

# DW-DGAT: Dynamically Weighted Dual Graph Attention Network for Neurodegenerative Disease Diagnosis

Chengjia Liang<sup>1</sup>, Zhenjiong Wang<sup>1</sup>, Chao Chen<sup>2</sup>, Ruizhi Zhang<sup>1</sup>, Songxi Liang<sup>1</sup>,  
Hai Xie<sup>3</sup>, Haijun Lei<sup>1\*</sup>, Zhongwei Huang<sup>4</sup>

<sup>1</sup>College of Computer Science and Software Technology, Shenzhen University, Shenzhen 518000, China

<sup>2</sup>School of Computer Science, University of Nottingham, Nottingham NG8 1BB, U.K.

<sup>3</sup>School of Cyberspace Security, Hainan University, Haikou 570228, China

<sup>4</sup>School of Computer Science, Hubei University of Technology, Wuhan 430068, China

{liangchengjia2022, 2021192003, 2022110191, 2022110319}@email.szu.edu.cn, chao.chen@nottingham.ac.uk,  
xieh@hainanu.edu.cn, lhj@szu.edu.cn, huangzhongwei@hbut.edu.cn

## Abstract

Parkinson's disease (PD) and Alzheimer's disease (AD) are the two most prevalent and incurable neurodegenerative diseases (NDs) worldwide, for which early diagnosis is critical to delay their progression. However, the high dimensionality of multi-metric data with diverse structural forms, the heterogeneity of neuroimaging and phenotypic data, and class imbalance collectively pose significant challenges to early ND diagnosis. To address these challenges, we propose a dynamically weighted dual graph attention network (DW-DGAT) that integrates: (1) a general-purpose data fusion strategy to merge three structural forms of multi-metric data; (2) a dual graph attention architecture based on brain regions and inter-sample relationships to extract both micro- and macro-level features; and (3) a class weight generation mechanism combined with two stable and effective loss functions to mitigate class imbalance. Rigorous experiments, based on the Parkinson Progression Marker Initiative (PPMI) and Alzheimer's Disease Neuroimaging Initiative (ADNI) studies, demonstrate the state-of-the-art performance of our approach.

## Code —

<https://github.com/AlexanderLeung9/DW-DGAT.git>

**Extended version** — <https://arxiv.org/abs/2601.10001>

## Introduction

Parkinson's disease (PD) and Alzheimer's disease (AD) are incurable neurodegenerative diseases (NDs) that primarily affect middle-aged and elderly individuals (Balestrino and Schapira 2020; Scheltens et al. 2021). Over time, patients succumb to complications arising from NDs. Early diagnosis is crucial for initiating treatment to slow disease progression. However, neuroimaging changes in the early stages of NDs are often too subtle for physicians to detect. Specifically, there exists an intermediate prodromal (PRO) stage between healthy controls (HC) and PD, and an early mild cognitive impairment (EMCI) stage between cognitively normal (CN) and AD. The tiny differences among these cohorts necessitate the use of highly sensitive neuroimaging modalities

for accurate diagnosis, among which magnetic resonance imaging (MRI) and diffusion tensor imaging (DTI) are two economical yet promising and complementary candidates.

T1-weighted MRI reflects static brain structures, including gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF). DTI is an advanced MRI that can reveal pathological alterations in WM and has shown promise in distinguishing early NDs from healthy subjects (Poewe et al. 2017; Mayo et al. 2019). However, fully exploiting DTI data remains challenging. First, DTI provides multiple 3D metrics, such as fractional anisotropy (FA), mean diffusivity (MD), local diffusion homogeneity based on Spearman's or Kendall's coefficient concordance (LDH-S or LDH-K) (Gong 2013), the axial diffusivity (AXD, equivalent to the first eigenvalue L1), the radial diffusivity (RDD, defined as the mean of L2 and L3), and the 1st to 3rd eigenvectors (V1~V3), among others. Although these metrics offer complementary insights into the pathological brain changes, directly using all of them results in significant memory consumption and prohibitive computational costs. Second, when conducting brain region-level analysis on these metrics, it can yield 2D brain regional networks and 1D brain regional statistics that differ structurally from the original 3D metrics. As a result, existing studies often utilize only a limited subset of DTI or MRI metrics due to the difficulty of fusing multiple structural forms and high-dimensional data. For example, Huang et al. (2025) employed GM from MRI along with L1 and V1 from DTI for PD diagnosis. Song et al. (2022) combined resting-state functional MRI (rs-fMRI) and FA deterministic networks for AD diagnosis.

Regarding neuroimaging data, there are two main types of analysis methods. The first type aims to extract features directly from brain images but faces challenges such as high memory consumption and computational cost in high-dimensional 3D images (Cobbinah et al. 2022; Yang et al. 2023) or incomplete perspectives in 2D slice images (Chen et al. 2024b). The second type segments the brain into multiple regions of interest (ROIs), constructs 2D brain networks, and applies graph-based analysis for feature extraction (Zhang et al. 2023b; Huang et al. 2024). Therefore, it reduces the ND diagnosis problem to a graph classification task. Although significant progress has been made, existing

\*Haijun Lei is the corresponding author.

methods still rely on limited information sources, often overlooking other structural forms and data modalities.

In addition to neuroimaging data, phenotypic data, such as sex, age, and clinical assessment scores, are commonly used in diagnosis by physicians. However, directly fusing these low-dimensional but highly informative data with high-dimensional neuroimaging data often compromises their diagnostic value. This raises a challenge in designing an architecture capable of accurately extracting sample features at both levels. Graph neural networks (GNNs) have shown promise in modeling inter-sample relationships, but they face two key challenges. First, graph construction is critical, as there is no unique inherent graph that precisely reflects the objective relationships between samples. Second, GNNs typically perform better in transductive learning settings than in inductive ones; therefore, transductive learning is widely adopted in the neurological field (Ghorbani et al. 2022; Song et al. 2022; Zhang et al. 2022; Liu et al. 2022; Song et al. 2025). However, transductive learning requires loading all samples during training, resulting in a significant memory burden and reduced flexibility for clinical use.

Finally, medical datasets often suffer from class imbalance, which leads conventional deep learning methods to misclassify minority-class samples as majority-class ones. Traditional strategies typically rely on over-sampling via data generation (Balakrishnan et al. 2023) based on minority classes, or on under-sampling via a bootstrapping method (Roy et al. 2023) or feature fusion (Yue et al. 2024) based on majority classes. However, data generation is complex for multi-metric and multi-form medical data that contains numerous subtle yet critical details; meanwhile, under-sampling risks severe overfitting in small-sized, high-dimensional datasets. Hence, cost-sensitive approaches have emerged as promising alternatives in medical image analysis (Kavitha and Kasthuri 2024). However, a single fixed cost function tailored to a specific dataset structure often fails to generalize to more diverse datasets.

