Bring Your Own Labels: Graph-Based Self-Supervised Learning for EEG in Computational Psychiatry

Udesh Habaraduwa[1],Julia Pottkämper[2], Marijn van Wingerden[1]

[1] Department of Cognitive Science & Artificial Intelligence, Tilburg University, Netherlands
[2] Psychiatry, Rijnstate Hospital, Arnhem
udesh86@gmail.com

Abstract. We test whether graph neural networks (GNNs) and self-supervision improve multi-class diagnosis on the TDBRAIN dataset by comparing a supervised graph model and SSL variants against EEGNet. Graph approaches raise macro-F1 from 0.24 to 0.42 on an imbalanced split and from 0.14 to 0.34 on a balanced subset; gains concentrate in minority classes, and class weighting outperforms undersampling. A well-tuned vanilla GNN performs on par with SSL-pretrained models and with recent reports, suggesting a strong inductive bias from electrode topology and band-power features. We release EEGLearn for reproducible processing and graph construction.

1 Motivation and Contribution

Psychiatric diagnosis lacks robust physiological markers. EEG is accessible but exhibits high inter-individual variability and a scarcity of ground-truth labels. GNNs are able to encode structure that grid CNNs ignore [4,3], and SSL can alleviate annotation bottlenecks [9]. In this work we introduce *EEGLearn*, a compact and reproducible pipeline for graph-EEG learning [1]; we conduct a controlled comparison between a vanilla GNN, SSL-based variants, and EEGNet [5]; and we demonstrate that simple graph priors combined with class weighting already achieve competitive macro-F1 on TDBRAIN [8].

2 Data and Pipeline

We evaluate on TDBRAIN [8], a resting-state EEG corpus (eyes open/closed) with 26 electrodes (10–10) spanning ADHD, MDD, OCD, SMC, and healthy controls. Two splits are used: The full imbalanced set (MDD 323, ADHD 175, SMC 119, OCD 49, healthy 47) and a balanced subset (45 per class). Our *EEGLearn* pipeline [1] standardizes cleaning and feature generation. We compare a supervised vanilla GNN, SSL pretraining (frequency permutation and spatial-region jigsaw [6]) with fine-tuning (FT), joint multi-task learning with uncertainty-weighted losses [2],

Table 1. 5-fold CV. Best macro-F1 in **bold**; 95% CIs in brackets. Prior work reports micro-F1 ≈ 0.46–0.55 [**?**] and macro-F1 ≈ 0.45 [7] on the same task.

Model/Study	Micro-F1	Macro-F1
	Dataset 1 (Imbalanced)	
EEGNet	0.38 [0.30, 0.47]	0.24 [0.18, 0.29]
SSL+FT	0.37 [0.33, 0.40]	0.29 [0.26, 0.32]
Jointly	**0.50 [0.46, 0.53]**	**0.42 [0.38, 0.46]**
Vanilla GNN	0.51 [0.47, 0.54]	0.38 [0.34, 0.41]
	Dataset 2 (Balanced)	
EEGNet	0.13 [0.03, 0.23]	0.14 [0.03, 0.25]
SSL+FT	0.28 [0.23, 0.34]	0.28 [0.23, 0.32]
Jointly	0.30 [0.25, 0.36]	0.29 [0.24, 0.34]
Vanilla GNN	**0.34 [0.29, 0.40]**	**0.34 [0.28, 0.40]**

and the EEGNet baseline [5]. Hyperparameters are tuned with Optuna; class imbalance is handled via weighted cross-entropy or undersampling; evaluation uses 5-fold cross-validation with micro- and macro-F1.
Context..

Key observations. (1) SSL pretexts diverge in difficulty: spatial jigsaw achieves high accuracy (0.78–0.89), while frequency permutation reaches only 0.40–0.57. (2) Graph models dominate EEGNet in macro-F1, especially for minority classes (e.g., healthy ≈ 0.55). (3) Class weighting works better than undersampling on average. (4) Vanilla GNN matches SSL+FT, implying strong inductive bias from scalp topology and band-power features [4,3].

3 Discussion and Outlook

SSL provided diagnostic insight but yielded limited downstream gains: spatial pretext task performance nears ceiling, while frequency pretext under-transfers. Train–validation gaps (GNN training F1 in the mid-0.9s) suggest overfitting. Regularization, larger unlabeled corpora, and EO/EC-specific encoders may help. Next steps include learned or connectivity-informed edges, contrastive SSL across epochs and conditions [9], and multimodal fusion. Overall, a simple graph pipeline with class-aware losses provides strong, reproducible baselines.

Code: *EEGLearn* (processing & datasets): [1].

References

1. Habaraduwa, U.: EEG-Graph-Learning. https://github.com/onedeeper/thesis/tree/main/eeg-graph-learning (Jun 2025), gitHub repository, accessed 7 June 2025
2. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7482–7491 (2018)
3. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)

4. Klepl, D., Wu, M., He, F.: Graph neural network-based eeg classification: A survey. IEEE Transactions on Neural Systems and Rehabilitation Engineering **32**, 493–503 (2024)
5. 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)
6. Li, Y., Chen, J., Li, F., Fu, B., Wu, H., Ji, Y., Zhou, Y., Niu, Y., Shi, G., Zheng, W.: Gmss: Graph-based multi-task self-supervised learning for eeg emotion recognition. IEEE Transactions on Affective Computing **14**(3), 2512–2525 (2022)
7. Smolders, T.: Self-Supervised Learning for EEG-Based Psychiatric Disorder Classification. Master's thesis, Tilburg University, Tilburg, The Netherlands (december 2024), thesis submitted for the degree of Master of Science in Data Science and Society
8. Van Dijk, H., Van Wingen, G., Denys, D., Olbrich, S., Van Ruth, R., Arns, M.: The two decades brainclinics research archive for insights in neurophysiology (tdbrain) database. Scientific data **9**(1), 333 (2022)
9. Weng, W., Gu, Y., Guo, S., Ma, Y., Yang, Z., Liu, Y., Chen, Y.: Self-supervised learning for electroencephalogram: A systematic survey. ACM Computing Surveys (2024)