severe class imbalance, and real-world recording constraints are explicitly enforced.

### B. Interpretability and Neurophysiological Context

Model attribution analysis using SHAP revealed structured topological patterns that differentiate gameplay engagement from rest. Gameplay epochs were consistently associated with increased contributions from gamma-band persistence landscape features at frontal electrodes, most prominently AF8 in $\mathbb{H}_1$. These features are consistent with enhanced dynamical organization during task engagement, consistent with prior associations between gamma-band activity and fast cortical processing, working memory, and executive control during goal-directed behavior [38]. Importantly, these observations represent model attribution patterns rather than direct neurobiological measurements.

In contrast, resting-state predictions were dominated by alpha-band entropy and persistence features at temporal electrodes, particularly TP9. This pattern is consistent with metastable disengagement and internally oriented processing, aligning with established roles of alpha rhythms in sensory inhibition and default-mode network dynamics [39].

Lateralized asymmetries were also observed across beta and gamma features, with stronger task-related contributions arising from left frontal regions compared to right-temporal sites. Such asymmetries are consistent with developmental neurophysiology, wherein frontoparietal and parietal–cerebellar pathways support visuomotor coordination and executive–motor integration during childhood [40]. Collectively, these attribution patterns suggest that ETF captures state-dependent reconfigurations of neural dynamics rather than incidental signal fluctuations or classifier artifacts.

### C. Developmental Trends and Effect Size Interpretation

To explore potential developmental associations, we examined the relationship between participant age and gamma-band topological features derived during gameplay. For each subject ($N = 43$), features were first averaged across all retained gameplay epochs to obtain a single subject-level estimate. The analysis focused on the mean persistence entropy of $\mathbb{H}_1$ loops in the frontal gamma band, which reflects the temporal stability of one-dimensional dynamical cycles.

A modest but statistically significant negative association was observed between age and frontal gamma-band $\mathbb{H}_1$ persistence (Pearson correlation, $r = -0.23$, $p = 0.0176$), indicating reduced persistence of high-frequency topological structure with increasing age. This analysis was conducted at the subject level and treated as exploratory; no correction for multiple comparisons across frequency bands was applied.

Although the effect size is small, correlations of this magnitude are typical in developmental EEG studies conducted under naturalistic conditions, particularly when using low-density mobile systems. Importantly, this association is not interpreted as a diagnostic marker, but rather as a developmental correlate reflecting gradual refinement of executive network organization.

From a neurodevelopmental perspective, this trend aligns with established models of synaptic pruning and large-scale network maturation, in which the brain transitions from dense, locally recurrent connectivity toward more selective and efficient circuitry across childhood and adolescence [37], [41]. Longitudinal EEG studies report age-related decreases in gamma-band activity that parallel reductions in synaptic density [41], while pediatric fMRI studies demonstrate a shift from locally redundant networks toward more distributed and integrated functional architectures [42]. Together, these processes may reduce the need for sustained high-frequency reverberatory activity, manifesting as shorter gamma-band topological persistence.

To assess whether the observed topological features were trivially driven by movement-related artifacts, we additionally computed correlations between gamma-band topological features and a motion proxy derived from high-frequency signal variance within each epoch. Correlations were computed at the epoch level and averaged across subjects. The resulting mean correlation was near zero (mean $r = 0.002$), indicating no systematic association between motion magnitude and topological persistence. This analysis serves as a control for motion-related confounding rather than as a hypothesis-driven statistical test. Finally, no significant differences in gamma-band persistence were observed across Minecraft gameplay frequency groups (one-way ANOVA,

$F = 0.66$, $p = 0.52$), suggesting that maturational factors, rather than task exposure, primarily influence the observed topological structure.

### D. Relation to EEG Microstate Dynamics

EEG microstate analysis models scalp activity as a sequence of short-lived, quasi-stable spatial configurations and is commonly used to summarize large-scale temporal organization in cognitive and developmental EEG [43]. Prior work reports that increased attentional demand is often accompanied by shifts in microstate expression and reduced dominance of default-mode–like patterns, whereas resting conditions show longer state lifetimes and greater default-mode occupancy [44].

We do not claim microstate validation in this work; the following comparison is interpretive only. Because microstates cannot be reliably estimated from the present four-channel montage, the microstate literature is used solely as a conceptual reference for interpreting state-dependent differences in temporal organization.

Consistent with the general contrast between task-driven and internally oriented dynamics, gameplay epochs were characterized by increased frontal gamma-band $\mathbb{H}_1$ persistence and reduced alpha-band structure at temporal sites, whereas rest epochs showed stronger alpha-band topological organization. These patterns qualitatively align with reports of task-related reductions in default-mode–like dynamics in microstate studies.