In this study, we fuse MRI and DTI metrics of different structural forms through a general-purpose and efficient strategy to maximize their utility in ND diagnosis. Based on the fused data, we design a single graph attention (SGA) module to learn long-range dependencies among ROIs, and a global graph attention (GGA) module based on phenotypic data to capture demographic relationships among subjects. Lastly, we incorporate a class weight generator (CWG) that dynamically generates class weights to guide the classifier within an cooperative training framework.

Our contributions are summarized as follows.

- We propose a data fusion (DF) method to efficiently fuse 1D, 2D and 3D multi-metric data for subsequent analysis.
- We design a dual graph attention network (DGAT) to learn micro- and macro-level graph features.
- We develop a dynamically weighted (DW) mechanism to alleviate the issue of class imbalance in medical datasets.
- The superiority and generalizability of the proposed DW-DGAT are validated on PD and AD datasets.

## Related Work

**Data fusion.** Although abundant multi-metric data can provide rich information for deep neural networks, efficiently fusing data with diverse structural forms for diagnosis remains challenging. For instance, some studies integrate multiple structural forms of data in a high-level embedding space using specialized architectures (Bi et al. 2024; Ding, Hsu, and Liu 2024), but these designs are typically limited to a few specific metrics and structural forms. Recently, Bourriez et al. (2024) introduced the ChAda-ViT network, which adaptively handles images with varying numbers of channels and projects them into a common embedding space. However, this method is restricted to 2D image data. A generalized approach capable of integrating multi-metric data across cubic, planar, and linear structures is still lacking.

**Brain network analysis.** Since the brain can be viewed as an ensemble of ROIs, many researchers have focused on analyzing graphs constructed from these regions and their interconnections. Convolutional layers and fully connected (FC) layers are widely used for message passing in ND-specific GNNs, aggregating information from node neighbors layer by layer (Kawahara et al. 2017; Xu et al. 2021; Kan et al. 2022a; Zhou et al. 2024), thereby enabling feature extraction from a single graph. Another important analytical technique is the attention mechanism (Kan et al. 2022b; Ying et al. 2021; Yu et al. 2024; Shehzad et al. 2025), which encodes nodes and edges via multi-head self-attention (MHSA) or graph attention networks (GATs), capturing global dependencies among ROIs. Additionally, graph pooling is widely employed in neuroscience to select key ROIs while discarding irrelevant ones, thus enabling versatile GNNs to adapt to brain networks (Peng et al. 2020; Li et al. 2021; Zhang et al. 2023a; Xu et al. 2024).

**Subject relationship analysis.** Different from brain network analysis, subject relationship analysis addresses the problem of node prediction and primarily focuses on phenotypic data. There are mainly two types of methods to leverage phenotypic data. The first type incorporates phenotypic data without any learnable parameters, preserving their original characteristics (Ghorbani et al. 2022; Liu et al. 2022). Consequently, the diagnostic performance heavily depends on the quality of the phenotypic data used to construct a graph representing subject relationships. In contrast, the second type of methods incorporates learnable parameters to model phenotypic data, aiming to construct a more accurate relationship graph among subjects (Song et al. 2022; Zhang et al. 2022; Song et al. 2025). However, these methods risk constructing suboptimal graphs when noisy node features are included in the adjacency graph. Recently, Chen et al. (2024a) proposed GCN-MHSA, which builds the adjacency graph solely from node features via an attention mechanism, in contrast to previous methods that rely on convolutional or affine operations.

**Class imbalance.** Medical datasets often suffer from the issue of class imbalance. To alleviate this, Ghorbani et al. (2022) proposed RA-GCN, a network based on a re-weighted cost-sensitive strategy that resembles the train-

ing paradigm of generative adversarial networks (GANs) (Goodfellow et al. 2020). However, GAN training is notoriously unstable and prone to collapse without appropriate hyperparameter tuning and regularization (Arjovsky and Bottou 2017). Although several techniques have been proposed to address these issues (Gulrajani et al. 2017), RA-GCN remains a relatively plain design. In our experiments, it occasionally exhibits instability and fails to converge. In this work, we propose a similar generative cooperative strategy but go a step further by explicitly addressing the instability and non-convergence issues inherent in GAN training.

## Datasets

We built two datasets from the Parkinson Progression Marker Initiative (PPMI) (Marek et al. 2018) and the Alzheimer's Disease Neuroimaging Initiative 3 (ADNI3) (Weiner et al. 2017), which are globally authoritative and publicly available ND data sources. The PPMI dataset contains three cohorts: HC, PRO, and PD, with 69, 72, and 175 subjects, and 121, 123, and 392 neuroimaging samples, respectively. The ADNI3 dataset contains another three cohorts: CN, EMCI, and AD, with 163, 118, and 29 subjects, and 234, 193, and 37 neuroimaging samples, respectively. All acquisitions are at time points of 0, 12, and 24 months.

We used the PANDA toolkit (Cui et al. 2013), which provides a fully automated pipeline based on the FMRIB Software Library (FSL) (Smith et al. 2004), to preprocess the data. As shown in Figure 1, the pipeline first registers DTI to T1-MRI, and then generates six 3D diffusion metrics with a voxel size of 2 mm in standard space. Furthermore, three 2D deterministic brain region connectivity networks—FA, FN (Fiber Number), and FL (Fiber Length)—are constructed, along with two 1D vectors containing statistical information on ROI surface size and ROI voxel size. Sex, age, years of education, race, and Montreal Cognitive Assessment (MoCA) scores are used as phenotypes in both datasets. Additionally, the PPMI dataset includes MDS-UPDRS Part II scores, and the ADNI3 dataset includes Mini-Mental State Examination (MMSE) scores. These three clinical assessment scores were selected because they are easy to administer by community health providers and patients.

## Methodology

The overall architecture of our proposed methodology is illustrated in Figure 1, which consists of four key components: DF, SGA, GGA, and CWG. The DF module fuses the 1D, 2D, and 3D data into a matrix  $\mathbf{X}$  whose rows represent ROIs and columns correspond to metrics. The SGA first constructs an ROI graph from  $\mathbf{X}$ , and then pools the graph and flattens it back into an ROI sequence. Next, it uses a small ViT encoder (Dosovitskiy et al. 2020) to compute the class token (CLS) for each sample. The GGA dynamically constructs a phenotypic adjacency graph based on the current batch of samples and captures sample-level graph features through two MHSA graph convolution (GC) layers. In the training phase, the CWG dynamically adjusts the class weights of each sample to balance the classifier output.

