Brain-State-Dependent Biomarker Discovery from Multi-Condition EEG: A Structured Signal Framework for Schizophrenia with Interpretable Graph Neural Analysis

Anonymous Authors[1]

Abstract

Multi-condition EEG is a structured physiological signal whose utility as a clinical digital biomarker is fundamentally brain-state-dependent—a property ignored by single-condition paradigms that assume state-invariant biomarkers. We propose a dual-stream GCN-Transformer on epoch-wise dynamic ROI-level PLV graphs across six cognitive paradigms on 60 subjects (40 schizophrenia, 20 healthy controls), using structured graph interventions to quantify state-dependent regional importance. Music perception yields the strongest digital biomarker signal (84.73%, $\pm 1.23\%$)—outperforming resting-state by 4.36% with $2\times$ lower variance—while we observe that occipital regions are dominant in five of six conditions and frontal regions emerge exclusively under post-task and working-memory paradigms, with the same region being informative in one state and suppressive in another. LOSO validation achieves $92.00 \pm 2.45\%$ (Warsaw) and $89.00 \pm 2.62\%$ (SU-SZ), outperforming all baselines, with a two-condition music-arithmetic protocol as a promising biomarker acquisition strategy.

1. Introduction

Physiological time series from the brain are structured complex signals whose discriminative properties—and therefore their value as clinical digital biomarkers—are modulated by cognitive state. EEG exemplifies this: the same patient produces fundamentally different inter-regional synchrony patterns at rest versus during music or arithmetic (Molina et al., 2025; Kutepov et al., 2020), meaning biomarker reliability is inseparable from recording context. Yet existing EEG-based psychiatric classifiers treat brain state as fixed, operating under a single condition and embedding the untested assumption of state-invariant biomarkers (Gosala et al., 2025; Alsharabi et al., 2025; Modak et al., 2025; Sharma & Meena, 2024).

The evidence against this assumption is substantial: resting-state recordings suppress discriminative frontal and occipital dynamics that are fully expressed under task conditions (González-Villar et al., 2022; Fernández et al., 2008; Repovš & Barch, 2012; Pei et al., 2025). This suggests that EEG biomarker inconsistency across studies reflects *condition-sampling variability*—a biomarker detectable under one brain state may be absent or reversed under another—not methodological noise.

We address this as a structured complex signal and digital biomarker problem: *which brain regions drive reliable biomarker signal, and how does this shift across brain states?* A dual-stream GCN-Transformer on epoch-wise dynamic PLV-coherence ROI graphs encodes both instantaneous inter-regional synchrony and its temporal evolution. Structured graph interventions—removing each region simultaneously from node features and topology—let us observe directly: occipital removal degrades performance under music but not arithmetic, suggesting it is a condition-specific biomarker driver, not a universal correlate. These observations explain EEG biomarker replication failures as condition-sampling confounds rather than noise. The principal contributions are:

1. A dual-stream GCN-Transformer on epoch-wise dynamic ROI-level PLV graphs that models multi-condition EEG as a structured complex signal, jointly capturing instantaneous inter-regional synchrony and its temporal evolution across six functionally distinct cognitive paradigms.

2. Empirical identification of music perception as the optimal digital biomarker acquisition context—4.36% higher accuracy and $2\times$ lower variance than resting-state EEG—with a clinically actionable two-condition music-arithmetic protocol that maximizes biomarker yield while minimizing recording burden.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

3. Intervention-based observation that occipital connectivity drives the schizophrenia biomarker signal under music but not arithmetic, providing a potential explanation for variability observed in prior EEG biomarker studies.

4. Cross-dataset validation across three independent cohorts confirming occipital PLV dysconnectivity as an acquisition-invariant biomarker substrate of schizophrenia.

2. Methodology

2.1. Dataset and Preprocessing

A clinical EEG dataset was acquired at ABVIMS–Dr. Ram Manohar Lohia Hospital, New Delhi (IEC Ref. No. TP(MD/MS)124/2024/IEC/ABVIMS/RMLH/137), comprising 60 subjects: 40 schizophrenia (SZ), 20 healthy controls (HC), across six paradigms: EO1, EC1, music (passive), arithmetic (serial subtraction), EO2, EC2.

Raw EEG (60 channels) was bandpass filtered at 0.5–50 Hz (4th-order Butterworth) and segmented into 8-second epochs with 7-second steps ($N_{\text{epochs}} = \lfloor (T_{\text{total}} - 8)/7 \rfloor + 1$). EOG and mastoid reference channels were excluded.

