

RSAD: Region-Specific Anomaly Detection in fMRI for Disease Diagnosis

Yusong Sun¹, Dongdong Chen¹, Mengjun Liu¹, Zhenrong Shen¹, Zhiyun Song¹, Yuqi Hu¹, Manman Fei¹, Xu Han¹, Zelin Liu¹, Xingkai Fang¹, Lu Bai², and Lichi Zhang¹(✉)

¹ School of Biomedical Engineering, Shanghai Jiao Tong University, Shanghai, China
lichizhang@sjtu.edu.cn

² School of Artificial Intelligence, Beijing Normal University, Beijing, China

Abstract. Functional magnetic resonance imaging (fMRI), a noninvasive neuroimaging technique for mapping neural activity, has demonstrated substantial potential in identifying brain disease. However, clinical applications face a critical challenge: patient data are typically scarce compared to abundant healthy control samples. This severe class imbalance significantly limits the performance of classification-based diagnostic models. To address this issue, we propose the Region-Specific Anomaly Detection (RSAD) framework, which formulates the brain disease identification as an anomaly detection task. We first employ pre-training to capture normal patterns of healthy data through a reconstruction task, and then develop the discrepancy score to enhance the model’s ability to perceive potential abnormalities, thereby improving the AD performance. Specifically, we design an affinity matrix learning module and an adaptive region of interest (ROI) masking strategy to improve the performance of signal representation learning. Additionally, we propose a region-based discrepancy score weighting strategy to amplify the distinction between potential abnormalities and healthy controls by assigning higher weights to key brain regions, thereby improving the model’s ability to detect anomalies. We conduct experiments across six different brain diseases, and the superior results demonstrate that RSAD effectively enables disease diagnosis, even with extreme class imbalance. Our code is available at <https://github.com/kylin1112/RSAD>.

Keywords: Anomaly detection · Self-supervised learning · rs-fMRI · Brain disease diagnosis.

1 Introduction

The increasing prevalence of brain diseases, such as depression [14] and dementia [16], significantly impacts both individual well-being and societal progress. Accurate screening and diagnosis are essential to enable effective interventions and equitable healthcare. Functional magnetic resonance imaging (fMRI), which measures blood-oxygen-level-dependent (BOLD) signals to reflect neural activity across brain regions, has emerged as an invaluable tool for diagnosing neurological and psychiatric conditions [11,26,4,5,6]. Most fMRI-based diagnostic

frameworks treat brain disease identification as a classification task, while facing two major challenges. First, task-specific models exhibit limited generalizability, typically focusing on a single disease, whereas clinical practice involves diverse brain diseases that such models cannot comprehensively address. Second, these models are hindered by a pronounced class imbalance. In practical scenarios, the number of healthy control samples frequently exceeds that of patient cases, further compromising the model’s diagnostic performance [23,24].

Analogous to a clinician using knowledge of healthy individuals to detect previously unnoticed abnormalities, anomaly detection (AD) models normal data patterns to identify rare instances of abnormality, providing a promising solution to these challenges. A general AD pipeline first learns to reconstruct normal samples and then detects potential anomalies by evaluating discrepancies between the input and its reconstruction. This approach has achieved substantial success in fields such as computer vision [25,10] and time series analysis [21,8]. Among these methods, autoencoders [18] and generative adversarial networks [10] are commonly utilized to capture a compact latent space that reflects the normal data distribution.

However, the existing AD-based approaches in clinical scenario face three challenges. First, existing AD methods in medical imaging are mostly tailored for 3D data (e.g., T1-weighted MRI [9] and CT scans [7]), whereas fMRI data inherently possess 4D spatiotemporal complexity, hindering the direct adaptation of conventional AD methodologies. Second, it is difficult to optimize the latent space of fMRI data simply by sequence reconstruction, due to the high noise and low information characteristic of BOLD signals. Third, existing methods typically employ reconstruction loss directly as the anomaly score, which exhibits limited sensitivity to subtle variations, particularly in conditions like brain diseases that lack distinct morphological features.

