

GBT: Geometric-oriented Brain Transformer for Autism Diagnosis

Zhihao Peng¹, Zhibin He², Yu Jiang¹, Pengyu Wang¹, and
Yixuan Yuan¹

¹ The Chinese University of Hong Kong, Sha Tin, Hong Kong

² Northwestern Polytechnical University, Xi'an 710072, China
yxyuan@ee.cuhk.edu.hk

Abstract. Human brains are typically modeled as networks of Regions of Interest (ROI) to comprehend brain functional Magnetic Resonance Imaging (fMRI) connectome for Autism diagnosis. Recently, various deep neural network-based models have been developed to learn the representation of ROIs, achieving impressive performance improvements. However, they (i) heavily rely on increasingly complex network architecture with an obscure learning mechanism, or (ii) solely utilize the cross-entropy loss to supervise the training process, leading to sub-optimal performance. To this end, we propose a simple and effective Geometric-oriented Brain Transformer (GBT) with the Attention Weight Matrix Approximation (AWMA)-based transformer module and the geometric-oriented representation learning module for brain fMRI connectome analysis. Specifically, the AWMA-based transformer module selectively removes the components of the attention weight matrix with smaller singular values, aiming to learn the most relevant and representative graph representation. The geometric-oriented representation learning module imposes low-rank intra-class compactness and high-rank inter-class diversity constraints on learned representations to promote that to be discriminative. Experimental results on the ABIDE dataset validate that our method GBT consistently outperforms state-of-the-art approaches. The code is available at <https://github.com/CUHK-AIM-Group/GBT>.

Keywords: Autism diagnosis · Regions of Interest · Representation learning · Brain transformer network

1 Introduction

Autism is an ineradicable neurodevelopmental disability characterized by challenges with social skills and communication [7, 10, 18, 20, 30], where functional Magnetic Resonance Imaging (fMRI) is a powerful neuroimaging tool that depicts human brains as networks of Regions of Interest (ROIs) to enhance the interpretation and assessment of Autism diagnosis. In brain fMRI connectome analysis, some ROIs can co-activate or co-deactivate simultaneously when performing cognitive-related tasks such as action, language, and vision. Based on

this pattern, brain ROIs can be classified into diverse functional modules to analyze diseases towards their diagnosis, progress understanding, and treatment. In recent decades, due to the powerful representation capability of deep learning, a series of brain-aware specialized deep neural networks have been designed for brain fMRI connectome analysis. Early works designed Convolutional Neural Network (CNN)-based brain networks with specific convolutional filters to consider the locality structure of the brain [14, 16, 26]. Recently, Graph Neural Network (GNN)-based brain networks [11, 17, 22–24] and Transformer (TF)-based brain networks [12, 13, 32] have attracted widespread attention due to their powerful graph representation learning capability, which is well fit the non-euclidean properties of brain fMRI data.

Although numerous deep neural networks are deliberately well-designed to conduct effective representation learning for Autism diagnosis and have achieved competitive classification performance [1–3, 6, 21, 25, 28, 29], there remain several unsolved challenges with two main limitations. Firstly, existing works heavily rely on increasingly complex network architecture with an obscure learning mechanism to implicitly learn representations of ROIs, which may stick into over-fitting problems due to the introduced inductive biases resulting from sophisticated structures. Thus, it is expected that the network could emphatically focus on the principal component information of brain fMRI data, allowing us to study the most relevant and representative graph representation to reveal the interactions and functional organization between different brain regions. Secondly, they solely utilize the cross-entropy loss to conduct supervised learning of the encoder and the classifier, ignoring the presumable blurred boundaries issue within cross-boundary samples, leading to sub-optimal performance. Such a restriction inspires us to identify and compare brain function differences between different individuals to achieve distinguishable guidance toward embedding representation learning, promoting the learned representation to be discriminative.