## Data Fusion

This module fuses three structural forms of metrics. Let  $R = 90$  be the ROI count based on the AAL-90 atlas. To merge the two 1D vectors, we compute the ratio of the number of surface voxels to the total voxel count in each ROI. For each 2D deterministic network, we first rescale all values to the range  $[0, 1]$  using a min-max function, and then apply the  $L_1$ -norm to each row to produce an  $R$ -dimensional vector. Let the 3D neuroimage of every metric be  $\{w_{i,j,k} | i, j, k \in \mathbb{N}, w_{i,j,k} \geq 0, w_{i,j,k} \in \mathbb{R}\}$ , and the brain template be  $\{t_{i,j,k} | t_{i,j,k} \in [1, R], i, j, k, t_{i,j,k} \in \mathbb{N}\}$ . We segment the ROIs by masking  $w_{i,j,k}$  with  $t_{i,j,k}$  along the ROI numbers  $\{r | r \in [1, R], r \in \mathbb{N}\}$ . Specifically, for  $i, j, k$  that satisfy  $t_{i,j,k} = r$ , the corresponding  $w_{i,j,k}$  are collected. For each ROI, we calculate its centroid coordinates  $(\bar{x}, \bar{y}, \bar{z})_r$ , centroid weight  $w_r$ , average weight  $\bar{w}_r$ , and maximum weight  $\hat{w}_r$ , as defined in Eq. 1, 2, 3, 4.

$$(\bar{x}, \bar{y}, \bar{z})_r = \frac{1}{\sum_{i,j,k} w_{i,j,k}} * \left( \sum_{i,j,k} i \cdot w_{i,j,k}, \sum_{i,j,k} j \cdot w_{i,j,k}, \sum_{i,j,k} k \cdot w_{i,j,k} \right),$$

$$i, j, k \in \{i, j, k | t_{i,j,k} = r\} \quad (1)$$

$$w_r = w_{\text{round}(\bar{x}), \text{round}(\bar{y}), \text{round}(\bar{z})} \quad (2)$$

$$\bar{w}_r = \frac{\sum_{i,j,k} w_{i,j,k}}{|\{w_{i,j,k}\}|}, \quad i, j, k \in \{i, j, k | t_{i,j,k} = r\} \quad (3)$$

$$\hat{w}_r = \max(\{w_{i,j,k}\}), \quad i, j, k \in \{i, j, k | t_{i,j,k} = r\} \quad (4)$$

Since the MRI-derived features exceed 1 in range, they are normalized by dividing both  $w_r$  and  $\bar{w}_r$  by  $\hat{w}_r$ , and scaling each  $\hat{w}_r$  by  $\max(\{\hat{w}_r\})$ . Next, we concatenate  $w_r$ ,  $\bar{w}_r$ , and  $\hat{w}_r$  to form a vector. Finally, we concatenate the 1D, 2D, and 3D features for each ROI, and then stack the ROI features into  $\mathbf{X} \in \mathbb{R}^{R \times F}$  for each sample.

## Single Graph Attention

The SGA module utilizes two crucial techniques to extract features of an ROI graph.

The first is graph pooling. Initially, a matrix is computed based on the pairwise Euclidean distances between the ROI features. The centrality distance of each ROI is then obtained by summing its distances to all other ROIs. Subsequently, we zero out the features for the 50% of ROIs with the greatest centrality distances, as these ROIs are assumed to have weaker connectivity and thus contribute less to ND diagnosis. Furthermore, a Gaussian kernel function is used to compute the mean similarity of each ROI based on its centrality distance, with the kernel width  $\sigma$  set to the average of  $R$  centrality distances. After appending the mean similarity to the metric features of each ROI, all ROIs are stacked into  $\mathbf{X}_1 \in \mathbb{R}^{R \times F'}$  according to their original indexing order.

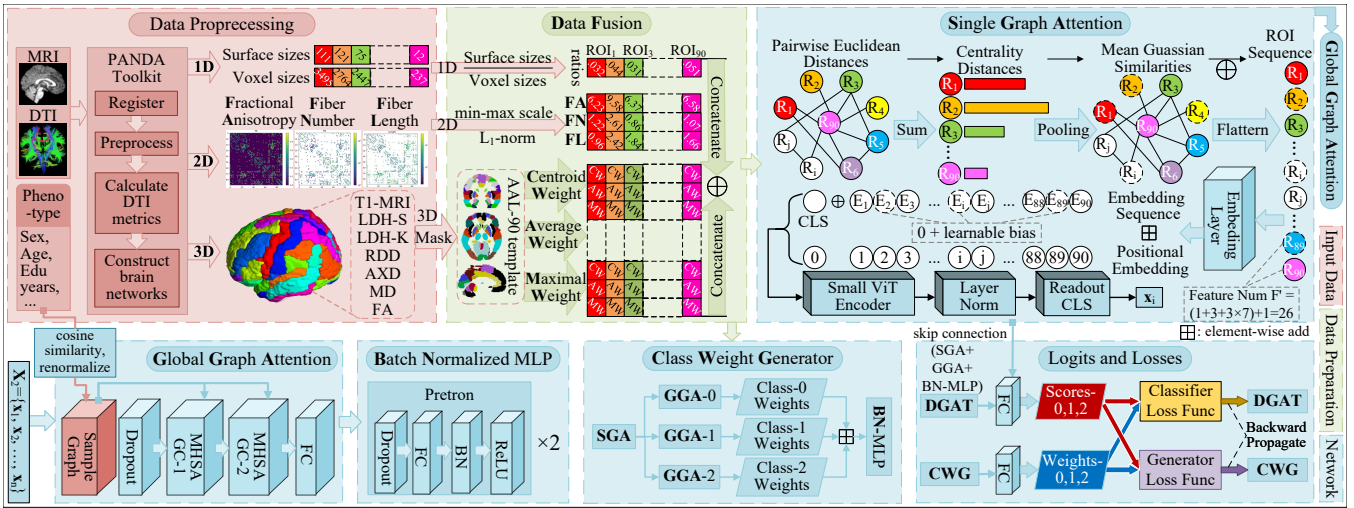


Figure 1: The architecture of the proposed methodology.

The second technique is the ViT encoder. SGA uses an embedding layer to project  $\mathbf{X}_1$  into  $\mathbf{Y}_1 \in \mathbb{R}^{R \times E}$ , where  $E = 384$ . A learnable CLS  $\mathbf{t}_1 \in \mathbb{R}^{1 \times E}$  is then prepended to  $\mathbf{Y}_1$ , yielding the extended ROI embeddings  $\mathbf{Y}_2 \in \mathbb{R}^{(1+R) \times E}$ . Next, learnable positional embeddings  $\mathbf{P}_2 \in \mathbb{R}^{(1+R) \times E}$  are added to  $\mathbf{Y}_2$ , resulting in  $\mathbf{Y}_3$ , which enables SGA to capture positional relationships among ROIs.  $\mathbf{Y}_3$  is fed into a small ViT encoder composed of 12 MHSA blocks to compute global attention. Layer normalization is applied afterward to stabilize the output. Finally, the CLS  $\mathbf{x}_i$  is read out as the representation of the sample.  $N$  sample representations form the sample matrix  $\mathbf{X}_2 = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]^T \in \mathbb{R}^{N \times E}$ .

## Global Graph Attention

We observed an important phenomenon in our experiments: conventional Graph Convolutional Network (GCN) tends to perform better when the underlying adjacency graph more accurately reflects the true relationships among samples. To this end, we enhance the GCN from two aspects.

