RefineNet: An Automated Framework to Generate Task and Subject-Specific Brain Parcellations for Resting-State fMRI Analysis

Naresh Nandakumar[1], Komal Manzoor[2], Shruti Agarwal[2], Haris I. Sair[2], and Archana Venkataraman[1]

[1] Dept. of Electrical and Computer Engineering, Johns Hopkins University, USA
[2] Dept. of Neuroradiology, Johns Hopkins School of Medicine, USA
nnandak1@jhu.edu

Abstract. Parcellations used in resting-state fMRI (rs-fMRI) analyses are derived from group-level information, and thus ignore both subject-level functional differences and the downstream task. In this paper, we introduce RefineNet, a Bayesian-inspired deep network architecture that adjusts region boundaries based on individual functional connectivity profiles. RefineNet uses an iterative voxel reassignment procedure that considers neighborhood information while balancing temporal coherence of the refined parcellation. We validate RefineNet on rs-fMRI data from three different datasets, each one geared towards a different predictive task: (1) cognitive fluid intelligence prediction using the HCP dataset (regression), (2) autism versus control diagnosis using the ABIDE II dataset (classification), and (3) language localization using an rs-fMRI brain tumor dataset (segmentation). We demonstrate that RefineNet improves the performance of existing deep networks from the literature on each of these tasks. We also show that RefineNet produces anatomically meaningful subject-level parcellations with higher temporal coherence.

Keywords: Rs-fMRI · Parcellation Refinement · Task optimization

1 Introduction

Resting-state fMRI (rs-fMRI) captures intrinsic neural synchrony, which provides insight into the functional organization of the brain [1]. Due to voxel-level variability, rs-fMRI data is often analyzed at the region level based on a predefined brain parcellation [2, 3]. Most parcellation schemes are based on group-level averages [4, 5]; however, it is well known that functional landmarks vary from person to person [6, 7], particularly for clinical populations [8].

Over the past decade, several methods have been proposed to obtain subject-specific functional boundaries from rs-fMRI data. The popular technique is independent component analysis (ICA), which estimates a set of spatially independent maps based on a linear decomposition of the voxel-wise rs-fMRI time series [9–12]. While ICA provides valuable subject-level information, the user cannot control the number or spatial compactness of the estimated components. For

this reason, it may not be possible to match region boundaries across subjects to draw group-level inferences. The more recent work of [13] uses a spatio-temporal deep residual network and ICA template priors to obtain subject-specific parcellations. While this method goes beyond traditional ICA to encode time-varying dynamics into network extraction, it has only been validated on a cohort of three subjects and might not generalize. The authors of [6] take a different approach by casting the problem as one of refining an existing parcellation, rather than creating a new parcellation from scratch. The proposed method iteratively reassigns the voxel memberships based on the maximum Pearson's correlation between the voxel time series and the mean time series of the current regions. The method in [8] builds upon this work by using a Markov random field prior to encourage spatial continuity during the iterative reassignment procedure. However, this method was only evaluated on a coarse initial parcellation (17 regions). Finally, the work of [7] uses a group sparsity prior as well as Markov Random Fields to generate subject-specific parcellations using an iterative graph partitioning approach. However, the authors do not show that the method improves performance on downstream tasks such as regression or classification.

In this paper we introduce RefineNet, the first deep learning approach for subject-specific and task-aware parcellation refinement using rs-fMRI data. RefineNet encodes both spatial and temporal information via a weight matrix that learns relationships between neighboring voxels and a coherence module that compares the voxel- and region-level time series. Importantly, RefineNet is designed as an all-purpose module that can be attached to existing neural networks to optimize task performance. We validate RefineNet on rs-fMRI data from three different datasets, each one designed to perform a different task: (1) cognitive fluid intelligence prediction (regression) on HCP [14], (2) autism spectrum disorder (ASD) versus neurotypical control (NC) classification on ABIDE [15], and (3) language localization using an rs-fMRI dataset of brain tumor patients. In each case, we attach RefineNet to an existing deep network from the literature designed for the given task. Overall, RefineNet improves the temporal cohesion of the learned region boundaries and the downstream task performance.

2 Methods

Fig. 1 illustrates our RefineNet strategy. The inputs to RefineNet are the 4D rs-fMRI data Z and the original brain parcellation $X^{(0)}$. We formulate a pseudo-prior, pseudo-likelihood and MAP style inference model to obtain the refined parcellation $X^{(e)}$. Following this procedure, RefineNet can be attached to an existing deep network to fine-tune $X^{(e)}$ for downstream task performance.

