SA-CFM: Structure-Aware Continuous Flow Matching for Normative Modeling of EEG Microstate

Shiang Hu*, Yifan Hu, Zhuofan Shi, and Zhao Lv

Key Laboratory of Intelligent Computing and Signal Processing of Ministry of Education, School of Computer Science and Technology, Anhui University, Hefei, China
shu@ahu.edu.cn

Abstract. Normative modeling quantifies individual-level deviations from healthy population variability, yet its application to EEG remains limited by two fundamental gaps: existing approaches treat electrophysiological features as independent univariate variables, discarding structured inter-feature dependencies; and stochastic generative models such as Variational Autoencoders (VAEs) cannot guarantee a deterministic, invertible mapping from observations to latent space, causing the same subject to yield different deviation scores across inference runs, a critical flaw for clinical reproducibility. We propose Structure-Aware Continuous Flow Matching (SA-CFM), which addresses both gaps via physiologically grounded tokenization of EEG microstate features and a Transformer-based conditional flow matching objective. The deterministic invertibility of the learned flow yields unique, reproducible subject-level deviation representations without disease labels, a property that stochastic alternatives cannot provide by construction. Evaluated across two independent Parkinson's disease cohorts under strictly subject-independent protocols, SA-CFM achieves unsupervised cognitive subtype separation of ΔMoCA up to 3.96 and diagnostic accuracy of 83.0% and 88.6% on two independent cohorts, substantially exceeding both end-to-end EEG classifiers and normative baselines, establishing EEG microstate normative modeling as a viable complement to neuroimaging approaches for characterizing functional network alterations in neurodegeneration. Code is available at https://github.com/gjyradl/SA-CFM.

Keywords: Normative Modeling · Continuous Flow Matching · EEG Microstate · Generative Models · Neurodegenerative Diseases

1 Introduction

Parkinson's disease (PD) exhibits profound inter-individual variability in cognitive trajectories: individuals sharing identical motor phenotypes may range from intact cognition to frank dementia, yet cognitive impairment affects up to 80% of

* Corresponding author.

patients over the disease course [28]. Traditional case-control paradigms, by collapsing patients into diagnostic groups, systematically obscure this heterogeneity and offer little guidance for individualized prognosis.

Normative modeling addresses this by quantifying individual deviations from a population reference distribution conditioned on biological covariates [20, 22], yielding subject-specific atypicality scores without requiring disease labels. In structural and functional MRI, deviation-based clustering has revealed neurobiologically distinct subtypes with divergent cognitive decline trajectories [12], demonstrating that individual deviation patterns carry richer clinical information than group-level statistics. EEG provides millisecond-scale access to the same large-scale functional networks probed by fMRI, at lower acquisition cost [11].

EEG microstates are transient quasi-stable scalp topographies lasting 60–120 ms [9], providing a network-level representation of rapid whole-brain coordination complementary to hemodynamic measures, with temporal and transitional statistics robustly linked to cognitive dysfunction in PD [21, 16]. The closest existing work constructs normative models of spectral power and functional connectivity [25], yet treats features as independent univariate variables, discarding the structured multivariate dependencies among microstate statistics—precisely the inter-microstate coupling patterns encoding network-level coordination.

Beyond univariate modeling, a more fundamental challenge concerns the inferential requirements of normative deviation scoring [19]: reliable subject-level atypicality quantification demands a unique, deterministic mapping from each observation to the reference space. Stochastic generative models, including VAEs and diffusion models, cannot satisfy this by construction [15, 8, 29], as the posterior $q(z|x)$ is a distribution rather than a function, introducing irreducible sampling variance into deviation estimates unrelated to true atypicality. Continuous normalizing flows [2], by contrast, learn a deterministic diffeomorphism $\phi : \mathcal{X} \to \mathcal{Z}$, guaranteeing unique and fully reproducible inversion independent of data dimensionality.

We propose Structure-Aware Continuous Flow Matching (SA-CFM), a conditional invertible generative framework addressing both gaps. Our contributions are:

1. A flow matching framework with physiologically grounded tokenization that models microstate features as interacting structural units rather than flat vectors.
2. An empirical demonstration that deterministic flow inversion yields more stable and reproducible deviation representations than stochastic VAE-based normative modeling.
3. Validation on two independent PD cohorts, achieving superior unsupervised cognitive subtyping and diagnostic classification compared to both normative baselines and end-to-end classifiers.