2.2. Structured Graph Representation

All 60 channels are grouped into seven anatomically motivated ROIs: Frontal, Fronto-Central, Central, Central-Parietal, Parietal, Parieto-Occipital, and Occipital (full electrode assignments in Appendix A). This parcellation converts raw high-dimensional EEG into a structured graph signal amenable to interpretable region-level analysis.

For each epoch, a 7×7 ROI-level adjacency matrix A_{spatial} is constructed dynamically. Channel-level hybrid connectivity combines coherence and PLV:

$$C_{ij} = \tfrac{1}{2}\left(\frac{|P_{xy}|^2}{P_{xx}P_{yy}} + \left| \frac{1}{N} \sum_t e^{i(\varphi_x(t) - \varphi_y(t))} \right| \right) \quad (1)$$

ROI-level entries average over all cross-region channel pairs:

$$A_{\text{spatial}}^{r_i r_j} = \begin{cases} 1.0 & r_i = r_j \\ \frac{1}{|\mathcal{C}_{r_i}||\mathcal{C}_{r_j}|} \sum_{p \in \mathcal{C}_{r_i}} \sum_{q \in \mathcal{C}_{r_j}} C_{pq} & \text{otherwise} \end{cases} \quad (2)$$

Crucially, each ROI's node feature is its *row* in this matrix—a 7-dimensional connectivity profile—so the adjacency serves dual purpose as both graph topology and node features, eliminating handcrafted signal features. A temporal adjacency $A_{\text{temporal}} \in \mathbb{R}^{13 \times 13}$ chains consecutive epochs (tridiagonal, $A^{ij} = 1$ iff $|i - j| = 1$) to capture sequential evolution of inter-ROI coupling.

Table 1. Per-condition LOSO classification performance across six cognitive states. Task conditions yield both higher accuracy and lower variance.

Condition	Acc. (%)	Prec. (%)	Rec. (%)	F1 (%)
Music	84.73±1.23	85.12±1.41	84.73±1.56	84.89±1.38
Subtraction	83.91±1.47	84.47±1.62	83.91±1.83	84.17±1.71
EO1	82.56±1.89	83.01±2.03	82.56±1.94	82.77±1.97
EO2	81.84±2.12	82.23±2.31	81.84±2.18	82.02±2.24
EC1	80.92±2.43	81.45±2.67	80.92±2.51	81.17±2.58
EC2	80.37±2.61	81.02±2.84	80.37±2.73	80.68±2.79

2.3. Dual-Stream GCN-Transformer

The architecture (Fig. 1) processes EEG through two parallel streams capturing complementary structured signal dimensions.

Spatial stream processes a single-epoch 7×7 ROI graph through $L = 3$ GCN-Transformer fusion layers, each combining:

$$H_{\text{GCN}}^{(l)} = \sigma\left(\tilde{A}_{\text{spatial}} H^{(l-1)} W_{\text{GCN}}^{(l)} \right) \quad (3)$$

$$H_{\text{Attn}}^{(l)} = \text{MHA}(Q, K, V, A_{\text{spatial}}) \quad (4)$$

where $\tilde{A}_{\text{spatial}} = D^{-1/2}(A_{\text{spatial}} + I)D^{-1/2}$ and MHA masks attention scores for disconnected ROI pairs. After $L = 3$ layers, global mean pooling yields $f_{\text{spatial}} \in \mathbb{R}^{128}$.

Temporal stream embeds each epoch's 7×7 matrix via a two-layer MLP to $\mathbb{R}^{128}$, forming a sequence of 13 embeddings processed through $L = 2$ GCN-Transformer layers using A_{temporal}. The center epoch representation yields $f_{\text{temporal}} \in \mathbb{R}^{128}$.

Fusion: $f_{\text{fused}} = \text{Dropout}(\text{GELU}(W_{\text{fuse}}[f_{\text{spatial}} \oplus f_{\text{temporal}}]))$, classified by a two-layer MLP with Softmax. Training: Adam (lr $= 10^{-4}$, batch=64), 100 epochs, cross-entropy loss. Total parameters: 1.25M; inference latency: 0.348 ± 0.043 ms/sample.