To address the issues above, we introduce the **Region-Specific Anomaly Detection (RSAD)** framework. Our main contributions are as follows: (1) We propose the RSAD framework, which adapts fMRI-based disease diagnosis to AD tasks and achieves disease diagnosis under extreme class imbalance. (2) We design the affinity matrix learning module and the adaptive Region of interest (ROI) masking strategy, capturing the relationships between brain regions and significantly enhancing the masked autoencoder framework’s ability to represent fMRI signals. (3) We introduce a region-specific discrepancy score weighting strategy to amplify the distinctions between potential anomalies and normal controls, thereby effectively improve the pre-trained model’s AD performance. (4) Our approach is evaluated on six distinct brain diseases, and the experimental results demonstrate its superiority over current SOTA methods.

2 Method

2.1 Problem Definition

We define the problem of disease identification as a task of AD. Consider an ROI-wise fMRI BOLD signal $\mathbf{x} \in \mathbb{R}^{N \times T}$, where N represents the number of ROIs

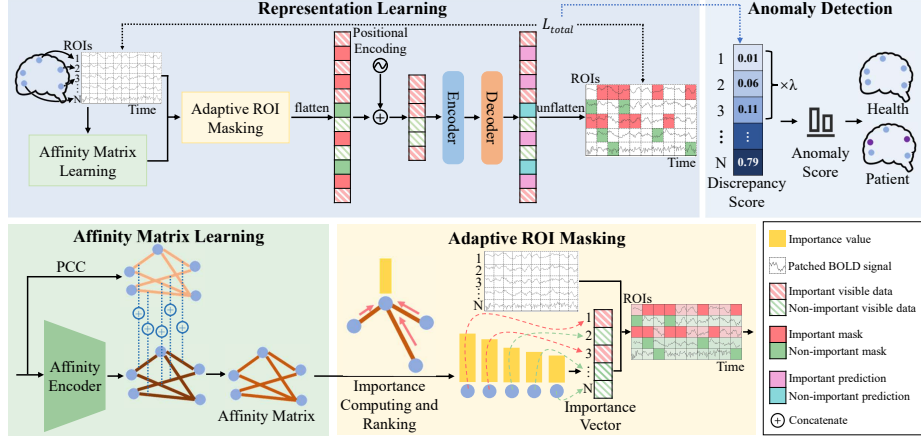


Fig. 1. The architecture of the region-specific anomaly detection (RSAD). RSAD consists of two steps: signal representation learning and anomaly detection. The first step involves the affinity matrix learning module and the adaptive ROI masking strategy.

defined by a brain atlas, and T denotes the number of time points in the fMRI data. During the training phase, a set of fMRI signals from m healthy individuals, denoted as $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$, is used to train the RSAD model, aiming to learn the distribution of healthy fMRI signals. In the AD phase, the pre-trained model is employed to classify each individual in the test set $\mathbf{T} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_j, \dots, \mathbf{x}_{j+k}\}$, where the set includes fMRI data from j unseen healthy individuals and k unseen patients, as either normal or abnormal. Our problem is categorized as self-supervised AD, as no label information is utilized.

2.2 Representation Learning

In the first stage of the framework, the Transformer masked autoencoder serves as the backbone architecture to model the signal distribution of the healthy control group through reconstruction tasks, aiming to learn its representation, as illustrated in Fig. 1. The affinity matrix learning module captures the relationships between brain regions, generates the affinity matrix, and determines brain importance scores to guide adaptive ROI masking. This process prompts the model to focus on the most representative brain regions of healthy individuals.