To this end, we propose a novel Geometric-oriented Brain Transformer (GBT) with the Attention Weight Matrix Approximation (AWMA)-based transformer module and the geometric-oriented representation learning module for brain fMRI connectome analysis. Specifically, the AWMA-based transformer module focuses on learning the most relevant and representative graph representation of brain ROIs by selectively removing the components of the attention weight matrix with smaller singular values. Moreover, the geometric-oriented representation learning module imposes low-rank intra-class compactness and high-rank inter-class diversity constraints on the learned embedding representation to make it discriminative. Notably, such a discriminative representation is capable of fitting the natural geometric properties of brain data that subjects having the same disorder share similar brain network patterns, which means disorder-specific representations across instances are preferable. Experimental results on the ABIDE dataset demonstrate that the proposed method GBT consistently outperforms state-of-the-art (SOTA) approaches, e.g., GBT outperforms the second-best comparisons with a 6.00% improvement on ACC. In




Fig. 1. Illustration of our method GBT, which mainly includes the AWMA-based transformer module and the geometric-oriented representation learning module.

addition, we conduct a series of ablation studies on each component of GBT to verify its effectiveness. In summary, our contributions are as follows:

- We propose a novel AWMA-based transformer module to learn the most relevant and representative representation for brain fMRI connectome analysis.
- We propose a novel geometric-oriented representation learning module that combines low-rank intra-class compactness and high-rank inter-class diversity to enhance the discriminative capability of the representation.
- Both proposed modules are plug-and-play for the existing brain transformer network architecture.
- Experimental results on ABIDE validate the effectiveness of our method GBT in Autism diagnosis, e.g., GBT achieves a 6.00% ACC improvement than SOTA approaches.

2 Proposed Method

In this section, we introduce our network in detail, where the overall network architecture is shown in Fig. 1. Specifically, our network consists of the following parts: an AWMA-based transformer module to conduct the matrix approximation toward the attention weight matrix of the transformer encoder and a geometric-oriented representation learning module with low-rank intra-class compactness and high-rank inter-class diversity constraints to improve the embedding learning capability of our network.

2.1 Problem Definition

In brain fMRI connectome analysis, human brains are popularly depicted as networks of ROIs given an atlas, where the set of neural connections is defined as the Pearson correlation coefficient matrix $\mathbf{X} \in \mathbb{R}^{v \times v}$ of pairwise ROIs between the blood-oxygen-level-dependent (BOLD) signal series. v is the number of ROIs. To fit the natural topological properties of brain data, one of the most commonly used paradigms is to employ a vanilla transformer to obtain the graph representation $\mathbf{Z}_l$ with l being the number of layers. For $\mathbf{Z}_i$ on the i -th layer, three full-connected layers, parametrized by layer-specific weight matrices $\mathbf{W}^Q \in \mathbb{R}^{v \times d^Q}$, $\mathbf{W}^K \in \mathbb{R}^{v \times d^K}$, and $\mathbf{W}^V \in \mathbb{R}^{v \times d^V}$, are introduced to obtain the query matrix $\mathbf{Q}$, the key matrix $\mathbf{K}$, and the value matrix $\mathbf{V}$, respectively. After that, the LeakyReLU activation function $\mathcal{F}^{leak}$ [19] is applied on the multiplication between $\mathbf{Z}_{i-1}$ and the corresponding weight matrices. The softmax function $\mathcal{F}^{soft}$ is then utilized to obtain the attention weight matrix $\mathbf{W}_i^A$ of the transformer encoder. Afterward, the output is normalized by a multi-head self-attention module and a concatenation operation, denoted as $\mathcal{F}^{mcat}$.

The corresponding expressions are formulated as

$$\begin{aligned} \mathbf{Z}_i &= \mathcal{F}^{mcat}(\mathbf{W}_i^A \mathbf{V}) \mathbf{W}_i, \quad \text{s.t.} \quad \mathbf{W}_i^A = \mathcal{F}^{soft}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d^K}}\right), \quad \mathbf{Z}_0 = \mathbf{X}, \\ \mathbf{Q} &= \mathcal{F}^{leak}(\mathbf{Z}_{i-1} \mathbf{W}^Q), \mathbf{K} = \mathcal{F}^{leak}(\mathbf{Z}_{i-1} \mathbf{W}^K), \mathbf{V} = \mathcal{F}^{leak}(\mathbf{Z}_{i-1} \mathbf{W}^V). \end{aligned} \quad (1)$$