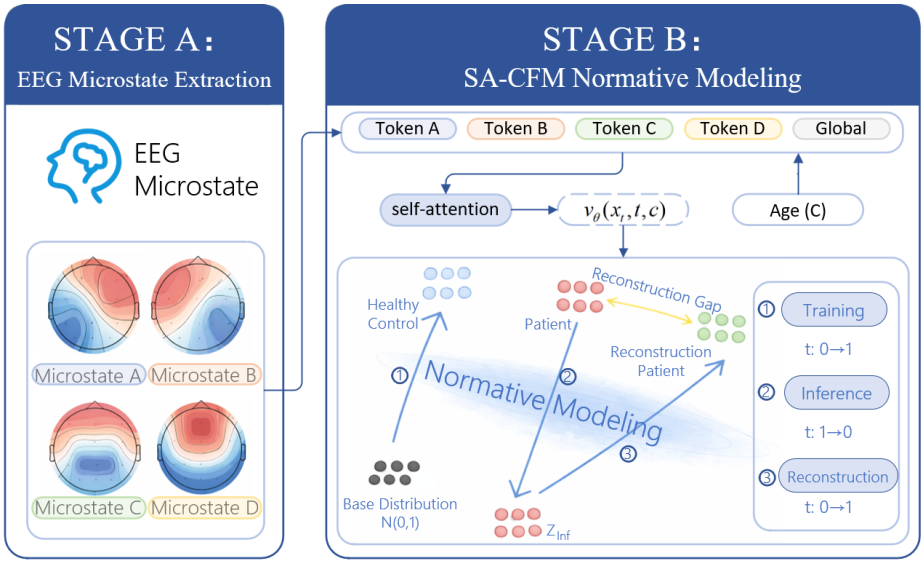

Fig. 1. Overview of SA-CFM. Stage A: Raw EEG is segmented into four canonical microstate topographies $\{A, B, C, D\}$, yielding a 27-dimensional feature vector x. Stage B: SA-CFM learns a normative flow conditioned on age c. The velocity field $v_\theta(x_t, t, c)$ is parameterized by a Transformer encoder that applies self-attention over microstate tokens, with age injected as a shared modulation signal. The flow supports training on healthy controls ①, deterministic inference to latent deviation z_{inf} ②, and normative reconstruction ③; the gap between patient and normative reconstruction quantifies atypicality.

2 Method

We propose SA-CFM, a normative modeling framework that learns a deterministic invertible mapping from the EEG microstate observation space to a standardized latent reference space, conditioned on subject age as the sole biological covariate (Fig. 1).

2.1 Physiologically Grounded Tokenization

We partition $x \in \mathbb{R}^{27}$ into five functional tokens $\{\tau_A, \tau_B, \tau_C, \tau_D, \tau_G\}$. Each map-specific token $\tau_k = [d_k, o_k, c_k, T_{k\to *}] \in \mathbb{R}^6$ $(k \in \{A, B, C, D\})$ groups duration, occurrence, coverage, and outgoing transition probabilities; the global token $\tau_G \in \mathbb{R}^3$ contains GEV, sequence entropy, and transition entropy. Each token is embedded as $h_i = f_{\text{emb}}^{(i)}(\tau_i) + p_i$, where p_i is a learnable type embedding, preserving inter-microstate coupling neglected by univariate approaches.

2.2 Continuous Flow Matching with Structured Velocity Field

We model the normative distribution via a continuous normalizing flow [17]:

$$\frac{dx_t}{dt} = v_\theta(x_t,\, t,\, c), \tag{1}$$

where v_θ is parameterized by a Transformer encoder operating on the token sequence $\{h_i\}$. Time t and age c are embedded through MLPs and added to each token as a shared modulation signal before self-attention:

$$\tilde{h}_i = h_i + g_t(t) + g_c(c). \tag{2}$$

The Transformer output tokens are concatenated and projected back to $\mathbb{R}^{27}$ to produce v_θ.