Affinity Matrix Learning. Fig. 1 illustrates the affinity matrix learning module with a vision transformer serving as the affinity encoder, designed to capture relationships between brain regions, inspired by Liu et al. [11]. The input to the module is $\mathbf{x} \in \mathbb{R}^{N \times T}$, and it's divided into two branches. First, the functional connectivity (FC) matrix $\mathbf{F}_{\text{pear}} \in \mathbb{R}^{N \times N}$ is computed using pearson correlation. We then use the attention map generated by the affinity encoder during forward propagation to represent the relative importance between brain regions [26]. The

final attention matrix $F_{\text{attn}} \in \mathbb{R}^{N \times N}$ is obtained by averaging the attention maps generated by each head in each layer of the affinity encoder. Notably, the matrix is symmetric to ensure interpretability. Finally, by adding the FC matrix F_{pear} and the attention matrix F_{attn} through the corresponding edges, we obtain the affinity matrix $F_{\text{affi}} \in \mathbb{R}^{N \times N}$. This learnable affinity matrix represents the relative relationships between different brain regions in healthy controls.

Adaptive ROI Masking. Random masking has demonstrated considerable success in the context of masked autoencoder, particularly in computer vision tasks. However, its effectiveness in fMRI is not straightforward. Different brain regions have distinct functional roles, and random masking may hinder the model’s ability to accurately capture the critical brain regions or networks essential for the healthy distribution of fMRI signals.

To address this issue, we propose an adaptive ROI masking strategy. Specifically, based on the affinity matrix generated by the affinity matrix learning module, we quantify the importance of each ROI by calculating $\alpha_t \in \mathbb{R}^N$, as shown in the following equation:

$$\alpha_t(j) = \frac{\sum_{i=1}^N F_{\text{affi}}(i, j)}{\sum_{i=1}^N \sum_{j=1}^N F_{\text{affi}}(i, j)}, \quad (1)$$

for $j = 1, 2, \dots, N$. This calculation provides the relative importance of each ROI [26], and the top p ROIs are selected as the important brain regions to guide the input data for masking. We apply different masking ratios R_i and R_n to the important and non-important brain regions, respectively. This strategy encourages the model to focus on learning the signal reconstruction of brain regions most crucial for the healthy distribution during training.

Reconstruction Loss. During the signal reconstruction phase, our model is inspired by BrainLM [3], which has demonstrated the effectiveness of masked autoencoder in learning fMRI representations. The objective is to predict the original signal of the masked patches. Specifically, the transformer encoder processes only the unmasked segments, as shown in the equation:

$$r = \text{Encoder}(M(\mathbf{X})), \quad (2)$$

where M denotes adaptive ROI masking and r represents the learned feature representation. Subsequently, learnable mask tokens are added to obtain r' . The decoder then reconstructs the masked portion:

$$\mathbf{X}' = \text{Decoder}(r'), \quad (3)$$

where \mathbf{X}' is the reconstructed fMRI signal. We use mean square error (MSE) as the reconstruction loss for the masked patches in \mathbf{X} and \mathbf{X}' , expressed as L_{rec} .

To further enhance representation learning, we incorporate variance invariance covariance (VICReg) [2] regularization into our reconstruction loss. VICReg

encourages more variety among the data in the batch by using a hinge loss that maintains the standard deviation above a threshold. Thus, the total loss of our model is updated as follows:

$$L_{total} = L_{rec} + \mu \cdot v(\mathbf{X}') + \gamma \cdot c(\mathbf{X}'), \quad (4)$$

where L_{total} is minimized over the sample batch, and $v(\mathbf{X}')$ and $c(\mathbf{X}')$ represent the variance and covariance losses at the batch level of \mathbf{X}' , respectively. The hyperparameters μ and γ control the relative contribution of each loss term.

2.3 Anomaly Detection

In the AD stage, the model classifies individuals as either normal or abnormal, as shown in Fig. 1. A discrepancy score is introduced to quantify the distributional differences between individuals and healthy controls, allowing the model to generate an anomaly score that identifies brain diseases.