2.2 AWMA-based Transformer Module

To avoid the over-fitting problem, we propose an AWMA-based transformer module for developing a robust brain network to learn the most relevant and representative graph representation, as shown in Figure 1 (a). Specifically, we replace the attention weight matrix $\mathbf{W}_i^A$ of the transformer encoder at the i -th layer with its rank k approximation $\mathbf{W}_i^{LR}$ being

$$\mathbf{W}_i^{LR} = \mathbf{L}_k \mathbf{\Sigma}_k \mathbf{R}_k^\top, \quad \text{s.t.} \quad \mathbf{W}_i^A = \mathbf{L} \mathbf{\Sigma} \mathbf{R}^\top, \quad (2)$$

where $\mathbf{L}$, $\mathbf{L}_k$, $\mathbf{\Sigma}$, $\mathbf{\Sigma}_k$, $\mathbf{R}^\top$, and $\mathbf{R}_k$ is the left-singular matrix, the matrix formed by the first k columns of $\mathbf{L}$, a diagonal matrix containing the singular values, the k -th principal sub-matrix of $\mathbf{\Sigma}$, the conjugate transpose of the right-singular matrix, and the matrix formed by the first k columns of $\mathbf{R}$, respectively. It is well-known that Eckart-Young-Mirsky theorem [8] has proven singular value decomposition is the optimal solution of the matrix approximation [31] by removing the components with smaller singular values. By combining the Eqs. (1) and (2), the learned representation can be formulated as

$$\mathbf{Z}_i = \mathcal{F}^{mcat}(\mathbf{W}_i^{LR} \mathbf{V}) \mathbf{W}_i. \quad (3)$$




Fig. 2. Visualization of the skewed synthetic data with two classes. (a) Synthetic data involving cross-boundary samples. (b) Geometric-oriented representation with low-rank intra-class compactness and high-rank inter-class diversity. To avoid overfitting, we used the rank k approximation with SVD, the dropout regularization technique with dropping-out neurons, and the pooling strategy with aggregating local information.

2.3 Geometric-oriented Representation Learning Module

To enforce the learned representation to be discriminative, we design a geometric-oriented representation learning module with low-rank intra-class compactness and high-rank inter-class diversity constraints, named the geometric-oriented loss function. Specifically, we impose the low-rank constraint on the learned embedding representation within the same class, making the intra-class samples compact. Moreover, we encourage the whole features with inter-class diversity based on a high-rank constraint. Specifically, let $\mathbf{Z}^c$ denote the subspace matrix formed by the columns of $\mathbf{Z}_l$ that lies in the c -th class, the geometric-oriented loss function is as follows,

$$\min_{\mathbf{Z}_l} \sum_{i=1}^c \text{rank}(\mathbf{Z}_i) - \text{rank}(\mathbf{Z}_l), \quad (4)$$

where the rank function is inherently non-convex, making it computationally complex to optimize directly. Thus, we relax the rank function to its convex formulation, also known as (a.k.a.) the nuclear norm, i.e.,

$$\min_{\mathbf{Z}_l} \sum_{i=1}^c \|\mathbf{Z}_i\|_* - \|\mathbf{Z}_l\|_*, \quad (5)$$

where the nuclear norm of a matrix is defined as the sum of its singular values. More details of the relationship between the minimum-rank solution and the nuclear norm minimization can be found in the published theoretical justifications [9, 27]. By minimizing Eq. (5), the learned representation can simultaneously achieve intra-class compactness and inter-class diversity, explicitly satisfying the brain fMRI data property that these subjects with the same disorder share similar brain network patterns. Figure 2 shows via two illustrative examples based on a synthetic dataset the visualization of using the geometric-oriented

