This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Adaptive Riemannian Geometry for Seizure Onset Zone Localization in Intracranial Electroencephalography

Ghazaleh Khodabandelou[1,2,*], Cemal Karakas[3], Marius George Linguraru[2,4], and Syed Muhammad Anwar[2,4]

[1] Université Paris-Est Créteil (UPEC), LISSI, F-94010 Créteil, France
[2] Sheikh Zayed Institute for Pediatric Surgical Innovation, Children's National Hospital, Washington DC, USA
[3] Division of Epilepsy and Clinical Neurophysiology, Children's National Hospital, Washington, DC, USA
[4] Departments of Radiology and Pediatrics, George Washington University, Washington DC, USA.
gkhodaband@childrensnational.org, ghazaleh.khodabandelou@u-pec.fr

Abstract. Accurate localization of the seizure onset zone (SOZ) is critical for presurgical epilepsy planning, yet automated channel-level SOZ detection from raw inter-ictal intracranial EEG (iEEG) remains underexplored. Existing outcome-prediction pipelines assume expert SOZ labels as input, creating a circular dependency that motivates end-to-end localization from raw recordings. We propose *Adaptive Riemannian Geometry on Symmetric Positive Definite Neural Manifolds* (ARG-SPD), a geometric deep learning framework for channel-level SOZ candidate detection. ARG-SPD encodes raw iEEG windows with a temporal convolutional neural network, projects features to Symmetric Positive Definite (SPD) covariance matrices on a Riemannian manifold, and classifies them via learnable SPD prototypes with affine-invariant geodesic distances. A Riemannian Focal Loss with spectral dispersion regularization addresses extreme class imbalance (∼8.8% SOZ channels) and numerical instability. Monte Carlo dropout provides uncertainty-aware triage, separating high-confidence SOZ candidates from ambiguous cases for expert review. On 38,708 interictal iEEG windows from 184 patients (OpenNeuro public dataset) under strict patient-level 5-fold cross-validation, ARG-SPD achieves $\mathbf{0.82 \pm 0.03}$ channel-AUC, outperforming SPDNet by +0.06 AUC (p<0.05) and a 1D-CNN baseline by +0.17. High-confidence predictions (78% of channels) reach 0.89 AUC, with spatial overlap of 78%±15% Dice against resection zones in a 47-patient subset. These results support ARG-SPD as a promising expert-in-the-loop decision-support tool for presurgical SOZ mapping, pending prospective independent validation.

Keywords: Seizure Onset Zone Localization · Intracranial EEG · Epilepsy Surgery Planning · Riemannian Geometry · SPD Manifolds · Imbalanced Learning · Uncertainty-Aware Decision Support · Medical AI.

1 Introduction

Drug-resistant epilepsy affects approximately 30% of patients on medication, and surgical resection of the epileptogenic zone—including the seizure onset zone (SOZ)—remains the most effective treatment. However, accurate SOZ identification from intracranial electroencephalography (iEEG) currently requires expert epileptologists to visually review hours of ictal and interictal recordings, a process that is labor-intensive, subject to inter-rater variability, and not standardized across centers [12,2].

Automated SOZ localization from raw inter-ictal iEEG is therefore a high-value clinical problem, yet it remains largely underexplored. A key obstacle is a *circular dependency*: recent outcome-prediction pipelines report strong patient-level area under the curve (AUC) [17], but they require expert SOZ annotations as input features, rather than estimating them. This highlights the need for a method that addresses the *preceding* localization step: estimating SOZ-likely channels from raw inter-ictal recordings without relying on expert labels at inference time.