Region-specific Discrepancy Score. Traditional AD methods usually rely on loss functions of the training phase as anomaly scores. However, for fMRI data, dependence on this single metric may yield suboptimal detection performance due to the functional complexity of brain regions. Distinct ROIs exhibit interconnected yet specialized roles, and global averaging across all ROIs risks obscuring critical biomarker contributions. To address this limitation, we leverage intrinsic functional characteristics of fMRI to develop a region-specific weighted discrepancy scoring strategy. Our core hypothesis posits that during pre-training, model learns normal sample distributions, with regions demonstrating lower reconstruction errors providing stronger discriminative power between normal and anomalous patterns. Accordingly, the pre-trained model computes masked patch MSE values for each ROI across healthy control individuals' input \mathbf{X} and reconstructed \mathbf{X}' data matrices. The top q ROIs with minimal MSE values are selected as representative normative regions. Subsequently, we define the discrepancy score $S_{dis} \in \mathbb{R}^N$ as the absolute difference between an individual's ROI-wise MSE matrix and the normal control pattern. For each subject in the test set \mathbf{T} , this generates $j + k$ discrepancy scores. Finally, scores corresponding to the q identified representative ROIs are weighted by coefficient λ , and the anomaly score S_{ano} is calculated as the normalized weighted sum:

$$S_{ano} = \frac{1}{N} \sum_{i=1}^N \lambda_i \cdot S_{dis}(i). \quad (5)$$

This regionally weighted score enhances the sensitivity of anomaly detection by focusing on diagnostically relevant biomarkers.

Table 1. AP Performance metrics for various disorders across different models (%).

Method	Dementia	Depressive	Parkinson	Bipolar	Manic	Anxiety	Mean
AE	51.47 \pm 3.00	52.92 \pm 0.01	60.70 \pm 0.23	50.38 \pm 0.39	72.95 \pm 0.27	59.70 \pm 0.12	58.02 \pm 8.51
GDN	83.92 \pm 4.55	45.51 \pm 0.77	59.04 \pm 2.23	52.02 \pm 2.10	65.56 \pm 4.60	47.55 \pm 0.83	58.93 \pm 13.34
TranAD	54.78 \pm 2.92	62.66 \pm 1.43	58.80 \pm 1.76	67.46 \pm 1.28	65.33 \pm 2.59	52.73 \pm 0.67	60.29 \pm 6.04
OmniAnomaly	68.37 \pm 2.70	61.76 \pm 2.03	58.81 \pm 0.93	67.52 \pm 2.90	66.07 \pm 5.01	54.01 \pm 0.76	62.76 \pm 6.17
Anomaly Transformer	75.82 \pm 1.65	55.15 \pm 0.83	64.49 \pm 3.28	67.14 \pm 1.29	69.51 \pm 4.71	51.07 \pm 0.44	63.86 \pm 8.87
BrainLM	84.13 \pm 6.32	65.14 \pm 1.00	73.78 \pm 3.00	69.40 \pm 2.64	61.68 \pm 3.16	61.40 \pm 0.45	69.25 \pm 8.73
RSAD (Ours)	92.09\pm2.34	74.45\pm2.64	81.81\pm2.67	76.92\pm2.00	75.95\pm9.74	67.30\pm1.66	78.09\pm8.47

Table 2. AUC Performance metrics for various disorders across different models (%).