First, we construct the adjacency matrix  $\mathbf{A}_2$  from the static phenotypic feature matrix  $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N]^T \in \mathbb{R}^{N \times P}$ . Specifically, we measure the pairwise distances between samples based on the transformed cosine similarity defined in Eq. 5. Then, we apply a Gaussian kernel function, with  $\sigma$  set to the median of all off-diagonal pairwise distances, to build a similarity matrix  $\mathbf{A}_1 \in \mathbb{R}^{N \times N}$ . Next, we remove self-loops from  $\mathbf{A}_1$  and apply the renormalization trick (Kipf and Welling 2016) to obtain  $\mathbf{A}_2$ .

$$d_{i,j} = \begin{cases} 1 - \frac{\mathbf{p}_i \cdot \mathbf{p}_j}{\|\mathbf{p}_i\|_2 \|\mathbf{p}_j\|_2}, & i \neq j \\ 1, & i = j \end{cases} \quad (5)$$

Second, we redesign the GC layer of the GCN by replacing its affine weight matrix with MHSA. Let  $(\mathbf{H}_k)_{i,j}$  be the element in the  $i$ -th row and  $j$ -th column of a matrix. We

compute the matrix  $\mathbf{H}_k \in \mathbb{R}^{N \times E}$  for the  $k$ -th sample as shown in Eq. 6.

$$(\mathbf{H}_k)_{i,j} = (\mathbf{A}_2)_{k,i} * (\mathbf{X}_2)_{i,j} \quad (6)$$

Then, we compute the neighbor aggregation  $\mathbf{Z}_k^h \in \mathbb{R}^{N \times d}$  for each attention head to adaptively adjust the edge weights of the adjacency graph, as defined in Eq. 7 and Eq. 8, where  $h \in [1, H]$ ,  $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{E \times E}$ , and  $d = E/H$ . We first obtain  $\mathbf{Q}_k, \mathbf{K}_k, \mathbf{V}_k \in \mathbb{R}^{N \times E}$  and then split them along the feature dimension into  $H$  heads, i.e.,  $\mathbf{Q}_k = [\mathbf{Q}_k^1, \mathbf{Q}_k^2, \dots, \mathbf{Q}_k^H]$  with  $\mathbf{Q}_k^h \in \mathbb{R}^{N \times d}$ . The per-head outputs are concatenated as  $\mathbf{Z}_k = [\mathbf{Z}_k^1, \mathbf{Z}_k^2, \dots, \mathbf{Z}_k^H] \in \mathbb{R}^{N \times E}$ , and the outputs from all samples are stacked to form  $\mathbf{Z} = [\mathbf{Z}_1; \mathbf{Z}_2; \dots; \mathbf{Z}_N] \in \mathbb{R}^{N \times N \times E}$ .

$$\mathbf{Z}_k^h = \text{softmax}\left(\frac{\mathbf{Q}_k^h (\mathbf{K}_k^h)^T}{\sqrt{d}}\right) \mathbf{V}_k^h \quad (7)$$

$$\mathbf{Q}_k, \mathbf{K}_k, \mathbf{V}_k = \mathbf{H}_k * (\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V) \quad (8)$$

Finally, the hidden representation of the GC layer is obtained by summing  $\mathbf{Z}$  along the second axis, i.e.,  $\mathbf{H} = \sum_j \mathbf{Z}_{i,j,k} \in \mathbb{R}^{N \times E}$ , followed by a layer normalization and a GELU activation (Hendrycks and Gimpel 2016). In addition, each of the two MHSA-GC layers outputs features with twice the dimensionality of its input, while the final FC layer reduces the dimensionality to half of its input.

## Class Weight Generator

The architecture of CWG is highly similar to DGAT, except that it consists of  $C$  GGAs, where  $C$  denotes the number of classes. The supporting graphs of GGAs are constructed by masking  $\mathbf{A}_2$  according to the corresponding class labels, allowing each GGA to focus exclusively on the samples within its class. The logits generated by the GGAs are summed element-wise, and then passed through a batch normalized (BN) MLP to produce the final class weights.

Compared to RA-GCN (Ghorbani et al. 2022), we develop two more stable and effective loss functions for both the generator and the classifier. The class weights are denoted as  $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_N]^\top \in \mathbb{R}^{N \times C}$ ; the DGAT scores as  $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N]^\top \in \mathbb{R}^{N \times C}$ ; and the one-hot labels as  $\mathbf{Y} \in \mathbb{N}^{N \times C}$ . We define the two losses  $L_1$  in Eq. 9 for DGAT and  $L_3$  in Eq. 14 for CWG, where  $\text{CE}_{\text{SGA}}$  is a standard cross-entropy loss computed from the SGA logits and  $\mathbf{Y}$ , and  $\alpha = 0.5$  is a hyperparameter.

$$L_1 = \min_{\text{DGAT}} \frac{1}{2} (L_2 + \text{CE}_{\text{SGA}}) \quad (9)$$

$$L_2 = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C \mathbf{Y}_{i,j} \cdot \mathbf{R}_{i,j} \cdot \log \mathbf{O}_{i,j} \quad (10)$$

$$\mathbf{O}_{i,j} = \frac{\exp(\mathbf{S}_{i,j} - \max(\mathbf{s}_i))}{\sum_{k=1}^C \exp(\mathbf{S}_{i,k} - \max(\mathbf{s}_i))} \quad (11)$$

$$\mathbf{R}_{i,j} = \frac{\mathbf{Q}_{i,j}}{\min(\mathbf{q}_i)} \quad (12)$$

$$\mathbf{Q}_{i,j} = \frac{\exp(-(\mathbf{W}_{i,j} - \min(\mathbf{w}_i)))}{\sum_{k=1}^C \exp(-(\mathbf{W}_{i,k} - \min(\mathbf{w}_i)))} + \epsilon \quad (13)$$

$$L_3 = \min_{\text{CWG}} \left( L_2 - \frac{\alpha}{N \cdot C} \sum_{i=1}^N \sum_{j=1}^C \mathbf{R}_{i,j} \cdot \log \mathbf{R}_{i,j} \right) \quad (14)$$

Eq. 11 and Eq. 13 subtract the extreme elements of vectors to align their values on the same side of the y-axis, thus avoiding division by large numbers in the softmax function, which may cause numerical instability. A small constant  $\epsilon$  equal to the double-precision machine epsilon is added to Eq. 13 to prevent gradient vanishing. To address class imbalance, Eq. 13 inverts the numerical relationship among class weights, while Eq. 12 penalizes the classifier for incorrect class predictions without affecting the other classes, thus ensuring function stability. Algorithm 1 outlines the training steps in an epoch for CWG and DGAT.