loss function as the objective function to qualitatively verify the methodology of the proposed geometric-oriented representation learning module. Finally, for the classification task, we first employ the MLP classifier $\mathcal{F}_\Theta$ to conduct supervised learning via the cross-entropy loss $\mathcal{F}^{ce}(\mathbf{Y}, \mathcal{F}_\Theta(\mathbf{Z}_l))$ with $\mathbf{Y}$ being the ground truth label, where the cross-entropy loss is the most widely used loss function for supervised learning of the transformer encoder and the classifier [1, 12]. Furthermore, we additionally exploit the proposed geometric-oriented loss Eq. (5) to supervise the whole training process, which can be written as

$$\min_{\mathbf{Z}_l} \sum_{i=1}^c |\mathbf{Z}_i|_* - |\mathbf{Z}_l|_* + \mathcal{F}^{ce}(\mathbf{Y}, \mathcal{F}_\Theta(\mathbf{Z}_l)). \quad (6)$$

3 Experiments

3.1 Datasets

Previous works [1, 12] commonly conduct the study of Autism on the open-source dataset Autism Brain Imaging Data Exchange (ABIDE), which is a collaborative initiative involving 17 international imaging sites [4]. Thus, we conduct experiments on ABIDE to understand the neural bases of Autism. Specifically, it aggregates and openly shares brain fMRI data of 1009 subjects where 516 individuals were diagnosed with Autism. The region definition is based on Craddock 200 atlas [5]. This dataset comprises structural and resting-state fMRI data and extensive phenotypic information, which are anonymous adhering to HIPAA guidelines and the 1000 Functional Connectomes Project/INDI protocols.

3.2 Compared Methods

We evaluate the effectiveness of our method by comparing it with one CNN-based brain network (BrainNetCNN [14]), three GNN-based brain networks (BrainGNN [17], FBNETGEN [11], BrainGB [6]), and three TF-based brain networks (Graphormer [32], BrainNetTF [12], and Com-BrainTF [1]).

3.3 Evaluation Metrics

We evaluate all the models on four commonly used evaluation metrics, i.e., the area under the receiver operating characteristic curve (AUC), accuracy (ACC), sensitivity (SEN), and specificity (SPE), where a higher value indicates a superior classification performance.

3.4 Training Procedure

For fair comparisons, we adopt [12] as our backbone network to conduct graph representation learning, where we employ a two-layer multi-head self-attention module with the number of attention heads being 4, the batch size being 64, and

Table 1. Performance comparisons, where the **bolded** and underlined values indicate the best and the second-best results, respectively.

Type	CNN-based	GNN-based			TF-based			
Method	BrainNetCNN [Neuroimage'17]	BrainGNN [MIA'21]	FBNETGEN [MIDL'22]	BrainGB [TMI'22]	Graphormer [NeurIPS'21]	BrainNetTF [NeurIPS'22]	Com-BrainTF [MICCAI'23]	Our
AUC	74.90±02.40	62.40±03.50	75.60±01.20	69.70±03.30	63.50±03.70	80.20±01.00	79.60±03.80	84.74±04.57
ACC	67.80±02.70	59.40±02.30	68.00±01.40	63.60±01.90	60.80±02.70	71.00±01.20	72.50±04.40	78.50±06.50
SEN	63.80±09.70	36.70±24.00	64.70±08.70	63.70±08.30	78.70±22.30	72.50±05.20	80.10±05.80	80.21±09.38
SPE	<u>71.00±10.20</u>	70.70±19.30	62.40±09.20	60.40±10.10	36.70±23.50	69.30±06.50	65.70±06.40	76.92±03.85

the epoch number being 200. We use Adam [15] with an initial learning rate of $1e-4$ and a weight decay of $1e-4$. For the datasets, we randomly split 70% for training, 10% for validation, and 20% for testing, which is the same as the splits/evaluation strategy as the SOTA ones [1, 12]. We perform the experiments five times and report the average results with their standard deviations (i.e., mean $\pm$ std). For compared approaches, we directly report the performances provided in the original papers [12] and [1]. The model is implemented with PyTorch on NVIDIA GeForce RTX 4090.