For training, we adopt Conditional Flow Matching with an Optimal Transport (OT) displacement interpolant [26]. Given healthy control samples $x_1 \sim p_{\mathrm{HC}}(x \mid c)$ and base noise $x_0 \sim \mathcal{N}(0, I)$, the interpolant is defined as:

$$x_t = (1-t)\, x_0 + t\, x_1, \quad t \sim \mathcal{U}[0,1], \tag{3}$$

with constant target velocity $u_t = x_1 - x_0$. The OT path minimizes transport cost and yields straight-line trajectories, enabling stable and efficient training. The flow matching objective is:

$$\mathcal{L}_{\mathrm{CFM}}(\theta) = \mathbb{E}_{t,\, x_0,\, x_1 \mid c}\Big[\big\| v_\theta(x_t, t, c) - (x_1 - x_0) \big\|^2\Big]. \tag{4}$$

2.3 Atypicality Quantification via Deterministic Inversion

After training on healthy controls, subject-specific deviations are obtained by integrating the learned ordinary differential equation (ODE) backward from $t = 1$ to $t = 0$:

$$z_{\mathrm{inf}} = \phi_{1\to 0}(x_{\mathrm{obs}}), \tag{5}$$

using a fixed-step RK4 solver with 50 integration steps.

3 Experiments

3.1 Datasets and Preprocessing

SA-CFM was trained on 1,419 healthy controls aggregated from five independent resting-state EEG datasets (Table 1), with 49 additional healthy subjects from PD-RS reserved for calibration. Two independently acquired PD datasets served as evaluation sets: PD-RS [1] and AEPD(NCT03808675), both providing cognitive annotations distinguishing cognitively normal (PD-CO-NO) from cognitively impaired (PD-CO-IM) subjects based on MoCA assessment [1]. All recordings underwent a unified pipeline: resampling to 128 Hz, 0.5–45 Hz bandpass filtering, and artifact removal via ADJUST-ICA [10]. Continuous EEG was

Table 1. Dataset summary. ⋆Used for unsupervised subtyping and supervised classification.

Purpose	Dataset	#Subj.	Age	Usage
Normative	BACA-RS [6]	607	20–70	Training
	CAUEEG [13]	459	23–88	Training
	MCEF-RS [3]	164	18–36	Training
	PEARL-Neuro [4]	78	50–63	Training
	SRM-RS [7]	111	17–71	Training
	PD-RS (HC)	49	52–86	Calibration
Evaluation	PD-RS (PD)	100	48–86	Stratification⋆
	AEPD	32	52–83	Stratification⋆

then segmented using a 4-second sliding window with 2-second overlap. EEG microstates were extracted via k-means clustering ($k = 4$) with 10 random restarts applied to GFP peak frames, yielding four canonical topographies $\{A, B, C, D\}$. Per-map duration, occurrence, coverage, and outgoing transition probabilities, together with global explained variance (GEV), sequence entropy, and transition entropy, form a 27-dimensional feature vector x. To mitigate inter-dataset heterogeneity, all datasets were processed with the same resampling, filtering, artifact-removal, GFP-peak selection, and microstate feature extraction pipeline before normative modeling.

3.2 Baselines and Metrics

We compare SA-CFM against two categories of baselines. Normative modeling baselines—all operating on the same 27-dimensional microstate feature vector x—include Generalized Additive Models for Location, Scale, and Shape (GAMLSS) [24] implemented with a Box–Cox t distribution for age-conditioned univariate normalization, and a conditional VAE (cVAE) [15] in which deviation scores are derived from the approximate posterior mean $\mu_\phi(x|c)$, with 10 Monte Carlo samples averaged to reduce but not eliminate sampling variance. All normative baselines share the same KAN-VAE downstream head for fair comparison. End-to-end EEG classifiers operate directly on raw EEG signals and include EEGNet [14], EEGInception [30], EEGConformer [23], Medformer [27], and MedGNN [5].

For unsupervised subtype discovery, we report Silhouette Score (S.S.), Davies-Bouldin Index (DBI), Spearman correlation ρ with MoCA scores (95% CI via 1,000 bootstrap resamples), MoCA gap between clusters (ΔMoCA), and clustering stability via Adjusted Rand Index (ARI) across seeds. For supervised diagnostic benchmarking, we report segment-level accuracy (Seg-Acc), and subject-level accuracy (Subj-Acc), where subject predictions are obtained by averaging softmax probabilities across segments.