## Experiments

### Implementation Details

The proposed method was implemented using Python 3.9, PyTorch 1.13.1+cu116, and CUDA drivers of v525 or v535 based on Ubuntu 20.04. All networks ran on a 24 GB GPU. The random seed was fixed at 231 for all experiments. GGA used  $H = 6$  on the PPMI dataset and  $H = 8$  on the ADNI3 dataset. Our network was trained for 500 epochs using the Adam optimizer (Kingma and Ba 2014) with a learning rate of 0.001. The dropout rate was set to 0.5. The batch size was set to 64 for both our network and baseline networks. All experiments were evaluated via ten-fold cross-validation (CV). To prevent data leakage, samples from the same subject at different time points were strictly assigned to either the training set or the test set.

The evaluation metrics include accuracy (ACC), balanced accuracy (BA), F1 score, and specificity (SPE). BA accounts for class imbalance, where 1.0 indicates perfect classification, and  $\frac{1}{C}$  or lower values mean a random guess.

---

### Algorithm 1: Training procedure of CWG and DGAT.

---

**Input:**  $\mathbf{X}'_1 = [\mathbf{X}'_1; \mathbf{X}'_2; \dots; \mathbf{X}'_N] \in \mathbb{R}^{N \times R \times F'}$ ,  $\mathbf{P}$ ,  $\mathbf{Y}$

**Output:** DGAT

- 1: **for** a batch of samples in all training samples **do**
  - 2:   Build the adjacency graph  $\mathbf{A}_2$  from  $\mathbf{P}$
  - 3:   Build masked graphs  $\mathbf{M}_0, \mathbf{M}_1, \mathbf{M}_2$  from  $\mathbf{Y}$
  - 4:    $\mathbf{A}_2^i = \mathbf{A}_2 * \mathbf{M}_i, i \in [0, C], i \in \mathbb{N}$
  - 5:    $\mathbf{W}'_1 = \text{CWG}_i.\text{forward}(\mathbf{X}'_1, \mathbf{A}_2^i) \in \mathbb{R}^{N \times C}$
  - 6:    $\mathbf{W}_1 = [\mathbf{W}'_1[:, i]]_{i=0}^{C-1} \in \mathbb{R}^{N \times C}$  // concat vectors
  - 7:    $\mathbf{S} = \text{DGAT}.\text{forward}(\mathbf{X}'_1, \mathbf{A}_2) \in \mathbb{R}^{N \times C}$
  - 8:    $L_3 = \text{CWG}.\text{cooperative\_loss}(\mathbf{W}_1, \mathbf{Y}, \mathbf{S})$
  - 9:   Descend gradients of CWG from  $L_3$
  - 10:   Update parameters of CWG
  - 11:    $\mathbf{W}'_2 = \text{CWG}_i.\text{forward}(\mathbf{X}'_1, \mathbf{A}_2^i), i \in \{0, 1, 2\}$
  - 12:    $\mathbf{W}_2 = [\mathbf{W}'_2[:, i]]_{i=0}^{C-1} \in \mathbb{R}^{N \times C}$  // class weights
  - 13:    $L_1 = \text{DGAT}.\text{cooperative\_loss}(\mathbf{S}, \mathbf{Y}, \mathbf{W}_2)$
  - 14:   Descend gradients of DGAT from  $L_1$
  - 15:   Update parameters of DGAT
  - 16: **end for**
  - 17: **return** DGAT
- 

### Comparative Experiments

We compared our method with three categories of networks across the two datasets. These include public vision networks (a three-layer MLP, VGG-19-BN (Simonyan and Zisserman 2015), ResNet-34 (He et al. 2016), and a small ViT (Dosovitskiy et al. 2020)), public GNNs (GCN (Kipf and Welling 2016), MA-GCNN (Peng et al. 2020), GATv2 (Brody, Alon, and Yahav 2021), ChebNetII (He, Wei, and Wen 2022), and GCN-MHSA (Chen et al. 2024a)), and ND-specific methods (BrainNetCNN (Kawahara et al. 2017), BrainGNN (Li et al. 2021), BrainNet-T (Kan et al. 2022b), LG-GNN (Zhang et al. 2022), and RA-GCN (Ghorbani et al. 2022)). Since the comparison networks cannot handle the full high-dimensional input, we provided them with core metrics—three 2D deterministic matrices—to ensure compatibility. Specifically, networks based on 2D input received three vertically stacked matrices. For 1D-input networks, we followed the practices of Song et al. (2022) and Huang et al. (2024), i.e., extracting the upper triangular elements of each symmetric matrix and then concatenating them into a 1D vector. The phenotypic graph was adopted for all public GNNs that rely on an adjacency graph but do not define a specific construction method. The input formats and other settings of ND-specific methods were kept as their original design. All models were trained until they approximately reached their minimum loss, with training capped at 500 epochs. Finally, we reported the average of their best evaluation performance across all folds, as presented in Table 1.

From Table 1, we observe that our method outperforms the second-best network by 7.57% on HC vs. PRO vs. PD and by 4.62% on CN vs. EMCI vs. AD in terms of accuracy. The best-performing vision network, GNN, and ND method are ViT-small, ChebNetII, and LG-GNN, respectively. Among them, ViT-small and ChebNetII achieve comparable performance, both surpassing LG-GNN. Notably,

Methods or Networks	HC vs. PRO vs. PD				CN vs. EMCI vs. AD			
	ACC(%)	BA(%)	F1(%)	SPE(%)	ACC(%)	BA(%)	F1(%)	SPE(%)
BrainGNN	64.15±3.96	41.02±5.45	54.83±6.56	70.51±2.72	56.90±3.10	44.15±8.19	53.76±6.67	72.08±4.10
BrainNet-T	63.52±3.47	38.10±5.44	53.11±8.53	69.05±2.72	56.47±3.29	40.86±2.57	54.04±2.57	70.43±1.29
BrainNetCNN	61.79±3.09	37.92±6.31	53.01±9.11	68.96±3.15	53.66±8.73	38.13±7.16	50.95±8.12	69.07±3.58
ChebNetII	66.99±4.13	44.01±5.74	58.17±5.57	72.00±2.87	62.90±2.76	50.15±8.34	58.36±5.57	75.07±4.17
GATv2	65.08±3.52	40.98±7.06	54.79±5.72	70.49±3.53	62.07±2.82	48.39±4.95	58.86±5.49	74.19±2.48
GCN	63.96±3.93	38.97±4.12	53.60±5.57	69.49±2.06	59.41±3.97	41.29±4.78	53.50±7.74	70.64±2.39
GCN-MHSA	61.54±4.21	33.33±0.00	46.97±5.18	66.67±0.00	50.86±3.22	34.19±2.75	40.03±4.83	67.10±1.38
LG-GNN	66.04±4.80	42.61±5.59	56.92±7.44	71.30±2.80	61.64±3.56	43.56±3.36	58.08±3.88	71.78±1.68
MA-GCNN	50.69±12.23	37.50±3.96	44.12±13.46	68.75±1.98	53.28±3.49	36.64±3.71	40.45±7.07	68.32±1.85
MLP	63.26±3.98	39.08±5.26	52.78±5.53	69.54±2.63	56.69±4.84	40.08±5.60	47.67±10.45	70.04±2.80
RA-GCN	61.54±4.21	33.33±0.00	46.97±5.18	66.67±0.00	55.07±3.92	37.02±4.30	44.12±9.12	68.51±2.15
ResNet-34	62.57±3.84	39.20±3.76	53.74±5.83	69.60±1.88	54.42±4.78	40.56±4.77	49.03±7.30	70.28±2.39
VGG-19-BN	63.96±2.81	38.61±3.73	53.51±2.55	69.31±1.87	60.64±4.28	42.41±3.70	55.93±5.21	71.21±1.85
ViT-small	<u>66.99±3.23</u>	<u>45.65±5.53</u>	<u>59.54±4.61</u>	<u>72.82±2.77</u>	<u>64.03±3.93</u>	46.71±6.23	<u>59.62±6.52</u>	73.35±3.11
<b>DW-DGAT</b>	<b>74.56±5.99</b>	<b>59.31±8.73</b>	<b>70.57±7.31</b>	<b>79.66±4.36</b>	<b>68.65±4.35</b>	<b>66.18±9.48</b>	<b>66.79±5.23</b>	<b>83.09±4.74</b>

