ATLAS-FREE FUNCTIONAL BRAIN CONNECTOME ANALYSIS VIA TASK-DRIVEN PARCELLATION

Keqi Han[⋆]*, Yao Su*[†]*, Songlin Zhao*[§]*, Charles Gillespie*[⋆]*, Boadie Dunlop*[⋆]*, Daniel Barron*[¶]*,*
Randy Hirschtick[‡]*, Liang Zhan*[#]*, Lifang He*[§]*, Xiang Li*[‡]*, Carl Yang*[⋆]

[⋆] Emory University, Atlanta, GA, USA
[†] Worcester Polytechnic Institute, Worcester, MA, USA
[§] Lehigh University, Bethlehem, PA, USA
[¶] Brigham and Women's Hospital, Boston, MA, USA
[‡] Massachusetts General Hospital, Boston, MA, USA
[#] University of Pittsburgh, Pittsburgh, PA, USA

ABSTRACT

Functional brain connectome analysis models brain connectivity as a network, where nodes represent brain regions and edges capture their functional interactions. Existing methods mainly rely on predefined atlases (*e.g.*, AAL) to partition the brain and extract regional signals. However, these predefined atlases lack task specificity and cannot capture individual variability, limiting their effectiveness in clinical prediction. To address these challenges, we propose AFCON, an atlas-free framework that jointly optimizes brain parcellation and connectome analysis in a task-driven manner. AFCON adaptively generates individualized, task-specific parcellations directly from fMRI data, ensuring better alignment with downstream predictions and offering enhanced interpretability. Moreover, we introduce two neurobiologically informed regularizers to ensure plausible parcellations: a spatial compactness regularizer to promote anatomical coherence, and a balanced distribution regularizer to mitigate extreme parcel size imbalances. Experiments on ADHD and ADNI datasets show that AFCON achieves robust predictions, and importantly, identifies disease-relevant brain regions that reveal disrupted functional connections, improving both interpretability and clinical relevance. The code is available at https://github.com/LearningKeqi/AFCON.

***Index Terms*—** Functional Brain Connectome Analysis, Atlas-Free, Task-Driven Parcellation

1. INTRODUCTION

Functional connectome analysis models the brain as a network of regions of interest (ROIs) connected by statistical dependencies [1, 2, 3, 4]. Such network representations have been shown to capture meaningful neural patterns that predict cognitive function and characterize alterations linked to neurological and psychiatric disorders [5, 6].

The effectiveness of brain connectome analysis hinges on how brain regions are defined as network nodes. The common practice is to select one of many well-established brain atlases (*e.g.*, Harvard–Oxford [7], AAL [8]) for node definition. However, choosing a suitable atlas for a particular research question is non-trivial, especially given the wide heterogeneity of brain functions and cognitive processes. A mismatched atlas may blur critical functional distinctions, diminishing the predictive power of the resulting connectome structure. Besides, most existing atlases, either derived from group-level anatomical or functional data, are not tailored to specific downstream task. Such generic, group-level parcellations may overlook task-specific functional topographies and fail to capture individual variability, which is increasingly recognized as essential for understanding brain-behavior relationships. To mitigate these issues, recent studies have proposed atlas-initialized adaptive parcellation strategies [9], which initialize from a well-established atlas and progressively adjust the atlas boundaries during model training. Such adjustment remains highly biased toward the predefined atlas, hindering the discovery of novel task-specific functional organization. In contrast, a growing body of work has begun to explore atlas-free connectivity analysis [10, 11], which completely removes the dependence on any predefined atlas and instead learns brain parcellations directly from neuroimaging data. Nevertheless, most of them have primarily focused on structural MRI (sMRI) and diffusion tensor imaging (DTI), while fully functional MRI (fMRI)-based approaches remain underexplored.