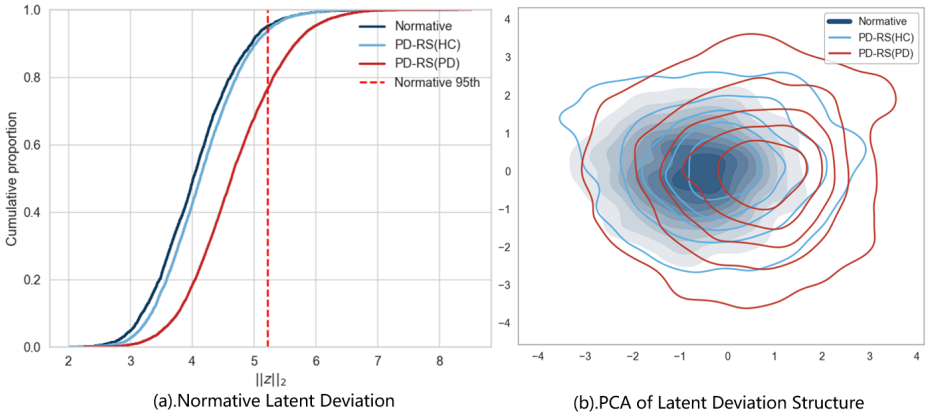

Fig. 2. Normative calibration and latent deviation structure. (a) Empirical CDF of $\|\mathbf{z}\|_2$ for the normative cohort, PD-RS(HC), and PD-RS(PD). Dashed line: normative 95th percentile. (b) 2D PCA projection of latent representations; normative dataset shown as density contours with PD groups overlaid.

3.3 Implementation Details

The SA-CFM velocity field uses a 3-layer Transformer encoder (4 heads, dimension 64) on five functional tokens, trained with AdamW ($\eta = 10^{-3}$, weight decay 10^{-5}) and cosine annealing over 200 epochs; ODE integration uses a fixed-step RK4 solver with 50 steps; increasing to 100 steps yields no measurable improvement in deviation scores. We adopt a KAN-VAE as the shared downstream head, where the VAE objective induces a structured latent space amenable to unsupervised subtype discovery, while the KAN layer provides expressive nonlinear feature transformation over the deviation representations [18]. This design unifies the clustering and classification pipelines under a single framework. The KAN-VAE downstream head is trained for unsupervised clustering (50 epochs) and supervised classification (up to 100 epochs with early stopping), with AdamW ($\eta = 5 \times 10^{-4}$) and KL annealing ($\beta_{\max} = 0.05$).

Unsupervised clustering was repeated with 10 random initializations for stability; supervised classification followed a strictly subject-independent 5-fold MCCV protocol with fixed seeds $\{5, 6, 7, 8, 9\}$.

4 Results and Discussion

4.1 Normative Calibration and Clinical Validity

As shown in Fig. 2(a), the normative cohort yields a smooth unimodal $\|z\|_2$ distribution, with PD-RS healthy controls overlapping substantially and PD patients shifting rightward beyond the 95th-percentile threshold. The PCA projection in Fig. 2(b) confirms that disease-related effects manifest as a structured

Table 2. Subtyping effectiveness across two PD datasets (KAN-VAE, k=2, 10 initializations). ρ: 95% CI via bootstrap; metrics as mean±std.

Method	S.S.	DBI	ρ (95% CI)	ΔMoCA	ARI
Dataset I: PD-RS					
Raw Features	0.31 ± 0.03	1.45 ± 0.08	0.13 (0.05, 0.28)	1.34 ± 0.64	0.12
GAMLSS	0.49 ± 0.02	0.72 ± 0.08	0.14 (0.04, 0.24)	0.81 ± 0.51	0.45
cVAE	0.38 ± 0.04	1.15 ± 0.07	0.28 (0.17, 0.39)	2.18 ± 0.42	0.33
SA-CFM (Ours)	**0.54 ± 0.02**	**0.63 ± 0.04**	**0.34 (0.16, 0.52)**	**3.32 ± 0.38**	**0.58**
Dataset II: AEPD					
Raw Features	0.33 ± 0.07	1.15 ± 0.05	0.21 (0.15, 0.44)	1.10 ± 0.13	0.08
GAMLSS	0.51 ± 0.03	0.65 ± 0.04	0.18 (0.02, 0.35)	0.82 ± 0.23	0.43
cVAE	0.35 ± 0.05	1.58 ± 0.09	0.24 (0.12, 0.43)	1.39 ± 0.31	0.39
SA-CFM (Ours)	**0.61 ± 0.01**	**0.48 ± 0.02**	**0.36 (0.20, 0.58)**	**3.96 ± 0.26**	**0.65**

Table 3. Diagnostic performance under 5-fold MCCV. Normative baselines share the same KAN-VAE head.