3. Structured ROI Importance Analysis

3.1. Intervention Protocol and Classification Performance

For each ROI r, we simultaneously zero its node feature vector and the corresponding rows/columns of A_{spatial}, eliminating the region from both feature space and connectivity topology. The accuracy change $\Delta\text{Acc} = \text{Acc}_{\text{removed}} - \text{Acc}_{\text{full}}$ is negative when a region is discriminatively important, positive when suppressive or redundant—providing a dual-direction characterization of regional contributions as a function of cognitive context.

Accuracy ranges from 80.37% (EC2) to 84.73% (Music): task conditions outperform all resting-state conditions by up to 4.36% and exhibit $2\times$ tighter variance (± 1.23–1.47%

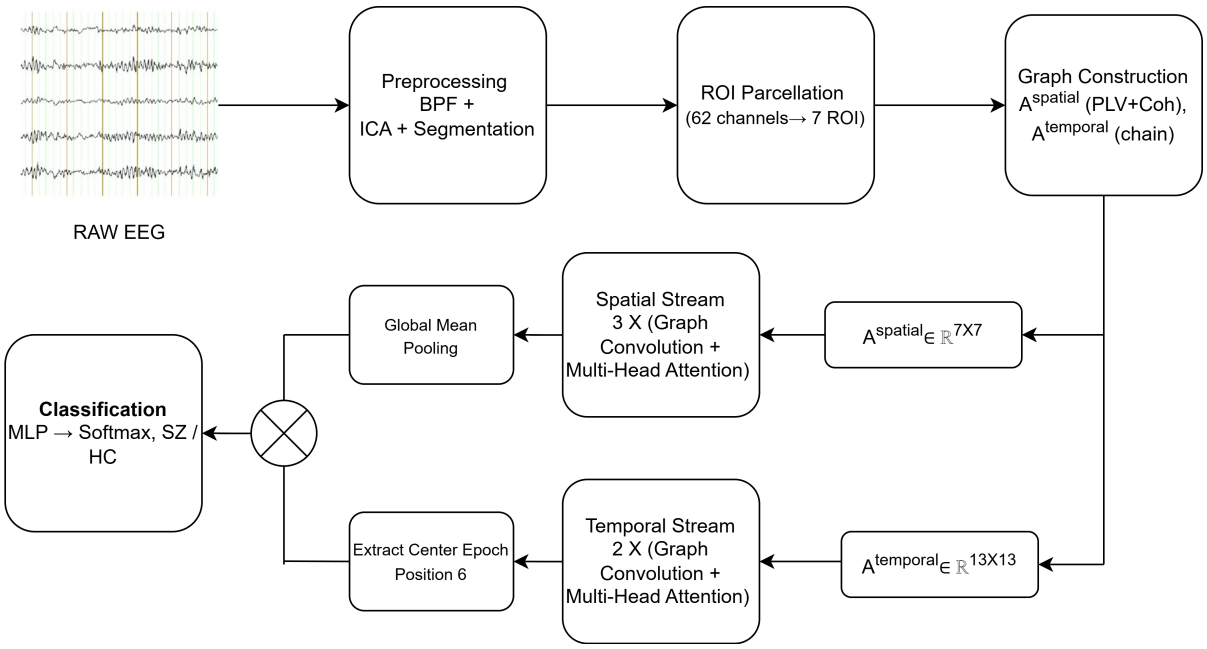

Figure 1. Proposed dual-stream graph transformer. Spatial stream processes instantaneous ROI connectivity; temporal stream processes a 13-epoch context window. Late fusion enables joint spatiotemporal classification.

Table 2. ROI importance: ΔAcc (%) per condition. Negative values indicate discriminatively important regions. The same region can be critical in one state and suppressive in another.

ROI	Mus	Sub	EO1	EO2	EC1	EC2
Frontal	+6.8	−3.8	+0.9	−1.8	+1.7	−11.0
Fr-Cent	+0.9	+7.2	+8.1	−1.2	−2.5	−4.9
Central	+4.9	+2.5	−1.4	+3.5	+4.1	−6.5
Cent-Par	+0.1	+2.3	+0.5	−1.6	+3.3	+0.1
Parietal	+10.4	+10.2	−1.7	−0.9	−0.9	−3.2
Par-Occ	+0.5	+5.5	−2.4	−0.6	−1.0	−5.5
Occipital	−10.8	+2.4	−6.0	−2.4	−2.9	−5.8

vs. ±2.43–2.61%), confirming that music yields the most stable digital biomarker signal.