In this work, we propose an end-to-end, **A**tlas-**F**ree functional brain **CON**nectome analysis (AFCON) framework driven by the downstream prediction task. Unlike methods that rely on predefined atlas, AFCON learns task-specific, individualized parcellations directly from fMRI data, which are optimized to enhance subsequent analyses. Specifically,

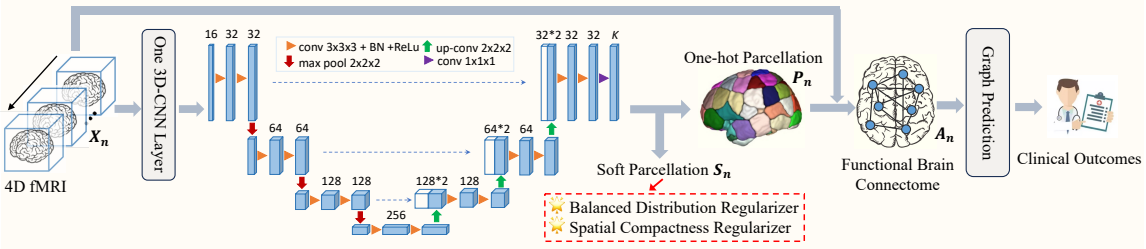

Fig. 1: The overall framework of AFCON. Each blue box corresponds to a multi-channel feature map. The number of channels is denoted on top of the box.

AFCON consists of two interconnected modules: an adaptive parcellation module, which segments the brain volume into non-overlapping ROIs, followed by a graph-based connectome analysis module, which constructs a brain graph based on the generated ROIs and makes predictions. By jointly optimizing both parcellation and downstream connectome analysis, the model adaptively discovers functionally meaningful regions that are closely aligned with the prediction task, thereby revealing potential biomarkers and enhancing interpretability. To further ensure neurobiologically plausible parcellations, we introduce two specialized regularizers grounded in established neuroscientific principles [12, 13]: a spatial compactness regularizer to promote anatomically coherent regions, and a balanced distribution regularizer to prevent extreme imbalances or domination by oversized ROIs. Experiments on ADHD [14] and ADNI [15] show that AFCON consistently achieves robust and competitive predictions while unveiling disease-relevant brain regions, enhancing both interpretability and clinical relevance.

2. METHOD

The overall framework of the proposed AFCON is illustrated in Fig. 1. Formally, the input is a set of N subjects, each represented by a four-dimensional fMRI volume $\mathbf{X}_n \in \mathbb{R}^{W\times H\times D\times T}$, where W, H, D denote the spatial dimensions and T represents the number of time points. This study focuses on the cerebral cortex by masking out non-cortical regions, so that only cortical signals remain in $\mathbf{X}_n$. The model outputs a clinical outcome $\hat{y}_n$ (*e.g.*, a brain disorder prediction). Besides, it produces a task-driven, subject-specific brain parcellation $\mathbf{P}_n$, which assigns each cortical voxel to one of K non-overlapping ROIs as an intermediate of the end-to-end pipeline. The generated parcellation highlights functionally relevant brain regions and offers insights into the underlying neurobiological processes.

Adaptive Brain Parcellation. Let $\mathbf{X}_n \in \mathbb{R}^{W\times H\times D\times T}$ denote the fMRI volume for subject n. Following the *time-as-channels* paradigm [16], we treat the temporal dimension as input channels and apply a 3D convolutional layer to jointly aggregate temporal dynamics while preserving the spatial topology of brain voxels. The extracted spatiotemporal features are then fed into a standard 3D U-Net [17], a widely used architecture for voxel-level segmentation in neuroimaging, whose encoder–decoder structure enables hierarchical integration of global and fine-grained functional patterns, facilitating the delineation of coherent brain regions. The 3D U-Net outputs voxel-level logits $\mathbf{L}_n \in \mathbb{R}^{W\times H\times D\times K}$, where K is the number of ROIs. To focus on the cerebral cortex, a binary mask $\mathbf{M} \in \{0,1\}^{W\times H\times D}$ is applied, retaining only cortical voxels as $\mathbf{L}_n^{cortex} \in \mathbb{R}^{Q\times K}$, where Q denotes the number of cortex voxels. The voxel-wise soft parcellation $\mathbf{S}_n \in \mathbb{R}^{Q\times K}$ is then obtained by applying a softmax operation along the ROI dimension, where each row of $\mathbf{S}_n$ sums to 1. Since downstream connectome analysis requires non-overlapping ROI labels, we use the Gumbel-Softmax to $\mathbf{L}_n^{cortex}$ to produce differentiable one-hot assignments $\mathbf{P}_n \in \{0,1\}^{Q\times K}$.