Table 1: Evaluation scores of all methods and networks based on ten-fold cross-validation. Bold font indicates the best scores, while underlined font indicates the second-best.

Modules	HC vs. PRO vs. PD				CN vs. EMCI vs. AD			
	ACC(%)	BA(%)	F1(%)	SPE(%)	ACC(%)	BA(%)	F1(%)	SPE(%)
baseline	63.05±3.62	36.15±5.09	50.64±6.86	68.08±2.54	56.03±3.23	40.58±4.46	52.13±6.50	70.29±2.23
DF	65.41±3.84	44.18±6.97	58.78±9.13	72.09±3.48	60.78±3.94	43.90±5.05	58.11±6.74	71.95±2.53
DF+SGA	67.45±3.63	46.62±4.25	61.50±5.45	73.31±2.12	61.42±4.29	47.27±7.35	59.48±9.80	73.63±3.67
DF+SGA+GGA	71.70±5.30	55.59±9.58	68.42±10.16	77.79±4.79	65.09±3.14	64.31±9.97	64.66±10.44	82.15±4.99
<b>complete</b>	<b>74.56±5.99</b>	<b>59.31±8.73</b>	<b>70.57±7.31</b>	<b>79.66±4.36</b>	<b>68.65±4.35</b>	<b>66.18±9.48</b>	<b>66.79±5.23</b>	<b>83.09±4.74</b>

Table 2: Ablation experiments on DW-DGAT. The baseline is a basic two-layer MLP based on three deterministic networks.

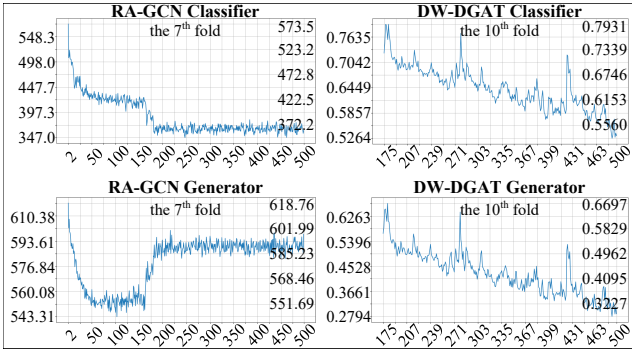


Figure 2: Loss histories of RA-GCN and DW-DGAT.

RA-GCN suffers from overfitting on HC vs. PRO vs. PD. To further investigate this, we compared the training loss histories of RA-GCN and DW-DGAT in their best-performing folds, as illustrated in Figure 2. The figure reveals that the adversarial training architecture of RA-GCN hampers effective convergence, with its loss stagnating at a high level. In contrast, the loss of DW-DGAT gradually decreases, demonstrating the stability and effectiveness of the proposed CWG.

To visualize the AUCs of all compared networks, we plotted their ROC curves. As shown in Figure 3, our ROC

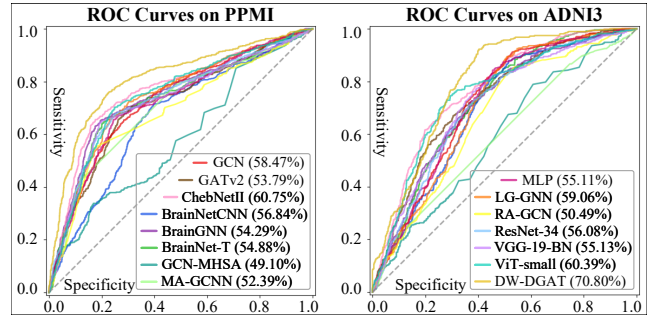


Figure 3: ROC curves and AUCs of all networks.

curves consistently dominate all others. Figure 4 presents the t-SNE visualizations of ViT-small, ChebNetII, and DW-DGAT. While DW-DGAT exhibits relatively less distinct inter-class boundaries, it achieves the most compact intra-class clusters, and identifies the largest number of samples from the minority classes.

### Ablation Experiments

We conducted ablation experiments to investigate the contributions of each module in the proposed methodology, as shown in Table 2. Notably, the most contributive module is

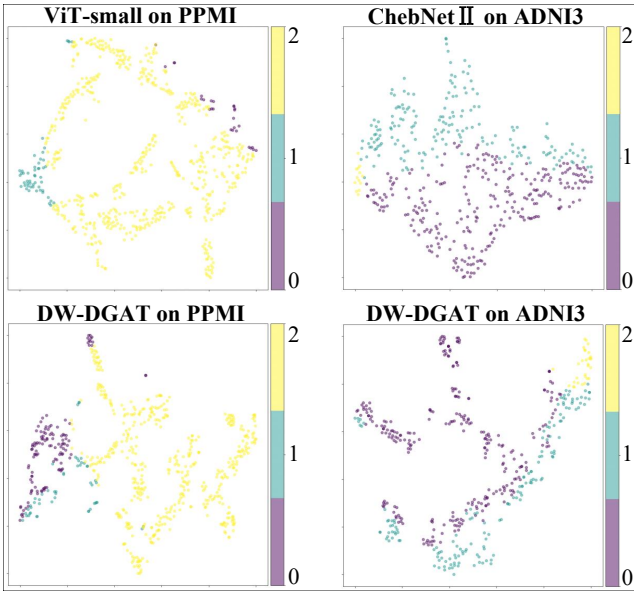


Figure 4: The t-SNE visualization results of three networks on two datasets, where class-0 in purple indicates normal.