Overall, microstate and topological analyses emphasize complementary levels of description. Microstates summarize recurrent scalp-level spatial configurations when sufficient spatial sampling is available, whereas ETF characterizes the geometric organization of band-limited single-channel dynamics in a manner that remains informative under low-density mobile EEG. Accordingly, this comparison is intended as conceptual alignment rather than a microstate-based validation of ETF.

### E. Clinical Implications and Real-World Applicability

ETF is motivated by real-world pediatric settings where traditional high-density EEG is impractical due to cost, setup complexity, and poor tolerance to movement. Potential applications include monitoring cognitive engagement during serious gaming or neurorehabilitation and supporting adaptive task design under ecologically valid conditions. Importantly, ETF is intended as a decision-support and monitoring framework rather than a diagnostic tool, enabling coarse but robust differentiation of cognitive states from low-density mobile EEG.

From a computational standpoint, ETF demonstrates favorable efficiency for resource-constrained settings. End-to-end feature extraction required an average of 676 ms to process a single 4.5 s EEG epoch using four channels (PL = 30, $K = 1$) with parallel CPU execution, corresponding to a real-time factor of approximately 0.15×. Although runtime was evaluated offline on a per-epoch basis, throughput exceeded real-time requirements, indicating feasibility without GPU acceleration. These measurements reflect full per-epoch processing, including filtering, phase-space reconstruction, persistent homology, and feature aggregation, and support near–real-time operation on standard multi-core hardware.

From a translational perspective, ETF is best viewed as a clinician- or therapist-facing monitoring tool rather than an autonomous decision system. In practical scenarios, ETF-derived metrics could be computed continuously during neurorehabilitation or serious gaming sessions to quantify engagement-related neural organization. Such measures may inform adaptive task difficulty or session pacing, or support early identification of disengagement, particularly in pediatric populations where sustained attention is variable.

### F. Methodological Limitations

Several limitations should be acknowledged. First, the use of a four-channel EEG system constrains spatial resolution, limiting access to distributed cortical networks and precluding source-level analysis. While this design choice enhances ecological validity and scalability in pediatric settings, ETF should be viewed as complementary to high-density EEG rather than a replacement.

Second, although nested cross-validation or a held-out subject test set could be employed for strict hyperparameter optimization, the present study treats persistence landscape resolution and truncation depth as robustness parameters rather than tunable hyperparameters. The absence of performance sensitivity across the evaluated parameter range supports the stability of the proposed representation and reduces the risk of configuration-specific overfitting.

Third, the binary classification paradigm adopted here does not capture finer-grained cognitive states such as flow, frustration, or partial disengagement during gameplay. Extension to multi-class or ordinal state decoding represents an important direction for future work.

Finally, interpretation of high-frequency EEG features in mobile pediatric contexts warrants caution. Although correlations between EEG amplitude and head-motion energy were negligible across participants (mean $r = 0.002$, median $r = 0.001$), gamma-band activity remains susceptible to subtle myogenic contamination [45]. While reliance on persistence-based topological features may provide partial robustness to transient artifacts, definitive separation of neurogenic and myogenic sources will require concurrent EMG monitoring or multimodal artifact characterization. Although quality-control procedures and the use of persistence-based descriptors may reduce sensitivity to transient artifacts, high-frequency EEG components in mobile pediatric recordings remain susceptible to myogenic contamination. Therefore, gamma-band findings are interpreted conservatively as gamma-band–filtered topological dynamics rather than direct evidence of neurogenic gamma oscillations.

### G. Future Directions

Future work should prioritize large-scale validation studies with target sample sizes of approximately 150–250 participants to assess generalizability and to establish age-adjusted normative ranges for ETF-derived features. Multi-center collaborations employing harmonized MoBI-style acquisition protocols will be critical for evaluating robustness across recording environments, hardware platforms, and population heterogeneity. Methodologically, extensions to ordinal or multi-class decoding frameworks may enable characterization of intermediate cognitive states beyond binary rest and gameplay conditions, such as graded engagement or partial disengagement. Such extensions would provide a more nuanced representation of cognitive dynamics in naturalistic settings. Finally, integrating ETF with complementary EEG analysis paradigms, including microstate analysis or connectivity-based representations, may support a multilevel description of pediatric brain dynamics. Linking transient scalp-level patterns to underlying phase-space topology could offer deeper insight into the temporal organization of neural activity during development.