To ensure neurobiologically plausible parcellation, we introduce two specific regularizers grounded in established neuroscientific principles as follows:

Spatial Compactness Regularizer. Neuroimaging evidence indicates that functionally homogeneous brain areas are typically spatially contiguous [12], while scattered parcellations could compromise both biological fidelity and interpretability. We design a spatial compactness regularizer that promotes local geometric coherence within each ROI. Let $\mathbf{coord}_v \in \mathbb{R}^3$ denote the 3D coordinates of the v-th voxel. For each subject n and ROI k, we define a weighted centroid as $\mathbf{c}_{n,k} = \frac{\sum_{v=1}^{Q} \mathbf{coord}_v \cdot s_{n,k}(v)}{\sum_{v=1}^{Q} s_{n,k}(v)}$. The spatial variance of ROI k around its centroid is given by $\mathrm{Var}_{n,k} = \frac{\sum_{v=1}^{Q} s_{n,k}(v)\,\|\mathbf{coord}_v - \mathbf{c}_{n,k}\|^2}{\sum_{v=1}^{Q} s_{n,k}(v)}$. Summing across all subjects and ROIs yields the final compactness loss $\mathcal{L}_{\text{compact}} = \sum_{n=1}^{N}\sum_{k=1}^{K} \mathrm{Var}_{n,k}$. Minimizing $\mathcal{L}_{\text{compact}}$ favors spatially coherent ROIs, aligning with the biological principle that functionally specialized regions tend to be localized rather than dispersed.

Balanced Distribution Regularizer. Highly imbalanced ROI sizes can distort functional connectivity measures by conflating distinct neural signals within oversized parcels [13]. Such effects reduce sensitivity to region-specific interactions, and smaller ROIs may be overshadowed in subsequent network analyses. We therefore propose a balanced dis-

tribution regularizer to prevent extreme imbalances by oversized ROIs. Given the soft parcellation $\mathbf{S}_n$, we first calculate the fractional size of ROI k as $p_{n,k} = \frac{\sum_{v=1}^{Q} s_{n,k}(v)}{Q}$, where $s_{n,k}(v)$ denotes the probability of voxel v being assigned to ROI k for subject n. To discourage any single ROI from dominating, we penalize deviations from a uniform distribution $\mathbf{u} = \left[\frac{1}{K}, \ldots, \frac{1}{K}\right]$ by minimizing KL divergence-based loss $\mathcal{L}_{\text{balance}} = \sum_{n=1}^{N} \mathrm{KL}(\mathbf{p}_n \,\|\, \mathbf{u})$, where $\mathbf{p}_n = [p_{n,1}, \ldots, p_{n,K}]$. Note that this term is a soft constraint alongside the prediction and spatial compactness objectives, encouraging comparable ROI sizes *rather than enforcing strict uniformity.*

Graph-based Connectome Analysis. The parcellation $\mathbf{P}_n$ is used to compute ROI-level signals by averaging voxel time courses within each ROI. From these averaged signals, a functional correlation matrix $\mathbf{C}_n$ is computed via Pearson correlation among ROIs. Following BrainGNN [5], only the top 10% of positive values from $\mathbf{C}_n$ are retained, resulting in a sparse functional brain network matrix $\mathbf{A}_n \in \mathbb{R}^{K \times K}$. We choose Graph Convolutional Network (GCN) as the predictor due to its simplicity and widespread application in graph-based learning tasks. It takes $\mathbf{A}_n$ and the initial node features $\mathbf{H}_n$ (defined as the connection profile, *i.e.*, the rows of $\mathbf{C}_n$) to produce the final prediction: $\hat{y}_n = \mathrm{GCN}(\mathbf{A}_n, \mathbf{H}_n)$. A supervised cross-entropy prediction loss, $\mathcal{L}_{\text{pred}}$, is then generated to jointly train the GCN and the parcellation module.

Overall Training Process. In the early stages of training, under-informed brain parcellations lead to low-quality and unreliable connectivity for the graph predictor, creating a *cold start* that can hamper subsequent learning. To address this issue, we employ a two-stage training strategy: we first optimize only the parcellation module using $\mathcal{L}_{\text{balance}}$ and $\mathcal{L}_{\text{compact}}$. After E epochs of this warm-up phase, we unfreeze the GCN, introduce the prediction loss $\mathcal{L}_{\text{pred}}$, and jointly optimize $\mathcal{L} = \mathcal{L}_{\text{pred}} + \alpha\mathcal{L}_{\text{balance}} + \beta\mathcal{L}_{\text{compact}}$, where α and β are weights of two regularizers respectively.