Three fundamental challenges make automated channel-level SOZ detection difficult: **(i) Class imbalance:** Typically only ∼8–9% of sEEG contacts (could exceed 100) lie within the SOZ (∼11:1 imbalance), which standard cross-entropy optimization fails to handle effectively. **(ii) Geometric structure:** Cross-channel iEEG covariances form Symmetric Positive Definite (SPD) matrices residing on a Riemannian manifold $\mathcal{S}_{++}(d)$. Euclidean-based deep learning ignores this intrinsic connectivity geometry [18,11]. **(iii) Multi-site heterogeneity:** Differences in electrode types (depth vs. grid), implantation strategies, and acquisition settings between centers create distribution shifts that degrade generalization.

Related work. Prior neural approaches to iEEG analysis have primarily addressed *adjacent* tasks: seizure detection and prediction [15,6,4,14], high-frequency oscillation (HFO) biomarker detection [7,10,12], scalp-EEG localization [3], and patient-level surgical outcome prediction [17,1]. For channel-level localization from raw interictal iEEG specifically, direct benchmarking is difficult because existing studies differ in labels, inputs (ictal vs. interictal), handcrafted feature pipelines, or multimodal anatomical data. SPDNet [11] provides the closest geometric baseline but was not evaluated on this task; we adapt and evaluate it under the same protocol. Graph-based methods [9,4] address related localization problems but rely on electrode coordinate graphs not available in all centers. ARG-SPD addresses channel-level SOZ detection from raw interictal iEEG under a strict patient-split protocol.

Contributions. ARG-SPD makes the following contributions: (i) a Riemannian Focal Loss with spectral dispersion regularization for imbalanced SPD classification; (ii) learnable SPD class prototypes induced by unconstrained parameters and compared via affine-invariant geodesic distances; (iii) strict patient-level evaluation on 184 patients and 38,708 interictal iEEG windows; and (iv) uncertainty-aware triage via MC Dropout and post-hoc spatial concordance analysis with resection zones. ARG-SPD achieves 0.82 ± 0.03 channel-AUC, significantly

outperforming geometric baselines (+0.06 over SPDNet, $p<0.05$), while providing calibrated uncertainty estimates for expert-in-the-loop presurgical planning.

2 Data and Evaluation Protocol

We use the publicly available OpenNeuro iEEG dataset [16], which includes N_p=185 subjects with clinical SOZ annotations. After preprocessing and validity checks (see below), we retain N_p=184 patients and N_w=38,708 iEEG windows. Recordings originate from two clinical sites (Detroit and UCLA), enabling cross-site evaluation.

Preprocessing. Each window spans 15 seconds at 512 Hz (T=7,680 samples per channel). Signals are bandpass-filtered (0.5–250 Hz, FIR, order 100) and normalized by per-window z-score without manual artifact rejection. Channels with zero variance or severe artifacts (>50% clipped samples) are excluded. SOZ-labeled windows constitute approximately 8.8% of the retained data.

Evaluation protocol. All evaluation uses strict patient-level 5-fold Group-K-Fold splitting: all windows and channels from a given patient appear exclusively in one fold, preventing leakage. Reported metrics include channel-AUC (area under ROC curve), channel-AUPRC, sensitivity at 90% specificity, and macro-averaged patient-AUC. Channel-level predictions are obtained by aggregating window-level probabilities (Eq. 1). Statistical significance is assessed by Wilcoxon signed-rank test on per-patient AUCs (α=0.05, Bonferroni-corrected).

3 Methodology

3.1 Problem Formulation and Notation

Let $\mathcal{D}=\{(\mathbf{x}_i, y_i, s_i, c_i)\}_{i=1}^{N_w}$ denote the set of iEEG windows, where N_w=38,708. Each window $\mathbf{x}_i \in \mathbb{R}^T$ contains T=7,680 samples (15 seconds at 512 Hz) from a single iEEG channel. The label $y_i \in \{0, 1\}$ indicates non-SOZ (0) or SOZ (1). Each window is associated with a recording site s_i, and a channel identifier $c_i=(p_i, m_i)$, where p_i is the patient identifier and m_i the electrode contact name. The goal is to estimate a SOZ probability $\hat{p}(c)$ for each channel c, using only raw interictal iEEG and no expert SOZ annotations at inference time.