Method	PD-RS		AEPD	
	Seg-Acc	Subj-Acc	Seg-Acc	Subj-Acc
EEGNet	61.20±2.89	58.00±2.74	55.02±17.25	54.29±15.65
EEGInception	61.65±8.70	59.00±8.94	52.36±9.33	51.43±12.78
EEGConformer	62.94±5.02	63.00±7.58	57.04±13.22	57.14±10.10
Medformer	59.57±8.53	57.00±8.66	54.82±6.65	54.29±11.95
MedGNN	59.44±9.97	61.00±14.75	60.54±9.60	62.86±6.39
GAMLSS	60.98±5.45	60.00±7.07	62.12±7.23	60.62±7.52
cVAE	71.52±3.71	72.00±4.47	69.78±7.32	71.87±7.65
SA-CFM	**79.88±4.62**	**83.00±5.70**	**83.94±6.48**	**88.56±6.39**

outward shift from the normative density core, consistent with a continuous deviation framework.

4.2 Pathological Stratification and Diagnostic Utility

Table 2 evaluates unsupervised stratification of cognitive impairment within PD. SA-CFM consistently achieves the highest cluster quality (S.S.: 0.54±0.02/0.61± 0.01; DBI: 0.63/0.48) and the largest cognitive gap between discovered subtypes (ΔMoCA = 3.32 ± 0.38 on PD-RS; 3.96 ± 0.26 on AEPD), substantially exceeding all baselines. Raw features yield negligible separation ($\Delta\text{MoCA} \leq 1.34$) due to age-related confounds; GAMLSS fails despite reasonable cluster geometry as univariate normalization discards inter-feature interactions; cVAE improves but stochastic inference limits reproducibility (ARI: 0.33/0.39). SA-CFM's deterministic inversion yields the most stable stratification (ARI: 0.58/0.65) and strongest MoCA correlation ($\rho = 0.34/0.36$), capturing cognitively meaningful structure without any label supervision.

As shown in Table 3, SA-CFM achieves consistent diagnostic gains across both cohorts, outperforming all baselines. The near-chance performance of end-

Table 4. Ablation study. Segment-level accuracy (%) across five MCCV runs.

Configuration	Components			Performance	
	Attention	Token	KAN	PD-RS	AEPD
(1) Naive CFM				72.25±5.19	70.22±7.67
(2) Attention CFM	✓			76.51±6.12	72.19±8.01
(3) SA-CFM (w/o KAN)	✓	✓		78.12±4.72	81.56±8.16
(4) SA-CFM (Full)	✓	✓	✓	**79.88±4.62**	**83.94±6.48**

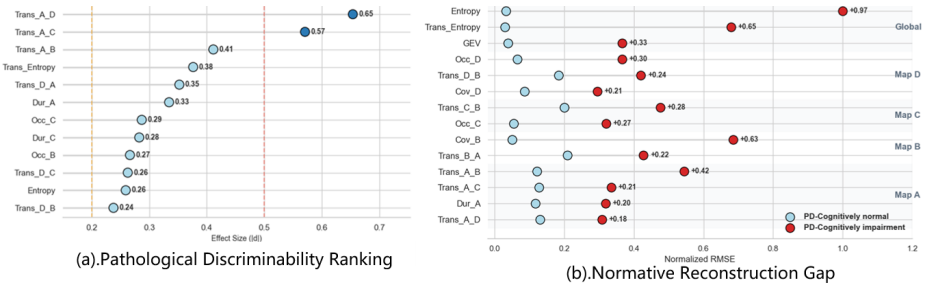

Fig. 3. Feature-level interpretability of normative deviations. (a) Cohen's $|d|$ between PD-CO-NO and PD-CO-IM; dashed lines mark small/medium effect-size thresholds ($|d| = 0.2/0.5$). (b) Per-feature reconstruction RMSE grouped by microstate map and global metrics; annotated gaps ($+\Delta$) denote excess error in PD-CO-IM.

to-end classifiers confirms that direct discriminative learning from EEG segments is ill-posed when inter-subject heterogeneity dominates intra-class structure; SA-CFM circumvents this by reframing classification as normative deviation quantification. Table 4 confirms the contribution of each component. Naive CFM incorporates neither self-attention nor physiologically grounded tokenization, serving as the most basic baseline. Adding self-attention over the feature sequence (Attention CFM) yields consistent improvements on both cohorts. Further introducing structured microstate tokenization (SA-CFM w/o KAN) brings additional gains, with the full SA-CFM incorporating all three components achieving the best performance across both datasets.