Method	Dementia	Depressive	Parkinson	Bipolar	Manic	Anxiety	Mean
AE	48.00 \pm 4.24	45.62 \pm 0.10	53.86 \pm 0.51	43.32 \pm 0.56	69.35 \pm 0.75	51.44 \pm 0.23	51.93 \pm 8.99
TranAD	45.40 \pm 4.56	58.35 \pm 0.84	55.38 \pm 2.07	62.05 \pm 2.01	52.78 \pm 1.75	51.94 \pm 0.67	54.32 \pm 5.63
GDN	81.40 \pm 3.62	42.02 \pm 1.02	56.71 \pm 2.26	49.36 \pm 2.04	56.33 \pm 2.91	46.34 \pm 1.67	55.36 \pm 13.36
Anomaly Transformer	67.20 \pm 3.03	50.84 \pm 0.27	58.68 \pm 1.77	60.25 \pm 2.15	58.82 \pm 4.93	42.85 \pm 0.57	56.44 \pm 8.18
OmniAnomaly	58.20 \pm 3.75	56.06 \pm 1.42	56.29 \pm 1.34	65.00 \pm 2.94	63.31 \pm 4.91	51.09 \pm 0.74	58.32 \pm 5.36
BrainLM	83.80 \pm 5.74	62.70 \pm 1.00	72.55 \pm 2.23	69.87 \pm 1.41	61.42 \pm 4.47	61.08 \pm 2.10	68.57 \pm 8.98
RSAD (Ours)	89.80\pm3.46	70.49\pm2.00	78.50\pm1.92	73.24\pm3.31	75.14\pm7.34	63.42\pm1.87	75.09\pm8.48

3 Experimental Results

3.1 Data Preprocessing and Experimental Settings

Datasets and Preprocessing. We construct a dataset based on the UK Biobank (UKB) [13] for the self-supervised pretraining of RSAD. The dataset includes resting-state fMRI recordings from 1,387 healthy individuals and 272 patients with six distinct brain disorders: 10 with Dementia, 63 with Depression, 47 with Parkinson’s Disease, 38 with Bipolar Disorder, 13 with Manic Episodes, and 101 with Phobic Anxiety. All participants are between 40 and 69 years old. For model training, we use 80% of the healthy individuals’ data (1,109 recordings), with 10% allocated for evaluation, and the remaining 10%, along with the patient data, for downstream performance assessment. All recordings undergo standard preprocessing, which included motion correction, normalization, temporal filtering, and ICA-based denoising, as detailed in previous studies [17,1]. To extract group-level time series, we employ the AAL-424 atlas [15], which divides the brain into 424 ROIs.

Implementation details. During training, random subsequences of 200 time steps are selected from each 490-time-step recording. These series are segmented into blocks of 20 time steps, resulting in 10 non-overlapping segments per subsequence, with a context length of 4,240 tokens. The dimensions of the affinity encoder, the encoder, and the decoder are all set to 512, with a depth of 2, 4 and 2, respectively. Tokens are masked with probabilities $R_i = 0.5$ and $R_n = 0.2$ for important and non-important brain regions, respectively. The hyperparameters μ , γ , λ , p and q are set to 1, 1, 500, 10%, 40%, respectively. The epoch with




Fig. 2. The reconstruction performance of RSAD and BrainLM. Ground truth, RSAD, and BrainLM are colored in blue, orange, and green, respectively, with the obvious differences highlighted (red). PCC represents pearson correlation coefficients between the reconstructed signals and ground truth.

the smallest L_{total} on the validation set is selected for evaluating performance on the test set. AD performance is assessed using two metrics: average precision (AP) and area under the receiver operating characteristic curve (AUC).