3.2. State-Dependent ROI Importance

Occipital ($\Delta\text{Acc} = -10.8\%$ under Music) is the most important region in five of six conditions yet suppressive under Subtraction (+2.4%). Frontal is the inverse: suppressive under Music (+6.8%) but dominant under EC2 (−11.0%, the largest ΔAcc overall) and Subtraction (−3.8%). Parietal is strongly suppressive under task but contributive during rest—a dissociation entirely invisible to single-condition protocols.

Wilcoxon signed-rank tests with Benjamini–Hochberg FDR correction (Appendix B) confirm robustness: Music and EC2 are the only conditions at ***-level—occipital ($p = 0.0056$) and frontal ($p = 0.0042$) respectively—with EC2 showing four ROIs at ** or better, confirming broad post-task dysconnectivity as a statistically genuine pattern.

4. Cross-Dataset Generalization

Evaluated on Warsaw (Olejarczyk & Jernajczyk, 2017) (28 subj., 4-ROI) and SU-SZ (Racz et al., 2025) (77 subj., 7-ROI) under eyes-closed resting, occipital ranks as the most important region on all three cohorts—an *acquisition-invariant digital biomarker*. Full per-ROI tables in Appendix C. LOSO accuracy comparison against all available baselines is in Appendix D: the proposed framework achieves 92.00±2.45% (Warsaw) and 89.00±2.62% (SU-SZ), outperforming all baselines by up to 24 points under identical subject-independent evaluation.

Table 3. Cross-dataset ROI importance under eyes-closed resting. Occipital is the most important region on all three independent cohorts.

Dataset	Most Important ROI	ΔAcc (%)
RML (60 subj.)	Occipital	-2.89
	Fronto-Central	-2.51
Warsaw (28 subj.)	Occipital	-2.63
	Frontal	-1.84
SU-SZ (77 subj.)	Occipital	-2.71
	Fronto-Central	-2.34

5. Discussion

Condition-sampling variability as an explanation for biomarker inconsistency. Structured graph interventions let us observe directly: removing occipital degrades performance under Music ($\Delta\text{Acc} = -10.79\%$, FDR $p = 0.0056$) but improves it under Subtraction ($+2.43\%$), suggesting occipital connectivity is a condition-specific driver of the schizophrenia biomarker signal—not a universal one. Conversely, removing frontal is most damaging under EC2 ($\Delta\text{Acc} = -10.96\%$, FDR $p = 0.0042$) and Subtraction, but beneficial under Music. We observe that a study recording only under passive auditory conditions identifies occipital as the primary biomarker region; one recording only under arithmetic identifies frontal—both locally correct, mutually contradictory, precisely the replication failure pattern pervading the EEG literature (Molina et al., 2025). Brain-state context is not a confound to control for but a structured moderator of which neural circuits drive biomarker validity.

Music as the optimal digital biomarker context. The performance gradient from EC2 (80.37%, $\pm 2.61\%$) to Music (84.73%, $\pm 1.23\%$) is not merely a classification improvement—it reflects a fundamental property of biomarker reliability. Task conditions, particularly music, reduce within-condition PLV variance by stereotyping inter-patient synchrony patterns: cognitive challenge decompensates circuits that are only marginally abnormal at rest, producing consistent, patient-invariant dysconnectivity signatures far more amenable to biomarker extraction than the heterogeneous compensatory strategies adopted during rest. The $2\times$ reduction in standard deviation ($\pm 2.61\%$ to $\pm 1.23\%$) directly quantifies this biomarker stabilization effect. Occipital importance across five of six conditions implicates thalamocortical mechanisms (Ramsay et al., 2019; Woodward & Heckers, 2016) as a robust, acquisition-invariant biomarker substrate—present at rest but maximally exposed under task. A promising clinical strategy is a *two-condition protocol* pairing music and arithmetic: music maximizes occipital biomarker signal; arithmetic maximizes frontal biomarker signal; together they may provide broader coverage of the state-dependent dysconnectivity landscape while minimizing patient recording burden.