3.5 Compared Results

Table 1 shows the comparisons between the proposed method and seven compared approaches with four metrics on the ABIDE dataset, where we have the following observations: (i) Our method GBT achieves almost the best performance on all metrics and outperforms the SOTA approaches. For example, GBT achieves a 6.00% improvement on ACC compared with the second-best method. (ii) The TF-based brain networks perform significantly better than the CNN-based and GNN-based brain networks on almost all the metrics. For instance, on SEN, the critical metric for diagnostic tests referring to true positive rate, our TF-based method outperforms the best GNN-based network FBNETGEN [11] and CNN-based network BrainNetCNN [14] by 15.51% and 16.41%, respectively. (iii) Compared with FBNETGEN, which uses group losses to extract GNN features, our main innovation is utilizing global rank-aware constraints to impose intra-class and inter-class rank constraints into the brain-aware transformer, effectively exploring the natural geometric properties of brain data and making a 14.52% SPE improvement. (iii) Our method GBT achieves better performance than the baseline BrainNetTF [12], demonstrating the effectiveness of the AWMA-based transformer module and the proposed geometric-oriented representation learning module. Concretely, GBT improves 4.54% on AUC, 7.50% on ACC, 7.71% on SEN, and 7.62% on SPE.

3.6 Ablation Study

We conduct comprehensive ablation studies, where the experimental results are listed in Table 2. Specifically, the first row (**I**) denotes the model [12], i.e., our backbone network, with the cross-entropy loss. The second row (**II**) denotes our

Table 2. Ablation studies of the proposed geometric-oriented representation learning module (i.e., low-rank and high-rank constraints) and the AWMA-based transformer module, where $\times$ and $\checkmark$ in each row indicate the non-use and use of the corresponding component, respectively. The best results are highlighted with **bold**.

	Low-rank	High-rank	AWMA	AUC	ACC	SEN	SPE
(I)	$\times$	$\times$	$\times$	80.20 $\pm$ 01.00	71.00 $\pm$ 01.20	72.50 $\pm$ 05.20	69.30 $\pm$ 06.50
(II)	$\times$	$\times$	$\checkmark$	81.19 $\pm$ 02.33	76.20 $\pm$ 00.98	77.79 $\pm$ 11.29	72.40 $\pm$ 09.17
(III)	$\times$	$\checkmark$	$\checkmark$	79.79 $\pm$ 02.63	71.60 $\pm$ 03.20	76.91 $\pm$ 10.48	66.82 $\pm$ 14.89
(IV)	$\checkmark$	$\times$	$\checkmark$	80.30 $\pm$ 03.13	76.25 $\pm$ 02.17	78.75 $\pm$ 12.95	71.17 $\pm$ 13.54
Our	$\checkmark$	$\checkmark$	$\checkmark$	84.74$\pm$04.57	78.50$\pm$06.50	80.21$\pm$09.38	76.92$\pm$03.85

backbone network with the AWMA-based transformer module and the cross-entropy loss. The third row (**III**) denotes a variant of our network that conducts supervised learning with the second-term high-rank loss and the cross-entropy loss, i.e., $\min_{\mathbf{Z}_l} (-|\mathbf{Z}_l|_* + \mathcal{F}^{ce}(\mathbf{Y}, \mathcal{F}_\Theta(\mathbf{Z}_l)))$. The fourth row (**IV**) denotes a variant of our network that utilizes the first-term low-rank loss and the cross-entropy loss to supervise the training process, i.e., $\min_{\mathbf{Z}_l} \sum_{i=1}^c |\mathbf{Z}_i|_* + \mathcal{F}^{ce}(\mathbf{Y}, \mathcal{F}_\Theta(\mathbf{Z}_l))$. The fifth row (**Our**) is our full method. From Table 2, we have the following observations: (i) The comparisons between (**I**) and (**II**) illustrate the effectiveness of the proposed AWMA-based transformer module among all metrics. For example, it achieves a 5.29% improvement on the SEN metric. (ii) By comparing the results in (**II**) and (**III**), we can observe that the utilization of the high-rank constraint has performance degradation. The possible reason is that the samples lack intra-class compactness, resulting in reduced discriminative power between different nodes. (iii) The comparisons of **Our** method with (**III**) and (**IV**) demonstrate that simultaneously considering the low-rank intra-class compactness and high-rank inter-class diversity could promote learning a discriminative representation, improving the classification performance. For example, **Our** method obtains 10.10 % performance improvement over **III** on SPE.