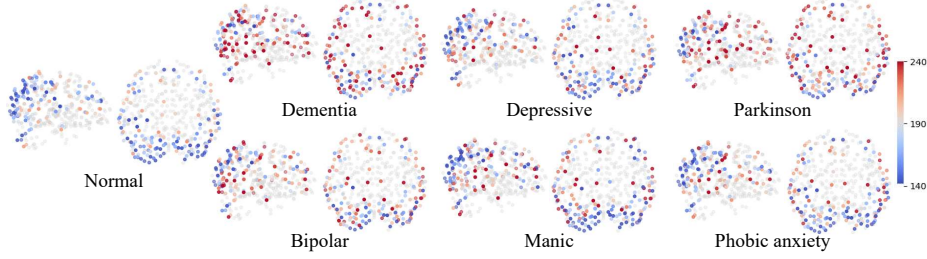
3.2 Results

Comparison with SOTA. We compare Su’s proposed autoencoder for fMRI AD detection [18], BrainLM [3]—the first foundation model for fMRI—and several representative time-series AD methods, including Anomaly Transformer [21], TranAD [20], OmniAnomaly [19], and GDN [8]. As shown in Table 1, RSAD outperforms all methods in terms of AP and AUC, demonstrating an 8% improvement in AP and a 6.5% improvement in AUC over the next best baseline. In particular, while multivariate time series AD methods focus on detecting anomalies at specific time points, our approach emphasizes identifying abnormal patterns and relationships between different ROIs, which are crucial for fMRI data analysis. For the fMRI foundation model, we optimize both the fMRI signal representation and the anomaly perception, resulting in significantly better performance than BrainLM. These results further underscore the effectiveness of the RSAD framework in detecting abnormalities in fMRI data.

Signal Reconstruction. The signal reconstruction results are illustrated in Fig. 2. To evaluate performance, a test-set individual is randomly selected for comparative analysis between our method and BrainLM, which employs pre-trained optimal weights for predictions. Our approach exhibits superior reconstruction quality despite BrainLM’s training dataset being approximately 60 times larger than RSAD’s. Moreover, our method achieves an average Pearson correlation coefficient (PCC) of 0.876, significantly outperforming BrainLM’s result with 0.728. These findings demonstrate that our framework more effectively captures the distribution and spatiotemporal characteristics of fMRI signals.

Table 3. Ablation study of the key components: (a) affinity matrix learning (AML), (b) adoptive ROI masking (ARoM), (c) discrepancy score (S_{dis}).

AML	ARoM	S_{dis}	Dementia		Depressive		Parkinson		Mean	
			AP(%)	AUC(%)	AP(%)	AUC(%)	AP(%)	AUC(%)	AP(%)	AUC(%)
×	×	×	82.40 \pm 6.13	82.20 \pm 5.05	63.72 \pm 2.76	61.79 \pm 1.44	71.78 \pm 2.67	71.29 \pm 1.33	72.63 \pm 8.04	71.76 \pm 8.51
×	✓	×	85.14 \pm 4.24	86.00 \pm 3.31	64.85 \pm 2.23	63.56 \pm 1.00	72.34 \pm 3.60	71.99 \pm 2.82	74.11 \pm 8.57	73.85 \pm 10.22
✓	✓	×	87.23 \pm 3.00	88.20 \pm 1.73	67.65 \pm 1.96	64.54 \pm 1.25	75.35 \pm 2.81	73.85 \pm 3.06	76.74 \pm 8.30	75.53 \pm 10.33
×	×	✓	87.58 \pm 4.35	86.80 \pm 3.74	72.05 \pm 1.73	66.75 \pm 2.44	79.48 \pm 1.73	76.91 \pm 1.41	79.70 \pm 6.55	76.82 \pm 8.13
×	✓	✓	89.53 \pm 2.64	88.60 \pm 4.00	72.76 \pm 2.23	68.60 \pm 1.00	80.26 \pm 2.64	77.57 \pm 2.00	80.85 \pm 7.11	78.26 \pm 8.30
✓	✓	✓	92.09\pm2.34	89.80\pm3.46	74.45\pm2.64	70.49\pm2.00	81.81\pm2.67	78.50\pm1.92	82.78\pm7.58	79.60\pm8.15

**Fig. 3.** Visualization of top q selected region-specific weighted discrepancy scores. The remaining brain regions are colored in light grey, and the colorbar indicates the magnitude of the weighted discrepancy score.

Ablation study. We present the results of the ablation study in Table 3. We have ablated three key components of the framework: AML, ARoM, and the discrepancy score, on three randomly selected brain diseases. The results indicate that the discrepancy score significantly improves AD performance. Additionally, both AML and ARoM contribute approximately 2% to the performance improvement, as they enhance the effectiveness of fMRI representation learning.