6. Conclusion

This work makes a case that brain-state context is a first-class structured variable in physiological time-series modeling—one that fundamentally determines the reliability, localization, and generalizability of digital biomarkers derived from EEG. We modeled multi-condition EEG as a structured complex signal problem: epoch-wise dynamic PLV graphs over seven anatomical ROIs, processed by a dual-stream GCN-Transformer jointly capturing instantaneous inter-regional synchrony and its temporal evolution across six cognitive paradigms. The result simultaneously advances three dimensions central to structured health data: *complex signal modeling*—handling high-dimensional, brain-state-varying physiological time series with neurophysiologically grounded graph structure; *digital biomarker discovery*—identifying music perception as the optimal acquisition context (4.36% higher accuracy, $2\times$ lower variance than resting-state) and prescribing a deployable two-condition protocol; and *trust and interpretability*—structured graph interventions reveal directly that occipital drives music-condition biomarker performance while frontal drives arithmetic and post-task performance, providing a potential explanation for variability observed in prior EEG biomarker studies. Occipital PLV dysconnectivity is confirmed as an acquisition-invariant biomarker substrate across three independent cohorts. Single-condition EEG protocols produce fundamentally incomplete biomarker characterizations; explicit modeling of brain-state context is a prerequisite for reliable, interpretable, and clinically deployable physiological time-series biomarkers.

Impact Statement

This work advances EEG-based psychiatric biomarker discovery with potential to support earlier schizophrenia detection; clinical deployment requires prospective validation and clinician oversight. All datasets are IEC-approved with informed consent; the ROI importance framework promotes interpretable, auditable model behavior.

References

N. Alsharabi, R. K. Mahendran, and G. Alshammari. EEG-SymNet: Multi-channel EEG signal-based schizophrenia diagnosis using channel recalibration and symmetric spatial temporal transformer network. *Eng. Sci. Technol., Int. J.*, 69:102116, 2025.

B. Gosala, A. R. Singh, H. Tiwari, and M. Gupta. GCN-LSTM: A hybrid graph convolutional network model

for schizophrenia classification. *Biomed. Signal Process. Control*, 105:107657, 2025.

S. Modak, K. Samanta, S. Halder, and S. Chatterjee. EEG rhythm-based functional brain connectivity for automated detection of schizophrenia employing deep learning. *IEEE Trans. Instrum. Meas.*, 74:2520309, 2025.

R. Sharma and H. K. Meena. Schizophrenia detection using interconnected graph-based features from EEG signals. *IEEE Trans. Instrum. Meas.*, 73:4009309, 2024.

R. Molina et al. Dynamic evolution of EEG complexity in schizophrenia across cognitive tasks. *Entropy*, 27(3):226, 2025.

I. E. Kutepov et al. Visualization of EEG signal entropy in schizophrenia. *Sci. Visualiz.*, 12(1):1–9, 2020.

J. D. González-Villar et al. Relation between EEG resting-state power and modulation of P300 task-related activity in theta band in schizophrenia. *Prog. Neuropsychopharmacol. Biol. Psychiatry*, 116:110530, 2022.

A. Fernández et al. Abnormal EEG complexity in patients with schizophrenia and depression. *Clin. Neurophysiol.*, 119(6):1343–1352, 2008.

G. Repovš and D. M. Barch. Working memory related brain network connectivity in individuals with schizophrenia and their siblings. *Front. Hum. Neurosci.*, 6:137, 2012.

O. Meiron, J. David, and A. Yaniv. Early auditory processing predicts efficient working memory functioning in schizophrenia. *Brain Sci.*, 12(2):212, 2022.

H. Pei, H. Li, C. Hou, Y. Liu, J. Liu, M. Duan et al. Fronto-occipital dyscommunication associates with brain hierarchy in schizophrenia. *Commun. Biol.*, 8(1):699, 2025.

E. Olejarczyk and W. Jernajczyk. Graph-based analysis of brain connectivity in schizophrenia. *PLoS One*, 12(11):e0188629, 2017.

F. S. Racz et al. Reduced temporal variability of cortical excitation/inhibition ratio in schizophrenia. *Schizophrenia (Heidelb)*, 11(1):20, 2025.

E. Aydemir, S. Dogan, M. Baygin, C. P. Ooi, P. D. Barua, T. Tuncer, and U. R. Acharya. CGP17Pat: Automated schizophrenia detection based on a cyclic group of prime order patterns using EEG signals. *Healthcare*, 10(4):643, 2022.

F. S. Racz and G. Csukly. Information leakage and performance overestimation in EEG-based schizophrenia detection. *medRxiv* (preprint), 2025.

I. S. Ramsay et al. An activation likelihood estimate meta-analysis of thalamocortical dysconnectivity in psychosis. *Biol. Psychiatry Cogn. Neurosci. Neuroimaging*, 4(9):859–869, 2019.