3. EXPERIMENTS

Datasets. We use two publicly available resting-state fMRI (rs-fMRI) datasets, both preprocessed with fMRIPrep pipeline. (1) *ADHD-200* [14]. It includes 569 subjects with 246 ADHD patients (43.2%) and the rest as healthy controls (HC). Each subject's fMRI data is a 4D volume of size $97 \times 115 \times 97 \times 64$. (2) *ADNI* [15]. The original dataset is highly imbalanced. Following [18], we randomly sample 100 Alzheimer's disease (AD) patients and 100 healthy controls for evaluation. Each subject's fMRI data is a 4D volume of size $91 \times 109 \times 91 \times 197$.

Baselines. We compare AFCON with representative connectome analysis models, including GCN [19], GAT [20], specialized graph deep learning methods tailored for brain analysis: BrainGNN [5], BrainGB [2], BrainNetTF [1], NeuroGraph [3], and a CNN-based model BrainNetCNN [4]. We also include RefineNet [9], an atlas-initialized adaptive parcellation method. *Since baseline methods rely on predefined atlases and the optimal atlas choice may vary across datasets and models, we evaluate each baseline using all four widely adopted atlases: Harvard-Oxford 48 [7], AAL 90 [8], Schaefer 200 [21], and HCP 360 [22].* Table 1 reports the results of the best atlas configuration for each baseline.

Table 1: Overall Prediction Performance (mean±std, %). Gray denotes best; underline denotes second best.

Model	ADHD			ADNI		
	ACC↑	AUC↑	F1↑	ACC↑	AUC↑	F1↑
GCN	$59.7_{\pm 6.2}$	$63.2_{\pm 6.9}$	$48.3_{\pm 12.7}$	$60.5_{\pm 9.1}$	$65.7_{\pm 9.2}$	$63.4_{\pm 8.3}$
GAT	$57.7_{\pm 2.9}$	$60.3_{\pm 3.7}$	$53.6_{\pm 10.0}$	$56.0_{\pm 2.5}$	$59.4_{\pm 9.1}$	$55.4_{\pm 5.7}$
BrainGNN	$53.2_{\pm 3.8}$	$55.2_{\pm 3.7}$	$50.3_{\pm 7.0}$	$51.0_{\pm 5.4}$	$52.3_{\pm 6.3}$	$53.2_{\pm 5.5}$
BrainNetCNN	$56.0_{\pm 3.3}$	$58.7_{\pm 6.4}$	$52.1_{\pm 6.7}$	$58.5_{\pm 4.6}$	$65.9_{\pm 7.8}$	$53.6_{\pm 19.1}$
BrainGB	$56.7_{\pm 2.7}$	$58.3_{\pm 4.4}$	$46.0_{\pm 6.1}$	$56.5_{\pm 5.8}$	$59.7_{\pm 4.9}$	$58.2_{\pm 7.2}$
BrainNetTF	$59.8_{\pm 5.4}$	$\underline{63.8_{\pm 7.7}}$	$45.0_{\pm 22.8}$	$59.5_{\pm 5.3}$	$62.3_{\pm 3.8}$	$59.7_{\pm 8.6}$
NeuroGraph	$56.5_{\pm 5.4}$	$59.4_{\pm 4.3}$	$57.6_{\pm 4.7}$	$56.5_{\pm 8.5}$	$57.6_{\pm 8.2}$	$58.9_{\pm 11.1}$
RefineNet	$56.5_{\pm 4.2}$	$61.2_{\pm 3.4}$	$44.0_{\pm 9.7}$	$57.0_{\pm 7.1}$	$61.3_{\pm 7.5}$	$54.8_{\pm 23.6}$
AFCON-48	$63.2_{\pm 2.7}$	$65.6_{\pm 2.1}$	$\underline{56.8_{\pm 3.5}}$	$62.5_{\pm 6.5}$	$\underline{66.1_{\pm 7.3}}$	$\underline{62.6_{\pm 5.7}}$
AFCON-90	$60.0_{\pm 4.9}$	$63.5_{\pm 3.1}$	$50.7_{\pm 12.0}$	$61.5_{\pm 4.1}$	$65.6_{\pm 5.6}$	$61.5_{\pm 5.9}$
AFCON-200	$\underline{61.8_{\pm 0.9}}$	$62.9_{\pm 2.0}$	$47.9_{\pm 6.8}$	$\underline{62.0_{\pm 5.1}}$	$66.7_{\pm 4.2}$	$59.8_{\pm 7.5}$
AFCON-360	$61.1_{\pm 3.9}$	$63.4_{\pm 5.7}$	$49.8_{\pm 11.6}$	$59.5_{\pm 3.3}$	$65.1_{\pm 2.7}$	$55.1_{\pm 6.8}$