4.3 Feature-Level Interpretability of Normative Deviations

Fig. 3(a) identifies inter-state transitions from Map A as the most cognitively discriminative features: Trans_A_D ($|d| = 0.65$) and Trans_A_C ($|d| = 0.57$) substantially exceed the medium effect-size threshold, implicating selective failure of sensory-to-executive and sensory-to-DMN dynamic switching, consistent with the network disconnection hypothesis of PD cognitive impairment. Their moderate reconstruction error in Fig. 3(b) indicates that discriminability reflects systematic directional divergence between subgroups rather than global departure from the normative manifold. In contrast, global descriptors Entropy and Trans_Entropy exhibit the largest reconstruction gaps (+0.97 and +0.65)

yet only moderate discriminability, suggesting that entropic collapse constitutes a disease-general signature shared across both PD subgroups rather than a cognitive-specific marker. This dissociation between transition-specific discriminability and global entropic burden demonstrates that structure-aware normative deviations capture a richer profile of PD cognitive heterogeneity than global atypicality scores alone.

5 Conclusion

We present SA-CFM, a structure-aware continuous flow matching framework for normative modeling of EEG microstate dynamics. Physiologically grounded tokenization and deterministic flow inversion jointly address the two fundamental limitations of existing approaches: discarding inter-feature dependencies and inferential instability. Validated on two independent PD cohorts, SA-CFM substantially outperforms stochastic baselines in unsupervised cognitive subtyping and diagnostic classification, with deviation representations correlating with cognitive decline severity.

Acknowledgments. This work was supported by the National Natural Science Foundation of China (Grant No. 62101003), the Natural Science Foundation of Anhui Province (Grant No. 2508085MF156), and the Humanities and Social Sciences Project of the Ministry of Education of China (Grant No. 22YJC190007).

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

References

1. Anjum, M.F., Espinoza, A.I., Cole, R.C., Singh, A., May, P., Uc, E.Y., Dasgupta, S., Narayanan, N.S.: Resting-state eeg measures cognitive impairment in parkinson's disease. npj Parkinson's Disease **10**(1), 6 (2024)
2. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. Advances in neural information processing systems **31** (2018)
3. Chenot, Q., Hamery, C., Truninger, M., Langer, N., Scannella, S., et al.: Investigating the relationship between resting-state eeg microstates and executive functions: A null finding. cortex **178**, 1–17 (2024)
4. Dzianok, P., Kublik, E.: Pearl-neuro database: Eeg, fmri, health and lifestyle data of middle-aged people at risk of dementia. Scientific Data **11**(1), 276 (2024)
5. Fan, W., Fei, J., Guo, D., Yi, K., Song, X., Xiang, H., Ye, H., Li, M.: Towards multi-resolution spatiotemporal graph learning for medical time series classification. In: Proceedings of the ACM on Web Conference 2025. pp. 5054–5064 (2025)
6. Getzmann, S., Gajewski, P.D., Schneider, D., Wascher, E.: Resting-state eeg data before and after cognitive activity across the adult lifespan and a 5-year follow-up. Scientific Data **11**(1), 988 (2024)
7. Hatlestad-Hall, C., Rygvold, T.W., Andersson, S.: Bids-structured resting-state electroencephalography (eeg) data extracted from an experimental paradigm. Data in Brief **45**, 108647 (2022)

8. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
9. Hu, S., Karahan, E., Valdes-Sosa, P.A.: Restate the reference for eeg microstate analysis. arXiv preprint arXiv:1802.02701 (2018)
10. Hu, S., Ruan, J., Valdes-Sosa, P.A., Lv, Z.: How do the resting eeg preprocessing states affect the outcomes of postprocessing? NeuroImage **310**, 121122 (2025)
11. Hu, S., Xiang, X., Huang, X., Lu, Y., Zhang, X., Yao, D., Valdes-Sosa, P.A.: Lifespan brain age prediction based on multiple eeg oscillatory features and sparse group lasso. Frontiers in Aging Neuroscience **17**, 1559067 (2025)
12. Huo, Y., Jing, R., Li, P., Chen, P., Si, J., Liu, G., Liu, Y.: Delineating the heterogeneity of alzheimer's disease and mild cognitive impairment using normative models of dynamic brain functional networks. Biological psychiatry **97**(11), 1034–1044 (2025)
13. Kim, M.j., Youn, Y.C., Paik, J.: Deep learning-based eeg analysis to classify normal, mild cognitive impairment, and dementia: Algorithms and dataset. NeuroImage **272**, 120054 (2023)
14. Lawhern, V.J., Solon, A.J., Waytowich, N.R., Gordon, S.M., Hung, C.P., Lance, B.J.: Eegnet: a compact convolutional neural network for eeg-based brain–computer interfaces. Journal of neural engineering **15**(5), 056013 (2018)
15. Lawry Aguila, A., Chapman, J., Janahi, M., Altmann, A.: Conditional vaes for confound removal and normative modelling of neurodegenerative diseases. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 430–440. Springer (2022)
16. Li, Z., Zhu, X.: Application of eeg microstates in parkinson's disease. Parkinsonism & related disorders **136**, 107869 (2025)
17. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022)
18. Liu, Z., Wang, Y., Vaidya, S., Ruehle, F., Halverson, J., Soljacic, M., Hou, T., Tegmark, M.: Kan: Kolmogorov–arnold networks. In: International conference on learning representations. vol. 2025, pp. 70367–70413 (2025)
19. Maccioni, L., Giacomel, A., Pinamonti, M., De Francisci, M., Del Pup, F., Giubergia, A., Vallini, G., Visani, V., Barzon, L., Dipasquale, O., et al.: Quantitative neuroimaging meets normative modelling: the last mile for precision medicine applications. NeuroImage p. 121637 (2025)
20. Marquand, A.F., Rezek, I., Buitelaar, J., Beckmann, C.F.: Understanding heterogeneity in clinical cohorts using normative models: beyond case-control studies. Biological psychiatry **80**(7), 552–561 (2016)
21. Michel, C.M., Koenig, T.: Eeg microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: a review. Neuroimage **180**, 577–593 (2018)
22. Rutherford, S., Barkema, P., Tso, I.F., Sripada, C., Beckmann, C.F., Ruhe, H.G., Marquand, A.F.: Evidence for embracing normative modeling. Elife **12**, e85082 (2023)
23. Song, Y., Zheng, Q., Liu, B., Gao, X.: Eeg conformer: Convolutional transformer for eeg decoding and visualization. IEEE Transactions on Neural Systems and Rehabilitation Engineering **31**, 710–719 (2022)
24. Stasinopoulos, D.M., Rigby, R.A.: Generalized additive models for location scale and shape (gamlss) in r. Journal of Statistical Software **23**, 1–46 (2008)
25. Tabbal, J., Ebadi, A., Mheich, A., Kabbara, A., Güntekin, B., Yener, G., Paban, V., Gschwandtner, U., Fuhr, P., Verin, M., et al.: Characterizing the heterogeneity

of neurodegenerative diseases through eeg normative modeling. npj Parkinson's Disease **11**(1), 117 (2025)
26. Tong, A., Fatras, K., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks, J., Wolf, G., Bengio, Y.: Improving and generalizing flow-based generative models with minibatch optimal transport. arXiv preprint arXiv:2302.00482 (2023)
27. Wang, Y., Huang, N., Li, T., Yan, Y., Zhang, X.: Medformer: A multi-granularity patching transformer for medical time-series classification. Advances in Neural Information Processing Systems **37**, 36314–36341 (2024)
28. Wüllner, U., Borghammer, P., Choe, C.u., Csoti, I., Falkenburger, B., Gasser, T., Lingor, P., Riederer, P.: The heterogeneity of parkinson's disease. Journal of neural transmission **130**(6), 827–838 (2023)
29. Yang, X., Chin, B., Silosky, M., Litwiller, D., Ghosh, D., Xing, F.: Learning with synthesized data for generalizable lesion detection in real pet images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 116–126. Springer (2023)
30. Zhang, C., Kim, Y.K., Eskandarian, A.: Eeg-inception: an accurate and robust end-to-end neural network for eeg-based motor imagery classification. Journal of Neural Engineering **18**(4), 046014 (2021)