4 Conclusion

This paper proposes a novel transformer-based brain network, GBT, to learn discriminative graph representations across brain ROIs. Specifically, GBT combines an AWMA-based transformer module to conduct the matrix approximation toward the attention weight matrix of the transformer encoder and a geometric-oriented representation learning module to consider the intra-class compactness and inter-class diversity, aiming to understand brain fMRI connectome for Autism diagnosis. Experimental results on ABIDE demonstrate its superiority over SOTA methods, validating the effectiveness of the proposed AWMA-based transformer module and geometric-oriented representation learning module. In the future, we will exploit the Neurosynth platform to conduct neuroimaging meta-analysis to promote biomarker discovery.

Acknowledgments. This work was supported by Hong Kong Research Grants Council (RGC) General Research Fund 14204321.

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

References

1. Bannadabhavi, A., Lee, S., Deng, W., Ying, R., Li, X.: Community-aware transformer for autism prediction in fmri connectome. In: Medical Image Computing and Computer-Assisted Intervention. pp. 287–297. Springer (2023)
2. Bessadok, A., Mahjoub, M.A., Rekik, I.: Symmetric dual adversarial connectomic domain alignment for predicting isomorphic brain graph from a baseline graph. In: Medical Image Computing and Computer-Assisted Intervention. pp. 465–474. Springer (2019)
3. Cai, M., Li, M., Xiong, Z., Zhao, P., Li, E., Tang, J.: An advanced deep learning framework for video-based diagnosis of asd. In: Medical Image Computing and Computer-Assisted Intervention. pp. 434–444 (2022)
4. Craddock, C., Benhajali, Y., Chu, C., Chouinard, F., Evans, A., Jakab, A., Khundrakpam, B.S., Lewis, J.D., Li, Q., Milham, M., et al.: The neuro bureau preprocessing initiative: open sharing of preprocessed neuroimaging data and derivatives. *Frontiers in Neuroinformatics* **7**(27), 5 (2013)
5. Craddock, R.C., James, G.A., Holtzheimer III, P.E., Hu, X.P., Mayberg, H.S.: A whole brain fmri atlas generated via spatially constrained spectral clustering. *Human Brain Mapping* **33**(8), 1914–1928 (2012)
6. Cui, H., Dai, W., Zhu, Y., Kan, X., Gu, A.A.C., Lukemire, J., Zhan, L., He, L., Guo, Y., Yang, C.: Braingb: A benchmark for brain network analysis with graph neural networks. *IEEE Transactions on Medical Imaging* **42**(2), 493–506 (2022)
7. D’Souza, N.S., Nebel, M.B., Crocetti, D., Wymbs, N., Robinson, J., Mostofsky, S., Venkataraman, A.: A deep-generative hybrid model to integrate multimodal and dynamic connectivity for predicting spectrum-level deficits in autism. In: Medical Image Computing and Computer-Assisted Intervention. pp. 437–447. Springer (2020)
8. Eckart, C., Young, G.: The approximation of one matrix by another of lower rank. *Psychometrika* **1**(3), 211–218 (1936)
9. Jain, P., Meka, R., Dhillon, I.: Guaranteed rank minimization via singular value projection. *Advances in Neural Information Processing Systems* **23** (2010)
10. Jones, W., Klin, A.: Attention to eyes is present but in decline in 2–6-month-old infants later diagnosed with autism. *Nature* **504**(7480), 427–431 (2013)
11. Kan, X., Cui, H., Lukemire, J., Guo, Y., Yang, C.: Fbnetgen: Task-aware gnn-based fmri analysis via functional brain network generation. In: International Conference on Medical Imaging with Deep Learning. pp. 618–637. PMLR (2022)
12. Kan, X., Dai, W., Cui, H., Zhang, Z., Guo, Y., Yang, C.: Brain network transformer. *Advances in Neural Information Processing Systems* **35**, 25586–25599 (2022)
13. Kang, E., Heo, D.W., Suk, H.I.: Prototype learning of inter-network connectivity for asd diagnosis and personalized analysis. In: Medical Image Computing and Computer-Assisted Intervention. pp. 334–343 (2022)
14. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* **146**, 1038–1049 (2017)