GGA, which improves accuracies by 4.25% on HC vs. PRO vs. PD and 3.67% on CN vs. EMCI vs. AD, totaling 7.92%. The second most contributive module is DF, which improves accuracies by 2.36% and 4.75% on the two tasks, respectively, totaling 7.11%. We attribute the success of GGA to its graph construction method and the MHSA-GC layers, which jointly enhance the relevance between phenotypic features and sample labels, enabling more precise message propagation. Specifically, the cosine similarity in graph construction mitigates the effects of large range variances in phenotypic features and normalizes the feature range to  $[0, 1]$ . This ensures that the  $\sigma$  does not deviate excessively from the extreme values on either end, providing a balanced compromise between extremes (Gretton et al. 2006). Moreover, since the phenotypic features may not strictly align with the underlying sample relationships, the MHSA-GC layers adaptively adjust these relationships during neighbor aggregation, yielding more accurate neighborhood structures. In summary, the progressively increasing evaluation scores demonstrate the effectiveness of the proposed modules.

## Discussion

### Temporal and Spatial Complexity Analysis

Based on the notations defined in the Methodology section, we analyze the temporal and spatial complexities of the prediction computations in the DW-DGAT classifier, which are crucial considerations for clinical applications.

**Temporal Complexity Analysis.** The temporal complexity of the embedding layer is  $O(N \cdot R \cdot F' \cdot E)$ . The temporal complexities of the MHSA and the feed-forward (FF) layer in the ViT encoder are  $O(N \cdot (R + 1)^2 \cdot E)$  and  $O(N \cdot (R + 1) \cdot E^2)$ , respectively. Therefore, for  $L =$

12 encoder layers, the temporal complexity of the SGA is  $O(N \cdot E \cdot (R \cdot F' + (R + 1)^2 + (R + 1) \cdot E))$ . The temporal complexities of the first and the second MHSA-GC layers are  $O(N^3 \cdot E \cdot 2E)$  and  $O(N^3 \cdot 2E \cdot 4E)$ , respectively. The temporal complexity of the last linear layer in the GGA is  $O(N \cdot 4E \cdot 2E)$ . Thus, the temporal complexity of the GGA is  $O(N^3 \cdot E^2)$ . In total, the temporal complexity is  $O(N \cdot E \cdot (R \cdot F' + (R + 1)^2 + (R + 1 + N^2) \cdot E))$ . While the total complexity contains an  $O(N^3)$  term, it reduces to  $O(E \cdot (R \cdot F' + (R + 1)^2 + (R + 2) \cdot E)) = O(E \cdot R \cdot (F' + R + E))$  when the batch size  $N$  is fixed. Under the reported experimental settings, the classifier requires 139.02 GFLOPs, while the generator requires 169.29 GFLOPs.

**Spatial Complexity Analysis.** The spatial complexity of the embedding layer is  $O(N \cdot R \cdot (F' + E))$ . The spatial complexities of the MHSA and the FF layer are  $O(N \cdot H' \cdot (1 + R) \cdot (1 + R + E))$  and  $O(N \cdot (1 + R) \cdot E)$ , respectively, where  $H'$  is the number of attention heads of the ViT encoder. Hence, for  $L$  encoder layers and  $F' < R \ll E$ , the spatial complexity of the SGA is  $O(N \cdot H' \cdot (1 + R) \cdot (1 + R + E))$ . The spatial complexities of the first and the second MHSA GC layers are  $O(N \cdot H \cdot N \cdot (N + E))$  and  $O(N \cdot H \cdot N \cdot (N + 2E))$ , respectively. The spatial complexity of the last FC layer in the GGA is  $O(N \cdot 2E)$ . Thus, the spatial complexity of the GGA is  $O(N^2 \cdot H \cdot (N + 2E))$ . Since  $H < H' \leq 2H$ , the spatial complexity is  $O(N \cdot H \cdot ((1 + R) \cdot (1 + R + E) + N \cdot (N + 2E)))$  in total. While the total complexity contains an  $O(N^3)$  item, it reduces to  $O(H \cdot ((1 + R) \cdot (1 + R + E) + 2E)) = O(H \cdot R \cdot (R + E))$  when the batch size  $N$  is fixed. Under the reported experimental settings, the classifier consumes approximately 2728 MB of GPU memory, whereas the generator consumes approximately 3288 MB.

## Conclusion

In this paper, we proposed a data fusion strategy to integrate multi-structured, high-dimensional metrics, providing a solid foundation for subsequent feature extraction. Next, we introduced DW-DGAT, a dual graph attention network with a class weight generator for ND diagnosis, which effectively captures reduced ROI graph features and fuses sample phenotypic information at both micro and macro levels, while mitigating class imbalance. A series of comparative and ablation experiments based on two datasets coming from the PPMI and the ADNI3 data sources demonstrated the superiority and generalizability of our proposed method, particularly the GGA module.

Nevertheless, our network is not without limitations. The 3D data fusion strategy does not account for the large differences in voxel numbers across ROIs, potentially neglecting latent features in larger ROIs.

## Acknowledgments

This work was supported partly by National Natural Science Foundation of China (No. 62276172, 62306106), National Natural Science Foundation of Guangdong Province (No. 2023A1515011378), Shenzhen Key Basic Research

Project (No. JCYJ20230808105602005), and China Scholarship Council.