3.2 Window-Level Training and Channel-Level Inference

ARG-SPD is trained at the window level for data efficiency, but the clinical target is channel-level SOZ candidate localization. For a channel c, let $\mathcal{W}(c)=\{i : c_i=c\}$ be the set of all windows from that channel. The model outputs a window-level SOZ probability $\hat{p}_i=f_\theta(\mathbf{x}_i)$. Channel-level predictions are obtained by averaging:

$$\hat{p}(c) = \frac{1}{|\mathcal{W}(c)|} \sum_{i \in \mathcal{W}(c)} \hat{p}_i. \tag{1}$$

All reported channel-level metrics are computed using $\hat{p}(c)$. Patient-level splitting ensures no windows from the same patient appear in both training and test folds.

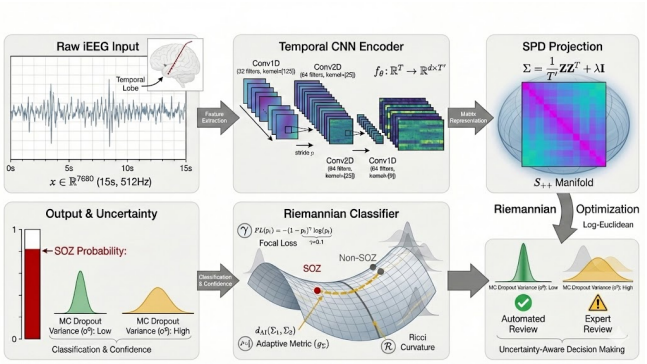

Fig. 1: ARG-SPD overview: raw iEEG $\mathbf{x}\in\mathbb{R}^{7680}$ is encoded by a temporal CNN, projected to an SPD covariance $\Sigma=\frac{1}{T'}\mathbf{Z}\mathbf{Z}^\top+\lambda\mathbf{I}$, and classified using learnable SPD prototypes with affine-invariant distances. MC Dropout (K=50) estimates uncertainty for expert-review triage.

3.3 Architecture Overview

ARG-SPD processes each iEEG window $\mathbf{x}_i \in \mathbb{R}^T$ through three sequential modules (Fig. 1): **(1)** a temporal CNN encoder that maps $\mathbf{x}_i$ to a feature map $\mathbf{Z}_i\in\mathbb{R}^{d\times T'}$; **(2)** an SPD projection module that maps $\mathbf{Z}_i$ to a covariance matrix $\Sigma_i\in\mathcal{S}_{++}(d)$; **(3)** a Riemannian classifier that computes SOZ probabilities from geodesic distances between Σ_i and two learnable SPD class prototypes, one for SOZ and one for non-SOZ. At inference time, MC Dropout over the CNN enables uncertainty estimation.

3.4 SPD Projection and Riemannian Geometry

Given feature map $\mathbf{Z}_i \in \mathbb{R}^{d\times T'}$ from the CNN encoder, the empirical covariance matrix is:

$$\Sigma_i = \frac{1}{T'}\mathbf{Z}_i\mathbf{Z}_i^\top + \lambda\mathbf{I}_d, \tag{2}$$

where λ=10^{-3} ensures $\Sigma_i \succ 0$ and improves numerical conditioning. The affine-invariant (AI) geodesic distance between two SPD matrices is:

$$d_{\mathrm{AI}}(\Sigma_1,\Sigma_2) = \left\|\mathrm{Log}\left(\Sigma_1^{-1/2}\Sigma_2\Sigma_1^{-1/2}\right)\right\|_F, \tag{3}$$

The inverse square root $\Sigma^{-1/2}$=$\mathbf{Q}\mathbf{D}^{-1/2}\mathbf{Q}^\top$ is computed via eigendecomposition with eigenvalue clamping to $[\epsilon,\infty)$, ϵ=10^{-6}, for numerical stability.