2.1 Spatial and Temporal Coherence Terms

Let V be the number of voxels in the rs-fMRI scan, and P be the number of regions in the original parcellation. We define $X \in \mathbb{R}^{V \times P}$ to RefineNet as a one-hot encoded label matrix, where $X_{v,p} = 1$ when voxel v is assigned to region

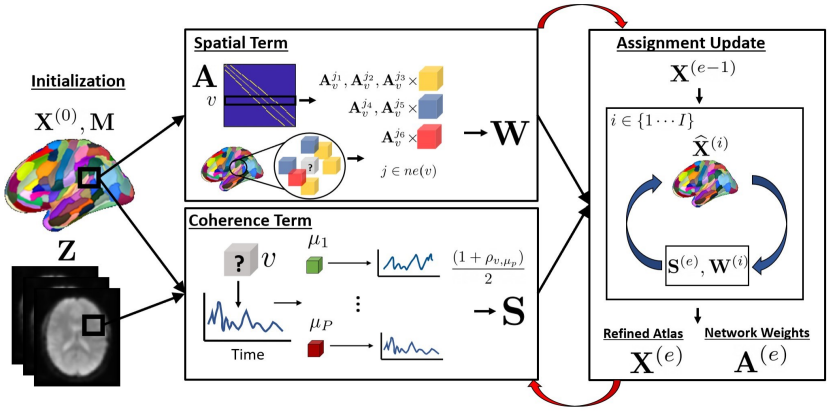

Fig. 1: Inputs: rs-fMRI Z, existing parcellation $X^{(0)}$ and neighbor mask M. Top: We show a six neighbor model for clarity. Our network parameter A learns voxel neighbor weights while the product (AX) multiplies these weights with the labels of the neighboring voxels. Bottom: The coherence term S uses the pearson correlation coeficient with each mean time series μ_p. Right: We obtain intermediate labels $\hat{X}^{(i)}$ I times before taking the mode and producing the next epoch's parcellation $X^{(e)}$, which is used during backpropagation to obtain $A^{(e)}$.

p and $X_{v,p} = 0$ otherwise. The core assumption of RefineNet is that voxels in close spatial proximity to each other are likely to belong to the same region [6, 8]. We encode this information via the intermediate activation $W \in R^{V \times P}$

$$W = \text{ReLU}(AX), \tag{1}$$

where the matrix $A \in R^{V \times V}$ enforces the local structure of the data. Formally, we obtain A as the Hadamard product of a sparse binary adjacency matrix $M \in R^{V \times V}$ that is nonzero only when the voxels are spatial neighbors and a learnable weight matrix $\hat{A} \in R^{V \times V}$ to encode spatially varying dependencies. Fig. 1 shows the nonzero weights in A_v being multiplied by the current labels of the neighbors of voxel v, where ne(v) denotes neighbors of voxel v.

At a high level, Eq. (1) acts as a proxy for the prior probability that voxel v belongs to region p based on the contribution of its neighbors currently assigned to region p, as governed by the spatially varying weights in A. Thus, our pseudo-prior term is designed to identify which neighbors are more important for voxel reassignment, which is important for boundary areas. Note that A is sparse by construction, which reduces both memory and computational overhead.

It is generally accepted that highly correlated voxels are more likely to be involved in similar functional processes, and if near each other, should be grouped into the same region [6, 7]. Let $Z \in R^{V \times T}$ denote the voxel-wise time series, where T is the duration of rs-fMRI scan. Thus, our pseudo-likelihood matrix $S \in R^{V \times P}$ that captures the un-normalized probability of voxel v being assigned to region

p is simply a shifted and scaled version of the Pearson's correlation coeficient between the voxel and mean region-wise time series, i.e., $S_{v,p} = \frac{\rho_{z_v,\mu_p}+1}{2}$.

Mathematically, given the voxel-to-region membership captured in X, we can compute the region-wise mean time series μ_p as follows:

$$\mu_p = \frac{\sum_{v} Z \cdot X_{v,p}}{\sum_{v=1}^{V} X_{v,p}}. \tag{2}$$

The correlation coeficient ρ_{z_v,μ_p} can also be obtained via matrix operations, allowing us to integrate the pseudo-likelihood term directly into a deep network.