Brain Region Analysis. To identify distinguishable brain regions for model recognition, we visualize regions associated with differences in discrepancy score weights and the values of these ROIs, highlighting representative ROIs that differentiate normal individuals from those with various brain diseases. As shown in Fig. 3, the discrepancy scores for patients' ROIs are significantly higher than those of normal individuals. Furthermore, these regions primarily concentrate in the Ventral Attention Network, Dorsal Attention Network, and Default Mode Network, which is consistent with previous literature [22,12], illustrating the effectiveness of the RSAD framework for brain diseases diagnosis.

4 Conclusion

In this paper, we present a novel fMRI-based AD framework, named RSAD. We enhance the representation of fMRI data in the reconstruction task by designing

an affinity matrix learning module and an adaptive ROI masking strategy. Additionally, we introduce a region-specific discrepancy score weighting strategy that is sensitive to potential anomalies. Comprehensive experiments demonstrate the effectiveness and interpretability of the RSAD framework in brain diseases diagnosis. Future work will focus on further investigating the relationship between brain diseases and the functional representations of the brain.

Acknowledgments. This research is supported by the National Natural Science Foundation of China (Grant No. 62471288).

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

References

1. Abdallah, C.G.: Brain networks associated with covid-19 risk: data from 3662 participants. *Chronic Stress* **5**, 24705470211066770 (2021)
2. Bardes, A., Ponce, J., LeCun, Y.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906* (2021)
3. Caro, J.O., Fonseca, A.H.d.O., Averill, C., Rizvi, S.A., Rosati, M., Cross, J.L., Mittal, P., Zappala, E., Levine, D., Dhodapkar, R.M., et al.: Brainlm: A foundation model for brain activity recordings. *bioRxiv* pp. 2023–09 (2023)
4. Chen, D., Liu, M., Shen, Z., Yao, L., Zhao, X., Song, Z., Yuan, H., Wang, Q., Zhang, L.: Exploring multiconnectivity and subdivision functions of brain network via heterogeneous graph network for cognitive disorder identification. *IEEE Transactions on Neural Networks and Learning Systems* (2024)
5. Chen, D., Liu, M., Wang, S., Li, Z., Bai, L., Wang, Q., Shen, D., Zhang, L.: Guiding fusion of dynamic functional and effective connectivity in spatio-temporal graph neural network for brain disorder classification. *Knowledge-Based Systems* **309**, 112856 (2025)
6. Chen, D., Yao, L., Liu, M., Shen, Z., Hu, Y., Song, Z., Wang, Q., Zhang, L.: Self-supervised learning with adaptive graph structure and function representation for cross-dataset brain disorder diagnosis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 612–622. Springer (2024)
7. Chen, L., You, Z., Zhang, N., Xi, J., Le, X.: Utrad: Anomaly detection and localization with u-transformer. *Neural Networks* **147**, 53–62 (2022)
8. Deng, A., Hooi, B.: Graph neural network-based anomaly detection in multivariate time series. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 35, pp. 4027–4035 (2021)
9. Han, C., Rundo, L., Murao, K., Noguchi, T., Shimahara, Y., Milacski, Z.Á., Koshino, S., Sala, E., Nakayama, H., Satoh, S.: Madgan: Unsupervised medical anomaly detection gan using multiple adjacent brain mri slice reconstruction. *BMC bioinformatics* **22**, 1–20 (2021)
10. Huang, H., Shen, Z., Wang, J., Wang, X., Lu, J., Lin, H., Ge, J., Zuo, C., Wang, Q.: Metaad: Metabolism-aware anomaly detection for parkinson’s disease in 3d 18 f-fdg pet. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 291–301. Springer (2024)