3.5 Riemannian Focal Loss for Imbalanced SOZ Detection

Class probabilities are computed from geodesic distances to two learnable SPD class prototypes $\mathbf{C}_{\mathrm{SOZ}}$ and $\mathbf{C}_{\mathrm{nonSOZ}}$. For class k, the logit is the negative geodesic

distance:

$$z_{i,k} = -d_{\mathrm{AI}}(\Sigma_i, \mathbf{C}_k), \qquad p^g_{i,k} = \frac{\exp(z_{i,k})}{\sum_{k'} \exp(z_{i,k'})}. \tag{4}$$

The training objective combines a focal loss term to address class imbalance with a spectral dispersion regularizer:

$$\mathcal{L}_{\mathrm{RFocal}} = -\frac{1}{N_w} \sum_{i=1}^{N_w} \alpha_{y_i} (1 - p^g_{i,y_i})^\gamma \log(p^g_{i,y_i}) + \frac{\beta}{N_w} \sum_{i=1}^{N_w} R(\Sigma_i), \tag{5}$$

where α_{y_i} is the class-specific weight, γ=2.0 is the focal modulation parameter that down-weights easy non-SOZ examples, and β=0.01 controls spectral regularization. The spectral dispersion regularizer $R(\Sigma_i)$ discourages ill-conditioned covariances (see below).

3.6 Learnable SPD Prototypes

To ensure that class prototypes remain in the SPD cone throughout training while avoiding the complexity of explicit manifold-constrained optimization, each prototype is induced by an unconstrained parameter matrix $\mathbf{P}_k \in \mathbb{R}^{d \times d}$:

$$\mathbf{C}_k = \mathbf{P}_k \mathbf{P}_k^\top + \epsilon \mathbf{I}_d, \quad k \in \{\mathrm{SOZ}, \mathrm{nonSOZ}\}. \tag{6}$$

Parameters $\mathbf{P}_k$ are updated via AdamW jointly with the CNN encoder. Optimization is performed in unconstrained Euclidean parameter space, while the induced prototypes $\mathbf{C}_k$ remain SPD-valid and all class comparisons use Riemannian distances on $\mathcal{S}_{++}(d)$.

3.7 Spectral Dispersion Regularization

The regularization term $R(\Sigma_i)$ penalizes ill-conditioned covariances by measuring the variance of log-eigenvalues:

$$R(\Sigma_i) = \frac{1}{d} \sum_{j=1}^{d} \big(\log \lambda_j - \bar{\log \lambda}\big)^2, \quad \bar{\log \lambda} = \frac{1}{d} \sum_{j=1}^{d} \log \lambda_j, \tag{7}$$

where λ_j are eigenvalues of Σ_i. This practical stability penalty prevents numerical failure in d_{AI} when Σ_i approaches singularity, and is complementary to the focal loss.

3.8 Uncertainty Quantification via Monte Carlo Dropout

To support clinical triage, we apply MC Dropout [8] at inference time. The CNN dropout masks are sampled K=50 times per window, yielding stochastic SOZ probability estimates $\{p_i^{(k)}\}_{k=1}^K$. Epistemic uncertainty is quantified as:

$$\mu_i = \frac{1}{K} \sum_{k=1}^{K} p_i^{(k)}, \qquad \sigma_i^2 = \frac{1}{K} \sum_{k=1}^{K} (p_i^{(k)} - \mu_i)^2. \tag{8}$$

High σ_i^2 indicates uncertain predictions; uncertainty thresholds τ_1=0.05 and τ_2=0.15 are calibrated on validation folds to define high-confidence, uncertain, and flagged subsets.

4 Experiments

4.1 Implementation Details

Architecture. The CNN encoder uses 3 convolutional layers (16, 32, 32 filters), kernel sizes [25, 9, 3], strides [6, 2, 1], BatchNorm, ReLU, and dropout (p=0.3). The SPD dimension is fixed to d=32 with regularization λ=10^{-3}. Prototype parameters $\mathbf{P}_k$ are initialized near identity so that initial induced prototypes are close to well-conditioned diagonal matrices.