2.2 RefineNet Training and Optimization

We adopt an iterative max product approach to derive our assignment updates. For convenience, let the index e denote the main epochs and the index i denote the refinement iterate. For each epoch e, we initialize the intermediate variable $\hat{X}^{(1)}$ with the assignment matrix $X^{(e-1)}$ from the previous iterate and compute the pseudo-likelihood matrix $S^{(e)}$ via the mean time series defined in Eq. (2). We then iteratively update $\hat{X}^{(i)}$ based on neighborhood information as follows:

$$\hat{X}^{(i+1)}_{v,p} = \begin{cases} 1 & \text{argmax}_p \; S^{(e)}_{v,:} \odot W^{(i)}_{v,:} \\ 0 & \text{else,} \end{cases} \tag{3}$$

where $W^{(i)} = \text{ReLU}\left(A\hat{X}^{(i)}\right)$ as in Eq. (1), and $\odot$ is the Hadamard product. The term $S^{(e)}$ remains constant throughout this iterative process from $i = \{1 \cdots I\}$ to act as the previous stationary point. The refined parcellation $X^{(e)}$ for epoch e is given by the majority vote over the intermediate region assignments $\{\hat{X}^{(i)}\}_{i=1}^{I}$. We employ this iterative approach over the pseudo-prior term to explore the space of intermediate label distributions for a robust assignment. We fix $I = 20$ in this work, as we empirically observed that this was large enough to provide robust reassignment.

We optimize the weights A in RefineNet via stochastic gradient descent to maximize the average temporal coherence with the newly assigned regions. Let V_p be the set of voxels assigned to region p. Our loss for backpropagation is

$$L^{RN} = -\frac{1}{V}\sum_{v=1}^{V} \frac{1}{|V_p|}\sum_{v \in V_p} \frac{(1+\rho_{z_v,\mu_p})}{2} \tag{4}$$

For clarity, our full training procedure is described in Algorithm 1.

2.3 Creating Task-Aware Parcellations with RefineNet

Once pretrained using Eq. (4), RefineNet can be attached to existing deep neural networks and re-optimized for performance on the downstream task. Our

Algorithm 1 RefineNet Training Procedure

```
procedure RefineNet(X, Z, M, E, I = 20)
    X^(0) ← X
    A^(0) ← Â, M                          ▷ Random initialization of weights in nonzero entries
    μ_1^(0) ··· μ_P^(0), S^(0) ← Z, X^(0)                                          ▷ Eq.(2)
    for e = 1 : E do
        X̂^(1) ← X^(e−1)
        for i = 1 : I do
            W^(i) ← A^(e−1), X̂^(i)                                                 ▷ Eq.(1)
            X̂^(i+1) ← S^(e−1), W^(i)                                               ▷ Eq.(3)
        X^(e) ← mode({X̂}_{i=1}^{I})
        μ_1^(e) ··· μ_P^(e), S^(e) ← Z, X^(e)                                      ▷ Eq.(2)
        L_RN ← X^(e), μ_1^(e) ··· μ_P^(e)                                          ▷ Eq.(4)
        A^(e) ← L_RN, SGD                 ▷ Backpropagation and gradient update
    return X^(E)
```

strategy is to pre-train RefineNet for 50 epochs using a learning rate of 0.001 before jointly training RefineNet with the network of interest. Here, we alternate between training just the network of interest for task performance and training both RefineNet and the network of interest in an end-to-end fashion. Empirically, we observed this strategy provides a good balance of task-optimization and preserving functional cohesion. Our second-stage loss function is a weighted sum of the downstream task and the RefineNet loss in Eq. (4):

$$L_{total} = L_{net} + \lambda L_{RN}, \tag{5}$$

where the hyperparameter λ can be chosen via a grid search or cross validation.

3 Experimental Results

We validate RefineNet on three different rs-fMRI datasets and prediction tasks. In each case, we select an existing deep network architecture from the literature to be combined with RefineNet. These networks take as input a $P \times P$ rs-fMRI correlation matrix. Fig. 2 illustrates the combined network architectures for each prediction task. We implement each network in Pytorch and use the hyperparameters and training strategy specified in the respective paper.

Our task-aware optimization (Section 2.3) alternates between by training the network of interest for e_a epochs while keeping RefineNet (and the input correlation matrices) fixed. We then jointly train both networks for e_a epochs while refining the parcellation, and thus, the correlation inputs between epochs.