N. D. Woodward and S. Heckers. Mapping thalamocortical functional connectivity in chronic and early stages of psychotic disorders. *Biol. Psychiatry*, 79(12):1016–1025, 2016.

H. Hughes et al. GABAergic dysfunction in postmortem dorsolateral prefrontal cortex: implications for cognitive deficits in schizophrenia. *Front. Cell. Neurosci.*, 18:1440834, 2024.

P. C. Tu, Y. M. Bai, C. T. Li, M. H. Chen, W. C. Lin, W. C. Chang, and T. P. Su. Identification of common thalamocortical dysconnectivity in four major psychiatric disorders. *Schizophr. Bull.*, 45(5):1143–1151, 2019.

N. D. Woodward, H. Karbasforoushan, and S. Heckers. Thalamocortical dysconnectivity in schizophrenia. *Am. J. Psychiatry*, 169(10):1092–1099, 2012.

E. Pomarol-Clotet, R. Salvador, S. Sarró, J. Gomar, F. Vila, A. Martínez et al. Failure to deactivate in the prefrontal cortex in schizophrenia: dysfunction of the default mode network? *Psychol. Med.*, 38(8):1185–1193, 2008.

R. Landin-Romero, P. J. McKenna, P. Salgado-Pineda, S. Sarró, C. Aguirre, C. Sarri et al. Failure of deactivation in the default mode network: a trait marker for schizophrenia? *Psychol. Med.*, 45(6):1315–1325, 2015.

M. L. Hu, X. F. Zong, J. J. Mann, J. J. Zheng, Y. H. Liao, Z. C. Li et al. A review of the functional and anatomical default mode network in schizophrenia. *Neurosci. Bull.*, 33(1):73–84, 2017.

A. M. Jimenez, J. Lee, J. K. Wynn, M. S. Cohen, S. A. Engel, D. C. Glahn et al. Abnormal ventral and dorsal attention network activity during single and dual target detection in schizophrenia. *Front. Psychol.*, 7:323, 2016.

A. ROI Electrode Assignments

Table 4. ROI Definitions and Constituent Electrodes

ROI	Electrodes
Frontal	Fp1, Fpz, Fp2, AF7, AF3, AF4, AF8, F7, F5, F3, F1, Fz, F2, F4, F6, F8
Fronto-Central	FT7, FC5, FC3, FC1, FCz, FC2, FC4, FC6, FT8
Central	T7, C5, C3, C1, Cz, C2, C4, C6, T8
Central-Parietal	TP7, CP5, CP3, CP1, CPz, CP2, CP4, CP6, TP8
Parietal	P7, P5, P3, P1, Pz, P2, P4, P6, P8
Parieto-Occipital	PO7, PO5, PO3, POz, PO4, PO6, PO8
Occipital	O1, Oz, O2

B. Statistical Validation of ROI Importance

Wilcoxon signed-rank tests with Benjamini–Hochberg FDR correction were applied across all seven ROIs per condition. The test is non-parametric and operates on paired per-fold accuracy differences, making it appropriate for LOSO distributions. BH-FDR controls the expected proportion of false discoveries, offering greater power than Bonferroni when multiple true effects are present.

Table 5. Most important and most suppressive ROI per condition with FDR-corrected significance (Wilcoxon signed-rank + Benjamini–Hochberg).

Cond.	ROI	ΔAcc (%)	FDR (p)	Sig.
Music	Occipital (imp.)	-10.79	0.0056	***
	Parietal (supp.)	$+10.37$	0.0581	ns
Sub	Frontal (imp.)	-3.82	0.0217	**
	Parietal (supp.)	$+10.17$	0.0623	ns
EO1	Occipital (imp.)	-6.00	0.0147	**
	Fr-Cent (supp.)	$+8.13$	0.0497	*
EO2	Frontal (imp.)	-1.77	0.2184	ns
	Central (supp.)	$+3.54$	0.1673	ns
EC1	Occipital (imp.)	-2.89	0.0794	ns
	Central (supp.)	$+4.08$	0.1438	ns
EC2	Frontal (imp.)	-10.96	0.0042	***
	Cent-Par (supp.)	$+0.12$	0.5923	ns

C. Full Cross-Dataset ROI Importance Tables

Warsaw Dataset (4-ROI)		*SU-SZ Dataset (7-ROI)*	
ROI	ΔAcc (%)	ROI	ΔAcc (%)
Occipital	-2.63	Occipital	-2.71
Frontal	-1.84	Fronto-Central	-2.34
Central	$+1.47$	Parieto-Occipital	-1.12
Parietal	$+3.91$	Parietal	-0.87
		Frontal	$+1.54$
		Central-Parietal	$+2.98$
		Central	$+3.84$

D. State-of-the-Art Comparison

All comparisons are restricted to LOSO validation—the most stringent subject-independent standard—ensuring no subject-level data leakage.