Implementation Details. Prediction tasks are binary classification evaluated by Accuracy, AUC, and F1. Data are split into 60%/20%/20% for training/validation/testing, and results are averaged over five runs. Models are trained for 100 epochs using Adam (batch size 4), selecting checkpoints by the best validation AUC. AFCON uses $K \in \{48, 90, 200, 360\}$ to align with well-established atlas scales. The 3D U-Net adopts a standard encoder–decoder (channels {32, 64, 128}), and the predictor is a 2-layer GCN (hidden size 32). The warm-up period $E = 20$ epochs. The regularizer weight α is tuned from $\{1, 2, 5\}$, β is tuned from $\{3e^{-3}, 5e^{-3}, 1e^{-2}, 2e^{-2}\}$. Hyperparameter tuning is also conducted by grid search for all baselines. Full configurations are available in the code repository.

Prediction Performance. As shown in Table 1, AFCON consistently achieves competitive or superior prediction performance compared to atlas-based baselines (both fixed and adaptively refined). In particular, AFCON with $K = 48$ yields the best overall performance on ADHD, achieving the highest accuracy and AUC while maintaining strong F1 scores. Similarly, on the ADNI dataset, AFCON achieves competitive performance, with $K = 48$ providing the best overall balance across ACC, AUC, and F1, and $K = 200$ yielding a strong AUC. These results highlight the potential of learning task-aware parcellations directly from fMRI data without relying on predefined atlases, allowing AFCON to better capture functionally relevant brain regions and construct more informative connectomes.

Quantitative Analysis of Learned Parcellation. To quantitatively evaluate AFCON's task-driven parcellation, we employ two standard metrics: *Region Homogeneity* and *Inter-subject Consistency* [23]. Region homogeneity is calculated as the average correlation between each voxel's time course and ROI's centroid time course, which is then averaged across all ROIs and test-set subjects. Inter-subject

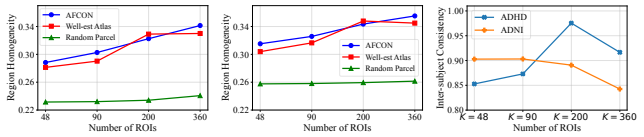

(a) Region Homogeneity on ADHD (b) Region Homogeneity on ADNI (c) Inter-subject Consistency

Fig. 2: Quantitative Analysis of Learned Parcellation

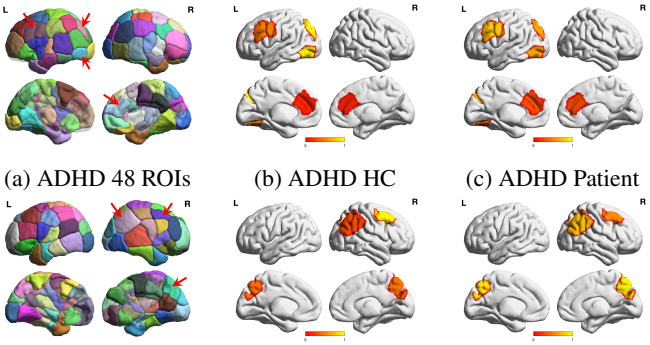

(a) ADHD 48 ROIs (b) ADHD HC (c) ADHD Patient

(d) ADNI 48 ROIs (e) ADNI HC (f) ADNI Patient

Fig. 3: Qualitative Analysis of Learned Parcellation. (a)(d) Task-driven parcellation obtained from ADHD and ADNI datasets, the red arrows indicate statistically significant ROIs. (b)(c) and (e)(f) The color of regions represents the significant ROIs' importance in the given group.

consistency is assessed by computing the Dice coefficient between every pair of parcellations in the test set and taking the average. We compare AFCON at multiple granularities with well-established atlases and random parcellations of the same resolution. As shown in Fig. 2(a)-(b), AFCON's parcellations achieve overall higher region homogeneity compared to both well-established atlases and random parcellations, reflecting more functionally cohesive ROIs. The strong inter-subject consistency shown in Fig. 2(c) indicates that AFCON not only accommodates individual variability but also captures a shared brain organization across subjects [23], which is crucial for robust group-level analysis.