3.1 Description of Networks and Data

M-GCN for Regression using HCP: We use the M-GCN model (rs-fMRI only) from [16] to predict the cognitive fluid intelligence score (CFIS). The

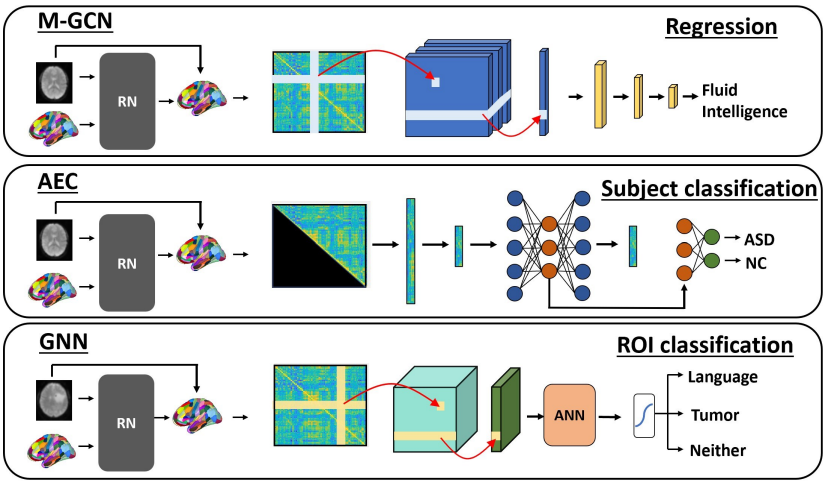

Fig. 2: Top: The M-GCN uses a graph convolution network applied to the connectivity matrix to predict fluid intelligence in HCP subjects. Middle: The AEC couples an autoencoder and a single layer perceptron to classify ASD vs. NC on ABIDE data. Bottom: The GNN uses graph convolutions to segment the language areas of eloquent cortex on a tumor dataset.

dataset contains 300 healthy subjects from the publicly available Human Connectome Project (HCP) S1200 release [14]. Standard rs-fMRI preprocessing was done according to [17], which handles motion, physiological artifacts, and registration to the MNI template. For simplicity, the CFIS values are scaled between (0 – 10) based on the training data of each fold. We report the mean absolute error (MAE) and correlation coeficient between the predicted and true scores.

AEC for Classification using ABIDE: We use the autoencoder/classifier (AEC) framework from [18] to predict subject diagnosis. The dataset contains 233 subjects (131 ASD, 102 NC) from the Autism Brain Imaging Data Exchange (ABIDE) II dataset [15]. The data was acquired across six different sites and preprocessed using the Configurable Pipeline for Analysis of Connectomes (CPAC) toolbox [19]. As per [18], the AEC network performs ASD vs. NC (neurotypical control) classification using the upper triangle portion of the rs-fMRI correlation matrix. We report the accuracy and area under the curve (AUC).

GNN for Language Localization in Tumor Patients: We use the GNN proposed by [20] to localize language areas of the brain in a lesional cohort. The dataset contains rs-fMRI and task fMRI data from 60 brain tumor patients. The data was acquired on a 3T Siemens Trio Tim system (EPI; TR = 2000 ms, TE = 30 ms, res = 1 mm^3 isotropic). The rs-fMRI was scrubbed using ArtRepair, followed by CompCorr for nuisance regression [21], bandpass filtering, and spatial smoothing. The task fMRI was used to derive “ground-truth" language labels for training and evaluation [22]. The tumor boundaries were obtained via expert

Table 1: Results across all experiments considered. Metric 1 represents MAE for regression and AUC for classification and localization while metric 2 represents correlation for regression and overall accuracy for classification and localization.