Training. Optimization uses AdamW with learning rate 10^{-4} (cosine annealing), batch size 16, class-balanced sampling, up to 100 epochs with early stopping (patience 15 on validation AUC). Fixed hyperparameters across all reported experiments: λ=10^{-3}, β=0.01, γ=2.0. Gradient clipping (max norm 1.0). Total wall-clock training time for 5-fold CV is approximately 50 hours on CPU; eigendecomposition cost scales as $\mathcal{O}(bd^3)$ per batch.

Reproducibility. All results are determined by the public OpenNeuro dataset [16], the preprocessing described in Sec. 2, strict patient-level Group-K-Fold splits, and the fixed hyperparameters listed above. Uncertainty thresholds (τ_1, τ_2) are calibrated on validation folds. All baselines are trained and evaluated under the same patient-level split, channel aggregation (Eq. 1), and metric protocol.

4.2 Main Results

Table 1 reports channel-level SOZ candidate detection results. *Importantly, patient-level outcome prediction studies [17,13] are listed only as clinical context, since they use expert-defined SOZ labels as input and address a different task.* Direct comparisons are restricted to methods trained under the same raw interictal, channel-level, patient-split protocol. ARG-SPD achieves 0.82 ± 0.03 channel-AUC, outperforming SPDNet by +0.06 AUC (p=0.003, Wilcoxon) and the 1D-CNN baseline by +0.17. Channel-AUPRC improves from 0.22 (1D-CNN) to 0.42 (ARG-SPD), which is particularly meaningful under the ∼8.8% class imbalance. Sensitivity at 90% specificity increases from 0.48 to 0.71, reflecting improved recovery of SOZ-likely channels at a clinically conservative false-positive operating point.

Uncertainty-based triage. The high-confidence subset (σ^2<0.05, retaining 78% of channels) reaches 0.89 AUC and 0.85 sensitivity at 90% specificity; this 0.89 AUC refers to that subset only. The most uncertain 7% of channels are flagged for mandatory expert review. This design is intended for expert-in-the-loop support, not autonomous decision-making.

Table 1: **Top**: Channel-level SOZ candidate localization from raw interictal iEEG under the same patient-level split. **Bottom**: Patient-level outcome prediction studies are included as clinical context only; they are *not directly comparable* as they use expert SOZ labels as input and address a different prediction task. Best results per section in bold. †$p<0.05$ vs. SPDNet (Wilcoxon signed-rank on per-patient AUCs, Bonferroni-corrected).

Method	Ch-AUC	Ch-AUPRC	Sens@Sp90%	Pat-AUC	Task / Input
Channel-level SOZ detection from raw interictal iEEG (same protocol):					
1D-CNN [19]	0.65±0.04	0.22±0.03	0.48	0.63±0.05	SOZ (channel), raw iEEG
CNN-Transformer [15]	0.65±0.09	0.24±0.05	0.48	0.62±0.08	SOZ (channel), raw iEEG
SPDNet [11]	0.76±0.04	0.31±0.04	0.58	0.74±0.04	SOZ (channel), raw iEEG
Log-Euclidean CNN	0.78±0.03	0.35±0.04	0.62	0.76±0.03	SOZ (channel), raw iEEG
ARG-SPD (Ours)	**0.82±0.03**†	**0.42±0.05**	**0.71**	**0.80±0.03**†	SOZ (channel), raw iEEG
w/o Focal (γ=0)	0.78±0.04	0.35±0.04	0.63	0.76±0.04	Ablation
w/o Dispersion (β=0)	0.79±0.03	0.36±0.04	0.65	0.77±0.03	Ablation
Fixed Prototypes	0.80±0.03	0.38±0.04	0.67	0.78±0.03	Ablation
Patient-level outcome prediction (clinical context only; different task and inputs):					
Partamian et al. [17]	–	–	–	0.86	Outcome, requires expert SOZ