11. Liu, M., Song, Z., Chen, D., Wang, X., Zhuang, Z., Fei, M., Zhang, L., Wang, Q.: Affinity learning based brain function representation for disease diagnosis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 14–23. Springer (2024)
12. Ma, W., Yao, Q., Hu, G., Ge, H., Xue, C., Wang, Y., Yan, Y., Xiao, C., Shi, J., Chen, J.: Reorganization of rich clubs in functional brain networks of dementia with lewy bodies and alzheimer’s disease. *NeuroImage: Clinical* **33**, 102930 (2022)
13. Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., Xu, J., Bartsch, A.J., Jbabdi, S., Sotiropoulos, S.N., Andersson, J.L., et al.: Multimodal population brain imaging in the uk biobank prospective epidemiological study. *Nature neuroscience* **19**(11), 1523–1536 (2016)
14. Moreno-Agostino, D., Wu, Y.T., Daskalopoulou, C., Hasan, M.T., Huisman, M., Prina, M.: Global trends in the prevalence and incidence of depression: a systematic review and meta-analysis. *Journal of affective disorders* **281**, 235–243 (2021)
15. Nemati, S., Akiki, T.J., Roscoe, J., Ju, Y., Averill, C.L., Fouda, S., Dutta, A., McKie, S., Krystal, J.H., Deakin, J.W., et al.: A unique brain connectome fingerprint predates and predicts response to antidepressants. *IScience* **23**(1) (2020)
16. Nichols, E., Steinmetz, J.D., Vollset, S.E., Fukutaki, K., Chalek, J., Abd-Allah, F., Abdoli, A., Abualhasan, A., Abu-Gharbieh, E., Akram, T.T., et al.: Estimation of the global prevalence of dementia in 2019 and forecasted prevalence in 2050: an analysis for the global burden of disease study 2019. *The Lancet Public Health* **7**(2), e105–e125 (2022)
17. Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L., Smith, S.M.: Automatic denoising of functional mri data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* **90**, 449–468 (2014)
18. Su, J., Shen, H., Peng, L., Hu, D.: Few-shot domain-adaptive anomaly detection for cross-site brain images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(3), 1819–1835 (2021)
19. Su, Y., Zhao, Y., Niu, C., Liu, R., Sun, W., Pei, D.: Robust anomaly detection for multivariate time series through stochastic recurrent neural network. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 2828–2837 (2019)
20. Tuli, S., Casale, G., Jennings, N.R.: Tranad: Deep transformer networks for anomaly detection in multivariate time series data. *arXiv preprint arXiv:2201.07284* (2022)
21. Xu, J., Wu, H., Wang, J., Long, M.: Anomaly transformer: Time series anomaly detection with association discrepancy. *arXiv preprint arXiv:2110.02642* (2021)
22. Yang, H., Chen, X., Chen, Z.B., Li, L., Li, X.Y., Castellanos, F.X., Bai, T.J., Bo, Q.J., Cao, J., Chang, Z.K., et al.: Disrupted intrinsic functional brain topology in patients with major depressive disorder. *Molecular psychiatry* **26**(12), 7363–7371 (2021)
23. Zeng, L., Shen, H., Liu, L., Wang, L., Li, B., Fang, P., Zhou, Z., Li, Y., Hu, D.: Identifying major depression using whole-brain functional connectivity: a multivariate pattern analysis. *Brain* **135**(5), 1498–1507 (2012)
24. Zeng, L., Wang, H., Hu, P., Yang, B., Pu, W., Shen, H., Chen, X., Liu, Z., Yin, H., Tan, Q., et al.: Multi-site diagnostic classification of schizophrenia using discriminant deep learning with functional connectivity mri. *EBioMedicine* **30**, 74–85 (2018)

25. Zhang, X., Mu, J., Zhang, X., Liu, H., Zong, L., Li, Y.: Deep anomaly detection with self-supervised learning and adversarial training. *Pattern Recognition* **121**, 108234 (2022)
26. Zhao, Y., Nie, D., Chen, G., Wu, X., Zhang, D., Wen, X.: Tardrl: Task-aware reconstruction for dynamic representation learning of fmri. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 700–710. Springer (2024)