## IX. Conclusion

This work introduces Enriched Topological Features (ETF) as a computationally efficient and interpretable framework for decoding pediatric EEG recorded in naturalistic settings. By integrating persistent homology, time-aggregated entropy, and phase-space embeddings, ETF captures structured, state-dependent neural dynamics that remain informative under low spatial resolution, substantial motion, and severe class imbalance.

Rather than optimizing accuracy alone, ETF emphasizes imbalance-aware evaluation and feature-level interpretability, enabling reliable recovery of minority-class (rest) sensitivity while maintaining stable performance across classifiers and channel configurations. These properties address key challenges in mobile neuroengineering, where robustness, transparency, and ecological validity are often prioritized over raw accuracy.

Taken together, the proposed framework provides a principled foundation for pediatric cognitive-state monitoring and longitudinal developmental research using low-density mobile EEG, with potential utility in real-world educational and neurorehabilitation contexts.

## References


[1] E. S. Norton *et al.*, "Social EEG: A novel neurodevelopmental approach to studying brain-behavior links and brain-to-brain synchrony during naturalistic toddler–parent interactions," *Dev. Psychobiol.*, vol. 64, no. 3, p. e22240, 2022.

[2] A. Giannadou *et al.*, "Investigating neural dynamics in autism spectrum conditions outside of the laboratory using mobile electroencephalography," *Psychophysiology*, vol. 59, no. 4, p. e13995, 2022.

[3] S. Georgieva *et al.*, "Toward the understanding of topographical and spectral signatures of infant movement artifacts in naturalistic EEG," *Front. Neurosci.*, vol. 14, p. 352, 2020.

[4] S. Makeig *et al.*, "Linking brain, mind and behavior," *Proc. IEEE*, vol. 100, no. Special Centennial Issue, pp. 528–541, 2009.

[5] K. Gramann *et al.*, "Cognition in action: Imaging brain/body dynamics in mobile humans," *Rev. Neurosci.*, vol. 25, no. 3, pp. 365–380, 2014.

[6] S. Ladouce *et al.*, "Understanding minds in real-world environments: Toward a mobile cognition approach," *Front. Hum. Neurosci.*, vol. 10, p. 694, 2017.

[7] A. S. Ravindran *et al.*, "Assaying neural activity of children during video game play in public spaces: A deep learning approach," *J. Neural Eng.*, vol. 16, no. 3, p. 036028, 2019.

[8] Z. Raeisi *et al.*, "EEG microstate biomarkers for schizophrenia: A novel approach using deep neural networks," *Cogn. Neurodyn.*, vol. 19, no. 1, pp. 1–26, 2025.

[9] Z. Raeisi *et al.*, "Enhanced classification of tinnitus patients using EEG microstates and deep learning techniques," *Sci. Rep.*, vol. 15, no. 1, p. 15959, 2025.

[10] R. A. Lashaki *et al.*, "EEG microstate analysis in trigeminal neuralgia: Identifying potential biomarkers for enhanced diagnostic accuracy," *Acta Neurol. Belg.*, pp. 1–21, 2025.

[11] M. Rahmati *et al.*, "Novel EEG-based diagnostic framework for major depressive disorder using microstate and entropy features," *Cogn. Neurodyn.*, vol. 19, no. 1, p. 116, 2025.

[12] M. Saggar *et al.*, "Towards a new approach to reveal dynamical organization of the brain using topological data analysis," *Nat. Commun.*, vol. 9, no. 1, p. 1399, 2018.

[13] X. Xu *et al.*, "Topological data analysis as a new tool for EEG processing," *Front. Neurosci.*, vol. 15, p. 761703, 2021.

[14] N. Aithal *et al.*, "Leveraging persistent homology for differential diagnosis of mild cognitive impairment," in *Proc. Int. Conf. Pattern Recognit.*, 2025, pp. 17–32.

[15] K. Gaurav *et al.*, "Characterizing neural activity during video game engagement using EEG sensor based topological dynamics analysis," *IEEE Sens. Lett.*, 2024.

[16] A. Zomorodian and G. Carlsson, "Computing persistent homology," in *Proc. 20th Annu. Symp. Comput. Geom.*, 2004, pp. 347–356.

[17] H. Edelsbrunner *et al.*, "Topological persistence and simplification," *Discrete Comput. Geom.*, vol. 28, pp. 511–533, 2002.

[18] P. Bubenik *et al.*, "Statistical topological data analysis using persistence landscapes," *J. Mach. Learn. Res.*, vol. 16, no. 1, pp. 77–102, 2015.

[19] N. Atienza *et al.*, "Persistent entropy for separating topological features from noise in Vietoris-Rips complexes," *J. Intell. Inf. Syst.*, vol. 52, pp. 637–655, 2019.