Table 6. LOSO accuracy comparison on Warsaw and SU-SZ datasets.

Method	Warsaw (%)	SU-SZ (%)
CGP17Pat (Aydemir et al., 2022)	84.33	–
CNN (Racz & Csukly, 2025)	58.06±7.45	67.66±10.95
Ensemble (Racz & Csukly, 2025)	64.73±13.58	64.91±8.96
SVM (Racz & Csukly, 2025)	56.38±35.77	72.84±31.01
Proposed (LOSO)	**92.00±2.45**	**89.00±2.62**

The proposed framework achieves $92.00 \pm 2.45\%$ on Warsaw (+7.67 over CGP17Pat) and $89.00 \pm 2.62\%$ on SU-SZ (+16–24 over all baselines), with substantially lower variance confirming that ROI-level PLV graphs provide neurophysiological inductive bias that generic classifiers cannot recover.

E. Per-Condition ROI Importance: Detailed Analysis

Music. Occipital ($\Delta\text{Acc} = -10.79\%$) is the overwhelmingly dominant contributor—consistent with posterior cortical roles in auditory-visual integration and predictive coding during music listening. Parietal ($+10.37\%$) and frontal ($+6.83\%$) are strongly suppressive, indicating their inclusion distorts rather than preserves condition-specific discriminative signals.

Subtraction. Frontal emerges as the sole strongly important region (-3.82%), consistent with prefrontal working memory and executive function demands. Occipital ($+2.43\%$) transitions to suppressive—a direct contrast to its music-condition dominance. Parietal ($+10.17\%$) and fronto-central ($+7.15\%$) show the largest positive ΔAcc values.

EO1. Occipital (-6.00%) leads, followed by parieto-occipital (-2.43%), parietal (-1.74%), and central (-1.40%). Broad posterior dominance reflects schizophrenia-related disruptions in visual cortex synchronization. Fronto-central ($+8.13\%$) is the largest suppressive region.

EO2. The most distributed profile: six of seven ROIs show negative or near-zero ΔAcc, with only central ($+3.54\%$) suppressive, suggesting post-task residual activations persist across multiple networks simultaneously.

EC1. Occipital (-2.89%) and fronto-central (-2.51%) co-dominate, suggesting thalamocortical loop activity manifests across both posterior sensory and frontal relay regions during eyes-closed rest.

EC2. Frontal dominates (-10.96%), with central (-6.46%), occipital (-5.78%), parieto-occipital (-5.49%), fronto-central (-4.90%), and parietal (-3.18%) all contributing. Only central-parietal ($+0.12\%$) is neutral—consistent with impaired default mode network reactivation and failure of post-task neural reset in schizophrenia.

F. Spatial Attention Weight Analysis

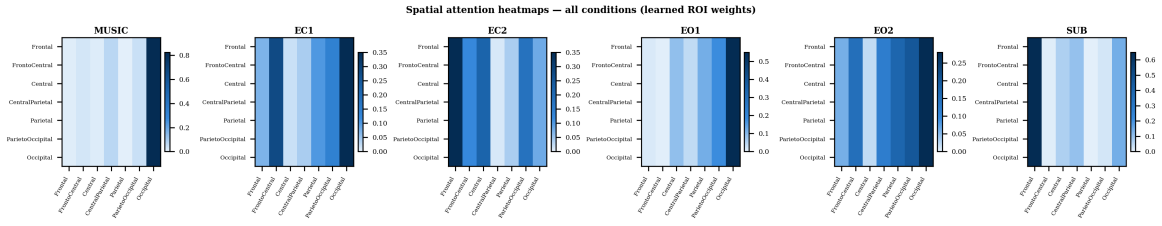

Figure 2. Spatial attention heatmaps (7×7) averaged across all attention heads and test subjects per condition. The highest-weighted ROI consistently matches the most important ROI from the removal analysis.

The graph-constrained multi-head attention mechanism learns to assign differential importance to inter-ROI connectivity

pairs, with scores masked for disconnected ROI pairs ensuring weights reflect genuine functional coupling. Weights are extracted per head per test subject, averaged across all four heads and LOSO folds, and normalized row-wise per condition.