15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
16. Li, X., Dvornek, N.C., Zhuang, J., Ventola, P., Duncan, J.S.: Brain biomarker interpretation in asd using deep learning and fmri. In: Medical Image Computing and Computer-Assisted Intervention. pp. 206–214. Springer (2018)
17. Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, J., Scheinost, D., Staib, L.H., Ventola, P., Duncan, J.S.: Braingnn: Interpretable brain graph neural network for fmri analysis. Medical Image Analysis **74**, 102233 (2021)
18. Loth, E., Spooren, W., Ham, L.M., Isaac, M.B., Auriche-Benichou, C., Banaschewski, T., Baron-Cohen, S., Broich, K., Boelte, S., Bourgeron, T., et al.: Identification and validation of biomarkers for autism spectrum disorders. Nature Reviews Drug Discovery **15**(1), 70–70 (2016)
19. Maas, A.L., Hannun, A.Y., Ng, A.Y., et al.: Rectifier nonlinearities improve neural network acoustic models. In: International Conference on Machine Learning. p. 3. Atlanta, GA (2013)
20. Nightingale, S.: Autism spectrum disorders. Nature Reviews Drug Discovery **11**(10), 745 (2012)
21. Parisot, S., Ktena, S.I., Ferrante, E., Lee, M., Guerrero, R., Glocker, B., Rueckert, D.: Disease prediction using graph convolutional networks: application to autism spectrum disorder and alzheimer’s disease. Medical Image Analysis **48**, 117–130 (2018)
22. Peng, Z., Liu, H., Jia, Y., Hou, J.: Attention-driven graph clustering network. In: ACM International Conference on Multimedia. pp. 935–943 (2021)
23. Peng, Z., Liu, H., Jia, Y., Hou, J.: Adaptive attribute and structure subspace clustering network. IEEE Transactions on Image Processing **31**, 3430–3439 (2022)
24. Peng, Z., Liu, H., Jia, Y., Hou, J.: Egrc-net: Embedding-induced graph refinement clustering network. IEEE Transactions on Image Processing **32**, 6457–6468 (2023)
25. Price, T., Wee, C.Y., Gao, W., Shen, D.: Multiple-network classification of childhood autism using functional connectivity dynamics. In: Medical Image Computing and Computer-Assisted Intervention. pp. 177–184. Springer (2014)
26. Rakić, M., Cabezas, M., Kushibar, K., Oliver, A., Llado, X.: Improving the detection of autism spectrum disorder by combining structural and functional mri information. NeuroImage: Clinical **25**, 102181 (2020)
27. Recht, B., Fazel, M., Parrilo, P.A.: Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. SIAM Review pp. 471–501 (2010)
28. Rosenberg, M.D., Finn, E.S.: How to establish robust brain–behavior relationships without thousands of individuals. Nature Neuroscience **25**(7), 835–837 (2022)
29. Wang, L., Li, G., Shi, F., Cao, X., Lian, C., Nie, D., Liu, M., Zhang, H., Li, G., Wu, Z., et al.: Volume-based analysis of 6-month-old infant brain mri for autism biomarker identification and early diagnosis. In: Medical Image Computing and Computer-Assisted Intervention. pp. 411–419. Springer (2018)
30. Wang, M., Zhang, D., Huang, J., Shen, D., Liu, M.: Low-rank representation for multi-center autism spectrum disorder identification. In: Medical Image Computing and Computer-Assisted Intervention. pp. 647–654. Springer (2018)
31. Ye, J.: Generalized low rank approximations of matrices. In: Proceedings of the twenty-first international conference on Machine learning. p. 112 (2004)
32. Ying, C., Cai, T., Luo, S., Zheng, S., Ke, G., He, D., Shen, Y., Liu, T.Y.: Do transformers really perform badly for graph representation? Advances in Neural Information Processing Systems **34**, 28877–28888 (2021)