Task	Model	Atlas	Metric 1	Metric 2	P-value
CFIS Prediction	Original	BNA246	2.20 ± 0.13	0.24 ± 0.029	
		CC200	2.24 ± 0.14	0.27 ± 0.045	
		AAL90	2.22 ± 0.16	0.23 ± 0.048	
	RefineNet Only	BNA246	2.22 ± 0.13	0.19 ± 0.026	0.64
		CC200	2.22 ± 0.18	0.22 ± 0.036	0.293
		AAL90	2.15 ± 0.14	0.25 ± 0.032	0.121
	Combined	BNA246	1.73 ± 0.14	0.3 ± 0.039	0.016*
		CC200	1.84 ± 0.12	0.34 ± 0.046	0.045*
		AAL90	1.91 ± 0.11	0.36 ± 0.04	0.078*
ASD vs. NC	Original	BNA246	0.65 ± 0.017	65.5 ± 1.57	
		CC200	0.66 ± 0.024	64.9 ± 2.12	
		AAL90	0.66 ± 0.029	64.5 ± 2.49	
	RefineNet Only	BNA246	0.63 ± 0.021	63.8 ± 1.78	0.74
		CC200	0.69 ± 0.016	66.6 ± 1.80	0.22
		AAL90	0.70 ± 0.021	67.5 ± 1.94	0.08*
	Combined	BNA246	0.69 ± 0.013	67.8 ± 1.60	0.062*
		CC200	0.72 ± 0.029	69.8 ± 1.76	0.022*
		AAL90	0.74 ± 0.023	71.8 ± 1.84	0.006*
Localization	Original	BNA246	0.74 ± 0.022	84.6 ± 0.09	
		CC200	0.75 ± 0.021	85.9 ± 0.92	
		AAL90	0.67 ± 0.023	82.32 ± 1.21	
	RefineNet Only	BNA246	0.75 ± 0.023	84.95 ± 0.91	0.261
		CC200	0.75 ± 0.018	84.6 ± 0.71	0.531
		AAL90	0.65 ± 0.021	81.8 ± 1.34	0.834
	Combined	BNA246	0.77 ± 0.021	85.9 ± 0.91	0.065*
		CC200	0.78 ± 0.017	86.9 ± 1.01	0.047*
		AAL90	0.68 ± 0.019	82.63 ± 1.09	0.312

segmentation. The GNN outputs a label (language, tumor, or neither) for each region. We report the overall accuracy and AUC for detecting the language class.

3.2 Task Performance

We compare three model configurations: (1) no refinement (original), (2) using just RefineNet to maximize temporal coherence (RefineNet only), and integrating RefineNet into the network as described in Section 2.3 (combined). We also apply three parcellations to each task: the Brainnetome atlas (BNA246) [23], the Craddocks 200 atlas (CC200) [24], and the Automated Anatomical Labelling (AAL90) atlas [25]. To prevent data leakage, we tune the hyperparameters λ in Eq. (5) and alternating training epoch for regression on 100 additional HCP subjects, yielding $\lambda = 0.2$ and $e_a = 5$.

Table 1 reports the quantitative performance for each model/atlas configuration. Metrics 1/2 refer to MAE/correlation for the regression task and AUC/accurary for the classification and localization tasks, respectively. We employ a ten repeated 10-fold cross validation (CV) evaluation strategy to quantify per-

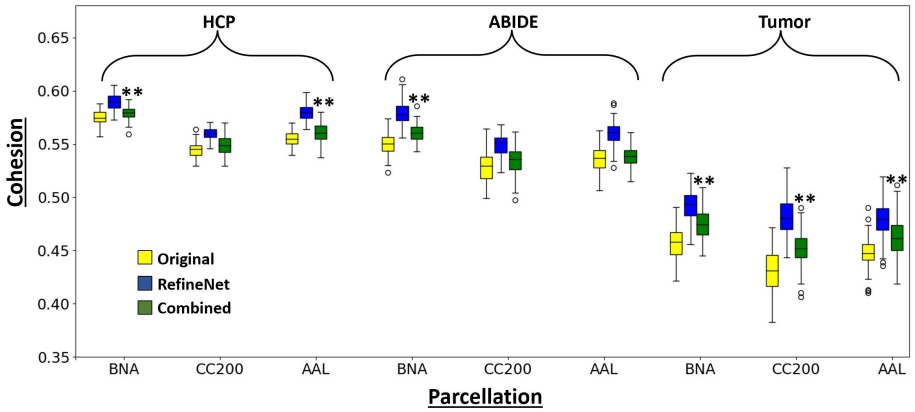

Fig. 3: Boxplots for region cohesion across the nine experiments. Yellow refers to the original model, blue refers to RefineNet only and green refers to combined. (**) denotes a significant increase from the original to combined parcellation.

formance variability. We report mean± standard deviation for each metric along with the FDR corrected p-value to indicate statistically improved performance in Metric 1 over the original model using the same parcellation [26]. As seen, the combined model provides statistically significant performance gains in eight out of nine experiments. In contrast, using RefineNet alone to strengthen func-tional coherence does not necessarily improve performance. Thus, our task-aware optimization procedure is crucial when considering downstream applications. Finally, we note that the AAL90 parcellation is likely too coarse for the language localization task, as reflected in the drastically lower performance metrics.