[20] A. S. Ravindran *et al.*, "Multi-modal mobile brain-body imaging (MoBI) dataset for assaying neural and head movement responses associated with creative video game play in children," IEEE Dataport, 2017, [Online]. Available: doi: 10.21227/H23W88.

[21] A. Widmann *et al.*, "Digital filter design for electrophysiological data: A practical approach," *J. Neurosci. Methods*, vol. 250, pp. 34–46, 2015.

[22] P. Welch, "The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms," *IEEE Trans. Audio Electroacoust.*, vol. 15, no. 2, pp. 70–73, 1967.

[23] F. Takens, "Detecting strange attractors in turbulence," in *Dynamical Systems and Turbulence, Warwick 1980*, 2006, pp. 366–381.

[24] A. M. Fraser and H. L. Swinney, "Independent coordinates for strange attractors from mutual information," *Phys. Rev. A*, vol. 33, no. 2, pp. 1134–1140, 1986.

[25] M. B. Kennel *et al.*, "Determining embedding dimension for phase-space reconstruction using a geometrical construction," *Phys. Rev. A*, vol. 45, no. 6, pp. 3403–3411, 1992.

[26] N. V. Chawla *et al.*, "SMOTE: Synthetic minority over-sampling technique," *J. Artif. Intell. Res.*, vol. 16, pp. 321–357, 2002.

[27] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 9, pp. 1263–1284, 2009.

[28] A. Shen *et al.*, "A data-centric and interpretable EEG framework for depression severity grading using SHAP-based insights," *J. Neuroeng. Rehabil.*, vol. 22, no. 1, pp. 1–16, 2025.

[29] F. Lotte *et al.*, "A review of classification algorithms for EEG-based brain–computer interfaces: A 10 year update," *J. Neural Eng.*, vol. 15, no. 3, p. 031005, 2018.

[30] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, pp. 5–32, 2001.

[31] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.

[32] M. W. Howard *et al.*, "Gamma oscillations correlate with working memory load in humans," *Cereb. Cortex*, vol. 13, no. 12, pp. 1369–1374, 2003.

[33] F. Aoki *et al.*, "Increased gamma-range activity in human sensorimotor cortex during performance of visuomotor tasks," *Clin. Neurophysiol.*, vol. 110, no. 3, pp. 524–537, 1999.

[34] L. Pei *et al.*, "Comparative analysis of multifaceted neural effects associated with varying endogenous cognitive load," *Commun. Biol.*, vol. 6, no. 1, p. 795, 2023.

[35] P.-M. Matta *et al.*, "Modulation of beta oscillatory dynamics in motor and frontal areas during physical fatigue," *Commun. Biol.*, vol. 8, no. 1, p. 687, 2025.

[36] W. Klimesch *et al.*, "EEG alpha oscillations: The inhibition–timing hypothesis," *Brain Res. Rev.*, vol. 53, no. 1, pp. 63–88, 2007.

[37] P. J. Uhlhaas *et al.*, "The development of neural synchrony reflects late maturation and restructuring of functional networks in humans," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 106, no. 24, pp. 9866–9871, 2009.

[38] M. Assem *et al.*, "High gamma activity distinguishes frontal cognitive control regions from adjacent cortical networks," *Cortex*, vol. 159, pp. 286–298, 2023.

[39] Y.-W. Kim *et al.*, "The importance of low-frequency alpha (8–10 Hz) waves and default mode network in behavioral inhibition," *Clin. Psychopharmacol. Neurosci.*, vol. 22, no. 1, pp. 53–62, 2023.

[40] E. Ünsal *et al.*, "From infancy to childhood: A comprehensive review of event- and task-related brain oscillations," *Brain Sci.*, vol. 14, no. 8, p. 837, 2024.

[41] S. D. McKeon *et al.*, "Age-related differences in transient gamma band activity during working memory maintenance through adolescence," *Neuroimage*, vol. 274, p. 120112, 2023.

[42] D. A. Fair *et al.*, "Functional brain networks develop from a "local to distributed" organization," *PLoS Comput. Biol.*, vol. 5, no. 5, p. e1000381, 2009.

[43] C. M. Michel and T. Koenig, "EEG microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: A review," *Neuroimage*, vol. 180, pp. 577–593, 2018.

[44] B. A. Seitzman *et al.*, "Cognitive manipulation of brain electric microstates," *Neuroimage*, vol. 146, pp. 533–543, 2017.

[45] A. J. Shackman *et al.*, "Electromyogenic artifacts and electroencephalographic inferences," *Brain Topogr.*, vol. 22, pp. 7–12, 2009.