Music. Occipital receives weight 0.826—near-exclusive posterior focus, consistent with $\Delta\text{Acc} = -10.79\%$ and pulvinar-to-occipital thalamocortical relay dysconnectivity.

EO1. Occipital leads (0.55) followed by parieto-occipital (0.18) and parietal (0.12)—broad posterior dominance during eyes-open resting.

EO2. Distributed: occipital (0.28), parieto-occipital (0.20), parietal (0.17), fronto-central (0.15)—consistent with post-task multi-network residual activations.

EC1. Occipital leads (0.35) but fronto-central rises substantially (0.28), reflecting emergence of thalamocortical loop dynamics during eyes-closed rest.

EC2. Near-inversion of Music: frontal dominates (0.35), central (0.22), occipital recedes to 0.08—fully corroborating $\Delta\text{Acc} = -10.96\%$.

SUB. Frontal receives the highest single-ROI weight across all conditions (0.65), confirming prefrontal connectivity as the primary discriminative signal under maximal working memory engagement.

Across all six conditions, the highest-weighted ROI matches the most important ROI from the removal analysis—providing convergent independent validation.

G. Ablation Study

Table 7. Baseline comparison and ablation: LOSO accuracy (%) across all six conditions.

Method	Music	Sub	EO1	EO2	EC1	EC2
SVM (RBF)	67.43±4.12	65.87±4.56	63.21±5.03	62.84±5.17	61.56±5.43	60.92±5.71
Random Forest	71.28±3.84	69.54±4.21	67.83±4.67	66.47±4.89	65.12±5.02	64.73±5.34
Transformer Only	76.54±3.21	74.83±3.47	72.64±3.89	71.93±4.12	70.47±4.38	69.81±4.52
Spatial Only	80.12±2.67	78.43±2.91	76.84±3.12	75.63±3.34	74.21±3.58	73.84±3.72
Temporal Only	78.34±2.89	76.92±3.14	75.12±3.43	74.27±3.61	72.83±3.84	72.14±3.97
Dual-Stream	**84.73±1.23**	**83.91±1.47**	**82.56±1.89**	**81.84±2.12**	**80.92±2.43**	**80.37±2.61**

Conventional classifiers show the weakest performance, confirming that flattened ROI connectivity vectors without graph-based propagation are insufficient for capturing spatiotemporal dysconnectivity. The standalone Transformer falls below both graph-based single-stream variants, validating PLV-coherence adjacency as essential structural bias. Spatial stream consistently outperforms temporal, reflecting primacy of instantaneous inter-ROI dysconnectivity. The full dual-stream outperforms both single-stream variants by 3.0–5.5% across all conditions.

H. Neurophysiological Discussion: Mechanistic Grounding

Occipital stability as a thalamocortical anchor. Occipital importance across five of six conditions—including non-visual paradigms—implicates thalamocortical rather than purely visual mechanisms. A meta-analysis across 20 experiments established thalamocortical hyperconnectivity in psychosis specifically including lateral occipital cortex (Ramsay et al., 2019), reflecting chronic abnormality in the pulvinar-to-occipital relay pathway (Woodward & Heckers, 2016). Parvalbumin interneuron dysfunction disrupts GABAergic inhibition in V1, producing gamma-band synchrony deficits independent of task context (Hughes et al., 2024). Notably, one resting-state fMRI study found no group differences in occipital-thalamic connectivity at rest, suggesting task conditions are necessary to expose this dysfunction—corroborating stronger occipital importance under music ($\Delta\text{Acc} = -10.79\%$) versus eyes-closed rest (EC1: $\Delta\text{Acc} = -2.89\%$).

Frontal dominance under EC2: post-task neural reset failure. Frontal dominance under EC2 ($\Delta\text{Acc} = -10.96\%$, FDR $p = 0.0042$)—the largest single ΔAcc across all 42 ROI-condition combinations—receives strong support from the DMN literature. In healthy subjects, transition from active cognition to rest triggers robust DMN reactivation and deactivation of task-positive frontal regions. In schizophrenia, this post-task reset is systematically impaired: SZ patients show failure

of deactivation in anterior cingulate and ventromedial PFC following task completion, a heritable trait marker present in unaffected first-degree relatives. EC2—recorded immediately after arithmetic, eyes closed—captures precisely this failure mode at the EEG connectivity level.