3.3 Parcellation Cohesion

Fig. 3 illustrates the average temporal cohesion of regions in the final parcellation, as computed on the testing data in each repeated CV fold. Once again, let μ_p denote the mean time series in each region p. We define the cohesion C as

$$C = \frac{1}{P}\sum_{v=1}^{V}\frac{1}{|V_p|}\sum_{v\in V_p}\rho_{z_v,\mu_p}. \quad (6)$$

Unsurprisingly, the parcellations recovered from just using RefineNet (with no downstream task awareness) achieve the highest cohesion. However, as shown in Table 1, these parcellations are not always suited to the prediction task. In contrast, the combined model produces more cohesive parcellations than the original atlas with statistically significant improvement denoted by (** **). Taken together, attaching RefineNet to an existing model achieves a good balance between functionally-cohesive grouping and task performance.

4 Conclusion

We present RefineNet, a flexible neural network module capable of obtaining meaningful subject-specific and task-aware parcellations. Our Bayesian-inspired approach considers both spatial contiguity and temporal coherence in reassignment. We show significant performance gains across three different datasets and prediction tasks when RefineNet is appended to existing networks from the literature. Finally, we show that even the task-driven refinement procedure produces more functionally cohesive parcellations than the origial atlas. Our work is a first of its kind, as other parcellation refinement methods are not able to be jointly trained with existing deep networks for task-awareness. In summary, our results show that RefineNet can be a promising tool for rs-fMRI analysis.

Acknowledgements: This work was supported by the National Science Foundation CAREER award 1845430 (PI: Venkataraman) and the Research & Education Foundation Carestream Health RSNA Research Scholar Grant RSCH1420.

References

1. B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde, "Functional connectivity in the motor cortex of resting human brain using echo-planar mri," Magnetic resonance in medicine, vol. 34, no. 4, pp. 537–541, 1995.
2. E. S. van Oort, M. Mennes, T. N. Schröder, V. J. Kumar, N. I. Z. Jimenez, W. Grodd, C. F. Doeller, and C. F. Beckmann, "Functional parcellation using time courses of instantaneous connectivity," Neuroimage, vol. 170, pp. 31–40, 2018.
3. M. Khosla, K. Jamison, A. Kuceyeski, and M. R. Sabuncu, "Ensemble learning with 3d convolutional neural networks for functional connectome-based prediction," NeuroImage, vol. 199, pp. 651–662, 2019.
4. B. Fischl, D. H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. Van Der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, et al., "Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain," Neuron, vol. 33, no. 3, pp. 341–355, 2002.
5. M. F. Glasser, T. S. Coalson, E. C. Robinson, C. D. 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, vol. 536, no. 7615, pp. 171–178, 2016.
6. D. Wang et al., "Parcellating cortical functional networks in individuals," Nature neuroscience, vol. 18, no. 12, p. 1853, 2015.
7. M. Chong, C. Bhushan, A. A. Joshi, S. Choi, J. P. Haldar, D. W. Shattuck, R. N. Spreng, and R. M. Leahy, "Individual parcellation of resting fmri with a group functional connectivity prior," NeuroImage, vol. 156, pp. 87–100, 2017.
8. N. Nandakumar, N. S. D'Souza, J. Craley, K. Manzoor, J. J. Pillai, S. K. Gujar, H. I. Sair, and A. Venkataraman, "Defining patient specific functional parcella-tions in lesional cohorts via markov random fields," in International Workshop on Connectomics in Neuroimaging, pp. 88–98, Springer, 2018.
9. F. Esposito, A. Aragri, I. Pesaresi, S. Cirillo, G. Tedeschi, E. Marciano, R. Goebel, and F. Di Salle, "Independent component model of the default-mode brain function: combining individual-level and population-level analyses in resting-state fmri," Magnetic resonance imaging, vol. 26, no. 7, pp. 905–913, 2008.