4.3 Ablation Study

Table 2 traces the progressive contribution of each ARG-SPD component. The largest single gain (+0.11 AUC) comes from replacing Euclidean temporal features with SPD covariance representations, confirming that preserving connectivity geometry is beneficial. Naive SPD projection without eigenvalue clipping converges in only 48% of folds due to numerical failures; its AUC/AUPRC are computed only on converged folds, and clamping restores 100% stability. The Log-Euclidean metric adds a small accuracy improvement (+0.02). The Riemannian Focal Loss, learnable prototypes, and spectral dispersion then further improve AUC from 0.78 to 0.82 and AUPRC from 0.35 to 0.42, with gains that are complementary across components.

Table 2: Progressive ablation of ARG-SPD components. Each row adds one component to the previous. *Unstable row: AUC/AUPRC over converged folds only.

Configuration	AUC	ΔAUC	AUPRC	Stability
(1) 1D-CNN baseline	0.65	–	0.22	100%
(2) + SPD projection (naive, no clamping)	0.76*	+0.11	0.31*	48%
(3) + Eigenvalue clamping (λ=10^{-3})	0.76	+0.11	0.31	100%
(4) + Log-Euclidean metric	0.78	+0.13	0.35	100%
(5) + Riemannian Focal Loss	0.80	+0.15	0.38	100%
(6) + Learnable prototypes	0.81	+0.16	0.40	100%
(7) + Spectral dispersion (full ARG-SPD)	**0.82**	**+0.17**	**0.42**	100%

4.4 Cross-Site Generalization

Leave-one-site-out (LOSO) evaluation provides a controlled cross-site assessment within the dataset. Detroit→UCLA achieves channel-AUC/AUPRC/Sens@Sp90%/Pat-AUC of 0.79/0.36/0.62/0.78; UCLA→Detroit achieves 0.77/0.34/0.58/0.76. The mean LOSO channel-AUC of 0.78 represents a moderate decrease from 0.82 under 5-fold CV. This degradation reflects the impact of electrode heterogeneity and center-specific acquisition differences, and indicates that further domain adaptation or multi-site training strategies are needed for robust cross-center deployment.

4.5 Spatial Concordance with Resection Zones

For the 47 patients with available resection zone (RZ) annotations, we compute spatial concordance between predicted SOZ channels (threshold $\hat{p}(c)>0.5$) and the actual RZ. Mean Dice overlap is 78%±15%, with significantly higher overlap in Engel I (seizure-free) outcomes (85%±8%) versus Engel II–IV (65%±18%; p=0.02). This provides complementary evidence of clinical relevance beyond classification metrics. The 0.89 AUC reported above refers to the 78% high-confidence subset, not to excluding only the 7% most uncertain channels.

5 Discussion and Conclusion

The core finding is that preserving the Riemannian geometry of iEEG covariance matrices is beneficial for SOZ candidate detection: moving from a flat 1D-CNN to SPD representations yields the largest single AUC gain in our ablation (+0.11), consistent with prior work on geometric deep learning for brain signals [11,18]. The subsequent gains from the Riemannian Focal Loss, learnable prototypes, and spectral dispersion are complementary, each addressing a specific challenge of the SOZ detection problem: class imbalance, expressive class representation, and numerical conditioning. ARG-SPD addresses a task that is *distinct* from patient-level outcome prediction. Methods such as Partamian et al. [17] achieve 0.86 patient-AUC on outcome prediction, but they require expert SOZ labels as input. ARG-SPD estimates SOZ-likely channels directly from raw interictal iEEG without any SOZ annotations at inference time, addressing the preceding localization step that such methods assume to be solved. The spatial concordance results with resection zones and the outcome-stratified Dice analysis provide evidence that ARG-SPD predictions align with clinically meaningful surgical boundaries. The uncertainty-based triage mechanism is important for practical usability. By separating high-confidence SOZ candidates from ambiguous cases requiring expert review, ARG-SPD supports a human-in-the-loop workflow in which the model assists, rather than replaces, clinical decision-making. This is the appropriate deployment paradigm given current AI limitations in this domain.
Limitations. (i) Static covariance: The 15-second window assumes stationarity; dynamic SPD representations [5] could capture evolving ictal-interictal