Qualitative Analysis of Learned Parcellation. Based on Table 1, AFCON achieves the best overall performance on both datasets when number of ROIs is 48. We therefore focus on this configuration to identify disease-related biomarkers. We first generate a group-level parcellation from test-set subjects, by aggregating individual parcellations through voxel-wise majority voting, then construct each subject's functional connectome from the consensus partition. We measure each ROI's importance via node degree and conduct t-tests ($p < 0.05$) comparing HC and patients. For ADHD, Fig. 3(a)–(c) highlight significant group differences in ROIs located around Precentral, Occipital, Fusiform, and Cingulum areas, corroborating scientific findings of altered structure in these regions for ADHD [24]. For ADNI, Fig. 3(d)–(f) reveal that the Precuneus, Inferior Parietal, and Middle Frontal regions exhibit significant differences between HC and patient groups, consistent with literature linking disrupted connectivity in these brain regions to AD pathology [25].

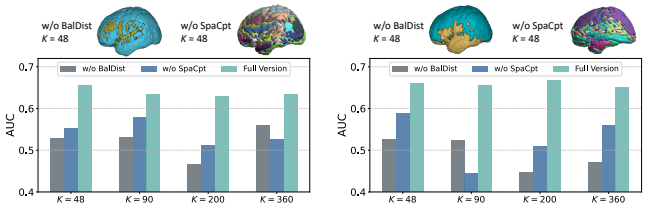

(a) ADHD (b) ADNI

Fig. 4: Ablation Study of the Proposed Regularizers.

Ablation Study. To validate the roles of proposed regularizers, we create two AFCON variants by removing either the balanced distribution regularizer (*w/o BalDist*) or the spatial compactness regularizer (*w/o SpaCpt*). As shown in Fig. 4, omitting any one of these regularizers leads to a notable drop in predictions. In addition, removing the balanced distribution regularizer causes the parcellation to be dominated by only a few oversized ROIs, whereas removing the spatial compactness regularizer produces fragmented, interspersed ROIs. Both outcomes deviate from the neuroscientific principles underlying feasible parcellations [12, 13], highlighting the importance of these regularizers in achieving meaningful parcellations for end-to-end brain connectome analysis.

4. CONCLUSION

We introduce AFCON, an atlas-free method for functional brain connectome analysis that bypasses the non-trivial atlas selection process. It achieves competitive predictions compared with atlas-based methods and reveals disease-relevant biomarkers in ADHD and AD cohorts, highlighting the potential for clinical diagnosis and treatment through task-specific brain parcellation. Future work will extend AFCON to subcortical regions and dynamic connectivity modeling.

5. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using de-identified human subject data made available in open access by ADHD and ADNI. Ethical approval was not required as confirmed by the license attached with the open access data.

6. ACKNOWLEDGMENTS

This work was partially supported by NSF (2319449, 2312502, IIS-2319451, MRI-2215789, IIS-2319450, IIS-2045848), NIH (R21AG087888, R01AG092661, 1R03AG078625-01, R01LM013519, RF1AG077820), DOE (DE-SC0025801), and Lehigh University (CORE, RIG).

7. REFERENCES

[1] X. Kan, W. Dai, H. Cui, Z. Zhang, Y. Guo, and C. Yang, "Brain network transformer," in *NeurIPS*, 2022.

[2] H. Cui, W. Dai, Y. Zhu, X. Kan, A. A. C. Gu, J. Lukemire, L. Zhan, L. He, Y. Guo, and C. Yang, "Braingb: A benchmark for brain network analysis with graph neural networks," *TMI*, 2023.

[3] A. Said, R. G. Bayrak, T. Derr, M. Shabbir, D. Moyer, C. Chang, and X. Koutsoukos, "Neurograph: benchmarks for graph machine learning in brain connectomics," in *NeurIPS*, 2023.