10. A. Tessitore, F. Esposito, C. Vitale, G. Santangelo, M. Amboni, A. Russo, D. Corbo, G. Cirillo, P. Barone, and G. Tedeschi, "Default-mode network connectivity in cognitively unimpaired patients with parkinson disease," Neurology, vol. 79, no. 23, pp. 2226–2232, 2012.
11. V. D. Calhoun and T. Adali, "Multisubject independent component analysis of fmri: a decade of intrinsic networks, default mode, and neurodiagnostic discovery," IEEE reviews in biomedical engineering, vol. 5, pp. 60–73, 2012.
12. H. I. Sair, N. Yahyavi-Firouz-Abadi, V. D. Calhoun, R. D. Airan, S. Agarwal, J. Intrapiromkul, A. S. Choe, S. K. Gujar, B. Caffo, M. A. Lindquist, et al., "Presurgical brain mapping of the language network in patients with brain tumors using resting-state f mri: Comparison with task f mri," Human brain mapping, vol. 37, no. 3, pp. 913–923, 2016.
13. B. Kazemivash and V. D. Calhoun, "A novel 5d brain parcellation approach based on spatio-temporal encoding of resting fmri data from deep residual learning," Journal of neuroscience methods, p. 109478, 2022.
14. D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugur-bil, W.-M. H. Consortium, et al., "The wu-minn human connectome project: an overview," Neuroimage, vol. 80, pp. 62–79, 2013.
15. A. Di Martino, D. O'connor, B. Chen, K. Alaerts, J. S. Anderson, M. Assaf, J. H. Balsters, L. Baxter, A. Beggiato, S. Bernaerts, et al., "Enhancing studies of the connectome in autism using the autism brain imaging data exchange ii," Scientific data, vol. 4, no. 1, pp. 1–15, 2017.
16. N. S. Dsouza, M. B. Nebel, D. Crocetti, J. Robinson, S. Mostofsky, and A. Venkataraman, "M-gcn: A multimodal graph convolutional network to inte-grate functional and structural connectomics data to predict multidimensional phenotypic characterizations," in Medical Imaging with Deep Learning, pp. 119–130, PMLR, 2021.
17. S. M. Smith, C. F. Beckmann, J. Andersson, E. J. Auerbach, J. Bijsterbosch, G. Douaud, E. Duff, D. A. Feinberg, L. Griffanti, M. P. Harms, et al., "Resting-state fmri in the human connectome project," Neuroimage, vol. 80, pp. 144–168, 2013.
18. J. Zhang, F. Feng, T. Han, X. Gong, and F. Duan, "Detection of autism spectrum disorder using fmri functional connectivity with feature selection and deep learning," Cognitive Computation, pp. 1–12, 2022.
19. C. Craddock, S. Sikka, B. Cheung, R. Khanuja, S. S. Ghosh, C. Yan, Q. Li, D. Lurie, J. Vogelstein, R. Burns, et al., "Towards automated analysis of connectomes: The configurable pipeline for the analysis of connectomes (c-pac)," Front Neuroinform, vol. 42, pp. 10–3389, 2013.
20. N. Nandakumar, K. Manzoor, J. J. Pillai, S. K. Gujar, H. I. Sair, and A. Venkataraman, "A novel graph neural network to localize eloquent cortex in brain tumor patients from resting-state fmri connectivity," in International Workshop on Connectomics in Neuroimaging, pp. 10–20, Springer, 2019.
21. Y. Behzadi, K. Restom, J. Liau, and T. T. Liu, "A component based noise correction method (compcor) for bold and perfusion based fmri," Neuroimage, vol. 37, no. 1, pp. 90–101, 2007.
22. W. D. Penny, K. J. Friston, J. T. Ashburner, S. J. Kiebel, and T. E. Nichols, Statistical parametric mapping: the analysis of functional brain images. Elsevier, 2011.
23. L. Fan, H. Li, J. Zhuo, Y. Zhang, J. Wang, L. Chen, Z. Yang, C. Chu, S. Xie, A. R. Laird, et al., "The human brainnetome atlas: a new brain atlas based on connectional architecture," Cerebral cortex, vol. 26, no. 8, pp. 3508–3526, 2016.

24. R. C. Craddock, G. A. 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, vol. 33, no. 8, pp. 1914–1928, 2012.
25. 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, vol. 15, no. 1, pp. 273–289, 2002.
26. R. R. Bouckaert and E. Frank, "Evaluating the replicability of significance tests for comparing learning algorithms," in Pacific-Asia Conference on Knowledge Discovery and Data Mining, pp. 3–12, Springer, 2004.