transitions. **(ii) No anatomical priors:** Electrode coordinates are not incorporated in the architecture; the spatial concordance analysis (Sec. 4.5) is post-hoc and not used for training. **(iii) Single-dataset validation:** All results are on one public iEEG cohort. The Detroit↔UCLA LOSO experiment evaluates within-dataset cross-site generalization, but does not substitute for independent external validation on a separate cohort. Additional precision/recall, calibration, and sampling-ratio sensitivity analyses are important future work. **(iv) Computational cost:** Eigendecomposition scales as $\mathcal{O}(d^3)$, requiring approximately 50 hours of CPU training for 5-fold CV. Approximate SPD distances or lower-rank covariance parameterizations could reduce inference cost.

In conclusion, we present ARG-SPD, a geometric deep learning framework for automated channel-level SOZ candidate detection from raw inter-ictal iEEG. By combining SPD covariance representations with affine-invariant geodesic distances, learnable SPD class prototypes, a Riemannian Focal Loss, spectral dispersion regularization, and MC Dropout uncertainty quantification, ARG-SPD achieves $\mathbf{0.82 \pm 0.03}$ channel-AUC under strict patient-level evaluation on 184 patients, significantly outperforming all evaluated baselines. High-confidence predictions reach 0.89 AUC, and SOZ candidates show 78%±15% Dice spatial overlap against surgical resection zones. **Prospects of Application:** ARG-SPD is designed as a decision-support tool for presurgical epilepsy evaluation. It provides ranked SOZ-likely channel candidates and uncertainty estimates from inter-ictal iEEG to assist—not replace—clinical review, potentially reducing manual screening burden. Prospective multi-center validation is required before routine clinical use.

Disclosure of Interests. The authors declare that they have no competing interests relevant to the content of this article.

References

1. Aharonyan, A.A., Amir, S.A., Wittayanakorn, N., Linguraru, M.G., Oluigbo, C., Anwar, S.M.: Multi-modal graph-based machine learning for predicting surgical outcome in epilepsy patients. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. Lecture Notes in Computer Science, vol. 15971, pp. 500–509 (2025)
2. Bartolomei, F., Lagarde, S., Wendling, F., McGonigal, A., Jirsa, V., Guye, M., Bénar, C.: Defining epileptogenic networks: Contribution of SEEG and signal analysis. Epilepsia **58**(7), 1131–1147 (2017). https://doi.org/10.1111/epi.13791
3. Craley, J., Jouny, C., Johnson, E., Hsu, D., Ahmed, R., Venkataraman, A.: Automated seizure activity tracking and onset zone localization from scalp EEG using deep neural networks. PLOS ONE **17**(2), e0264537 (2022). https://doi.org/10.1371/journal.pone.0264537
4. Díaz-Montiel, A.A., Zhang, R., Lankarany, M.: Optimal graph representations and neural networks for multichannel time series data in seizure phase classification. Scientific Reports **15**(1), 19552 (2025). https://doi.org/10.1038/s41598-025-01882-7
5. Fan, M., Tang, F., Guo, Y., Zhao, X.: Riemannian dynamic generalized space quantization learning. Pattern Recognition **132**, 108932 (2022). https://doi.org/10.1016/j.patcog.2022.108932