[4] J. Kawahara, C. Brown, S. Miller, B. Booth, V. Chau, R. Grunau, J. Zwicker, and G. Hamarneh, "Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment," *NeuroImage*, 2017.

[5] X. Li, Y. Zhou, N. Dvornek, M. Zhang, S. Gao, J. Zhuang, D. Scheinost, L. Staib *et al.*, "Braingnn: Interpretable brain graph neural network for fmri analysis," *Medical Image Analysis*, 2021.

[6] K. Han, Y. Su, L. He, L. Zhan, S. Plis, V. Calhoun, and C. Yang, "Rethinking functional brain connectome analysis: Do graph deep learning models help," *arXiv*, 2025.

[7] FSL-Team, "Fsl - fmrib software library," Available at: https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases, 2025.

[8] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot, "Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain," *Neuroimage*, 2002.

[9] N. Nandakumar, K. Manzoor, S. Agarwal, H. I. Sair, and A. Venkataraman, "Refinenet: an automated framework to generate task and subject-specific brain parcellations for resting-state fmri analysis," in *MICCAI*, 2022.

[10] M. Cole, K. Murray, E. St-Onge, B. Risk, J. Zhong, G. Schifitto, M. Descoteaux, and Z. Zhang, "Surface-based connectivity integration: An atlas-free approach to jointly study functional and structural connectivity," *Human Brain Mapping*, 2021.

[11] C. H. Gomez, L. Dodero, A. Gozzi, V. Murino, and D. Sona, "Atlas-free connectivity analysis driven by white matter structure," in *ISBI*, 2017.

[12] T. Blumensath, S. Jbabdi, M. F. Glasser, D. C. Van Essen, K. Ugurbil, T. E. Behrens, and S. M. Smith, "Spatially constrained hierarchical parcellation of the brain with resting-state fmri," *NeuroImage*, 2013.

[13] R. C. Craddock, G. James, P. E. Holtzheimer III, X. P. Hu, and H. S. Mayberg, "A whole brain fmri atlas generated via spatially constrained spectral clustering," *Human Brain Mapping*, 2012.

[14] ADHD-Consortium, "The adhd-200 consortium: a model to advance the translational potential of neuroimaging in clinical neuroscience," *Frontiers in systems neuroscience*, 2012.

[15] S. G. Mueller, M. W. Weiner, L. J. Thal, R. C. Petersen, C. Jack, W. Jagust, J. Q. Trojanowski, A. W. Toga, and L. Beckett, "The alzheimer's disease neuroimaging initiative," *Neuroimaging Clinics of North America*, 2005.

[16] X. Wang, X. Liang, Z. Jiang, B. A. Nguchu, Y. Zhou, Y. Wang, H. Wang, Y. Li, Y. Zhu, F. Wu *et al.*, "Decoding and mapping task states of the human brain via deep learning," *Human brain mapping*, 2020.

[17] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *MICCAI*, 2015.

[18] X. Kong, P. S. Yu, X. Wang, and A. B. Ragin, "Discriminative feature selection for uncertain graph classification," in *SDM*, 2013.

[19] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *ICLR*, 2017.

[20] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *ICLR*, 2018.

[21] A. Schaefer, R. Kong, E. Gordon, T. Laumann, X. Zuo, A. Holmes, S. Eickhoff, and T. T. Yeo, "Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri," *Cerebral Cortex*, 2017.

[22] M. F. Glasser, T. S. Coalson, E. C. Robinson, C. Hacker, J. Harwell, E. Yacoub, K. Ugurbil, J. Andersson, C. F. Beckmann, M. Jenkinson *et al.*, "A multi-modal parcellation of human cerebral cortex," *Nature*, 2016.

[23] P. Moghimi, A. T. Dang, Q. Do, T. I. Netoff, K. O. Lim, and G. Atluri, "Evaluation of functional mri-based human brain parcellation: a review," *Journal of neurophysiology*, 2022.

[24] M. Aricò, E. Arigliani, F. Giannotti, and M. Romani, "Adhd and adhd-related neural networks in benign epilepsy with centrotemporal spikes: A systematic review," *Epilepsy & Behavior*, 2020.

[25] H. Zhang, S. Wang, B. Liu, Z. Ma, M. Yang, Z. Zhang, and G. Teng, "Resting brain connectivity: changes during the progress of alzheimer disease," *Radiology*, 2010.