6. Fan, X., Xu, P., Zhao, Q., Hao, C., Zhao, Z., Wang, Z.: A domain adaption approach for EEG-based automated seizure classification with temporal-spatial-spectral attention. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. Lecture Notes in Computer Science, vol. 15005, pp. 14–24 (2024). https://doi.org/10.1007/978-3-031-72086-4_2
7. Fedele, T., van 't Klooster, M., Burnos, S., Zweiphenning, W., van Klink, N., Leijten, F., Zijlmans, M., Sarnthein, J.: Automatic detection of high frequency oscillations during epilepsy surgery predicts seizure outcome. Clinical Neurophysiology **127**(9), 3066–3074 (2016). https://doi.org/10.1016/j.clinph.2016.06.009
8. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In: Proceedings of the 33rd International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 48, pp. 1050–1059 (2016)
9. Grattarola, D., Livi, L., Alippi, C., Wennberg, R., Valiante, T.A.: Seizure localisation with attention-based graph neural networks. Expert Systems with Applications **203**, 117330 (2022). https://doi.org/10.1016/j.eswa.2022.117330
10. Hassan, M., Mheich, A., Biraben, A., Merlet, I., Wendling, F.: Identification of epileptogenic networks from dense EEG: A model-based study. In: 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. pp. 5610–5613 (2015). https://doi.org/10.1109/EMBC.2015.7319664
11. Huang, Z., Van Gool, L.: A Riemannian network for SPD matrix learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 31 (2017)
12. Jacobs, J., Staba, R., Asano, E., Otsubo, H., Wu, J.Y., Zijlmans, M., Mohamed, I., Kahane, P., Dubeau, F., Navarro, V., Gotman, J.: High-frequency oscillations (HFOs) in clinical epilepsy. Progress in Neurobiology **98**(3), 302–315 (2012). https://doi.org/10.1016/j.pneurobio.2012.03.001
13. Kuroda, N., Sonoda, M., Miyakoshi, M., Nariai, H., Jeong, J.W., Motoi, H., Luat, A.F., Sood, S., Asano, E.: Objective interictal electrophysiology biomarkers optimize prediction of epilepsy surgery outcome. Brain Communications **3**(2), fcab042 (2021). https://doi.org/10.1093/braincomms/fcab042
14. Li, Z., Yeo, K., Gifford, W., Marcuse, L., Fields, M., Yener, B.: Spatio-temporal attention network for epileptic seizure prediction. arXiv preprint arXiv:2511.02846 (2025). https://doi.org/10.48550/arXiv.2511.02846
15. Li, Z., Li, W., Zhu, N., Han, J., Liu, T., Chen, B., Yan, Z., Zhang, T.: Epileptic seizure detection in SEEG signals using a unified multi-scale temporal-spatial-spectral transformer model. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. Lecture Notes in Computer Science, vol. 15011, pp. 254–264 (2024). https://doi.org/10.1007/978-3-031-72120-5_24
16. OpenNeuro: Open ieeg dataset (2024). https://doi.org/10.18112/openneuro.ds005398.v1.0.1, openNeuro dataset ds005398, Version 1.0.1
17. Partamian, H., Jahromi, S., Corona, L., Perry, M.S., Tamilia, E., Madsen, J.R., Bolton, J., Stone, S.S.D., Pearl, P.L., Papadelis, C.: Machine learning on interictal intracranial EEG predicts surgical outcome in drug resistant epilepsy. npj Digital Medicine **8**(1), 138 (2025). https://doi.org/10.1038/s41746-025-01531-3
18. Pennec, X., Fillard, P., Ayache, N.: A Riemannian framework for tensor computing. International Journal of Computer Vision **66**(1), 41–66 (2006). https://doi.org/10.1007/s11263-005-3222-z
19. Usman, S.M., Khalid, S., Aslam, M.H.: Epileptic seizures prediction using deep learning techniques. IEEE Access **8**, 39998–40007 (2020). https://doi.org/10.1109/ACCESS.2020.2976866