## References

- Arjovsky, M.; and Bottou, L. 2017. Towards principled methods for training generative adversarial networks. *arXiv preprint arXiv:1701.04862*.
- Balakrishnan, A.; Medikonda, J.; Namboothiri, P. K.; and Natarajan, M. 2023. Mahalanobis metric-based oversampling technique for Parkinson's disease severity assessment using spatiotemporal gait parameters. *Biomedical Signal Processing and Control*, 86: 105057.
- Balestrino, R.; and Schapira, A. 2020. Parkinson disease. *European Journal of Neurology*, 27(1): 27–42.
- Bi, X.-A.; Yang, Z.; Huang, Y.; Xing, Z.; Xu, L.; Wu, Z.; Liu, Z.; Li, X.; and Liu, T. 2024. CE-GAN: Community evolutionary generative adversarial network for Alzheimer's disease risk prediction. *IEEE Transactions on Medical Imaging*, 43(11): 3663–3675.
- Bourriez, N.; Bendidi, I.; Cohen, E.; Watkinson, G.; Sanchez, M.; Bollot, G.; and Genovesio, A. 2024. ChAdaViT: Channel adaptive attention for joint representation learning of heterogeneous microscopy images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 11556–11565.
- Brody, S.; Alon, U.; and Yahav, E. 2021. How attentive are graph attention networks? *arXiv preprint arXiv:2105.14491*.
- Chen, J.; Xie, H.; Cai, S.; Song, L.; Geng, B.; and Guo, W. 2024a. GCN-MHSA: A novel malicious traffic detection method based on graph convolutional neural network and multi-head self-attention mechanism. *Computers & Security*, 147: 104083.
- Chen, L.; Zhou, Y.; Zhang, X.; Zhang, Z.; and Zheng, H. 2024b. MvRNA: A new multi-view deep neural network for predicting Parkinson's disease. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 93–104.
- Cobbinah, B. M.; Sorg, C.; Yang, Q.; Ternblom, A.; Zheng, C.; Han, W.; Che, L.; and Shao, J. 2022. Reducing variations in multi-center Alzheimer's disease classification with convolutional adversarial autoencoder. *Medical Image Analysis*, 82: 102585.
- Cui, Z.; Zhong, S.; Xu, P.; He, Y.; and Gong, G. 2013. PANDA: A pipeline toolbox for analyzing brain diffusion images. *Frontiers in Human Neuroscience*, 7: 42.
- Ding, J.-E.; Hsu, C.-C.; and Liu, F. 2024. Parkinson's disease classification using contrastive graph cross-view learning with multimodal fusion of SPECT images and clinical features. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissensborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Ghorbani, M.; Kazi, A.; Baghshah, M. S.; Rabiee, H. R.; and Navab, N. 2022. RA-GCN: Graph convolutional network for disease prediction problems with imbalanced data. *Medical Image Analysis*, 75: 102272.
- Gong, G. 2013. Local diffusion homogeneity (LDH): An inter-voxel diffusion MRI metric for assessing inter-subject white matter variability. *PLoS One*, 8(6): e66366.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Gretton, A.; Borgwardt, K.; Rasch, M.; Schölkopf, B.; and Smola, A. 2006. A kernel method for the two-sample-problem. *Advances in Neural Information Processing Systems*, 19: 513–520.
- Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. C. 2017. Improved training of Wasserstein GANs. *Advances in Neural Information Processing Systems*, 30: 5769–5779.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- He, M.; Wei, Z.; and Wen, J.-R. 2022. Convolutional neural networks on graphs with Chebyshev approximation, revisited. *Advances in Neural Information Processing Systems*, 35: 7264–7276.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*.
- Huang, X.; He, Q.; Ruan, X.; Li, Y.; Kuang, Z.; Wang, M.; Guo, R.; Bu, S.; Wang, Z.; Yu, S.; et al. 2024. Structural connectivity from DTI to predict mild cognitive impairment in de novo Parkinson's disease. *NeuroImage: Clinical*, 41: 103548.
- Huang, Z.; Wang, K.; Chen, C.; Chen, J.; Wan, J.; Yang, Z.; Zhou, R.; and Gan, H. 2025. Incomplete multi-view data learning via adaptive embedding and partial  $L_{\{2,1\}}$  norm constraints for Parkinson's disease diagnosis. *IEEE Journal of Biomedical and Health Informatics*, 10.1109/JBHI.2025.3576786.
- Kan, X.; Cui, H.; Lukemire, J.; Guo, Y.; and Yang, C. 2022a. FBNetGen: Task-aware GNN-based fMRI analysis via functional brain network generation. In *International Conference on Medical Imaging with Deep Learning*, 618–637.
- Kan, X.; Dai, W.; Cui, H.; Zhang, Z.; Guo, Y.; and Yang, C. 2022b. Brain network transformer. *Advances in Neural Information Processing Systems*, 35: 25586–25599.
- Kavitha, M.; and Kasthuri, M. 2024. Enhanced cost-sensitive ensemble learning for imbalanced class in medical data. *J. Electrical Systems*, 20(7s): 1043–1053.
- Kawahara, J.; Brown, C. J.; Miller, S. P.; Booth, B. G.; Chau, V.; Grunau, R. E.; Zwicker, J. G.; and Hamarneh, G. 2017. BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage*, 146: 1038–1049.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.

- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Li, X.; Zhou, Y.; Dvornik, N.; Zhang, M.; Gao, S.; Zhuang, J.; Scheinost, D.; Staib, L. H.; Ventola, P.; and Duncan, J. S. 2021. BrainGNN: Interpretable brain graph neural network for fMRI analysis. *Medical Image Analysis*, 74: 102233.
- Liu, J.; Du, H.; Guo, R.; Bai, H. X.; Kuang, H.; and Wang, J. 2022. MMGK: Multimodality multiview graph representations and knowledge embedding for mild cognitive impairment diagnosis. *IEEE Transactions on Computational Social Systems*, 11(1): 389–398.
- Marek, K.; Chowdhury, S.; Siderowf, A.; Lasch, S.; Coffey, C. S.; Caspell-Garcia, C.; Simuni, T.; Jennings, D.; Tanner, C. M.; Trojanowski, J. Q.; et al. 2018. The Parkinson’s progression markers initiative (PPMI)—establishing a PD biomarker cohort. *Annals of Clinical and Translational Neurology*, 5(12): 1460–1477.
- Mayo, C. D.; Garcia-Barrera, M. A.; Mazerolle, E. L.; Ritchie, L. J.; Fisk, J. D.; Gawryluk, J. R.; and Initiative, A. D. N. 2019. Relationship between DTI metrics and cognitive function in Alzheimer’s disease. *Frontiers in Aging Neuroscience*, 10: 436.
- Peng, H.; Li, J.; Gong, Q.; Ning, Y.; Wang, S.; and He, L. 2020. Motif-matching based subgraph-level attentional convolutional network for graph classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 5387–5394.
- Poewe, W.; Seppi, K.; Tanner, C. M.; Halliday, G. M.; Brundin, P.; Volkman, J.; Schrag, A.-E.; and Lang, A. E. 2017. Parkinson disease. *Nature Reviews Disease Primers*, 3(1): 1–21.
- Roy, S.; Roy, U.; Sinha, D.; and Pal, R. K. 2023. Imbalanced ensemble learning in determining Parkinson’s disease using keystroke dynamics. *Expert Systems with Applications*, 217: 119522.
- Scheltens, P.; De Strooper, B.; Kivipelto, M.; Holstege, H.; Chételat, G.; Teunissen, C. E.; Cummings, J.; and van der Flier, W. M. 2021. Alzheimer’s disease. *The Lancet*, 397(10284): 1577–1590.
- Shehzad, A.; Zhang, D.; Yu, S.; Abid, S.; Kumar, D. K.; and Xia, F. 2025. Multiscale graph transformer for brain disorder diagnosis. *IEEE Transactions on Consumer Electronics*, 71(2): 5535–5546.
- Simonyan, K.; and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. *3rd International Conference on Learning Representations (ICLR 2015)*, 1–14.
- Smith, S. M.; Jenkinson, M.; Woolrich, M. W.; Beckmann, C. F.; Behrens, T. E.; Johansen-Berg, H.; Bannister, P. R.; De Luca, M.; Drobnjak, I.; Flitney, D. E.; et al. 2004. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, 23: S208–S219.
- Song, X.; Shu, K.; Yang, P.; Zhao, C.; Zhou, F.; Frangi, A. F.; Xiao, X.; Dong, L.; Wang, T.; Wang, S.; et al. 2025. Knowledge-aware multisite adaptive graph transformer for brain disorder diagnosis. *IEEE Transactions on Medical Imaging*, 44(6): 2370–2383.
- Song, X.; Zhou, F.; Frangi, A. F.; Cao, J.; Xiao, X.; Lei, Y.; Wang, T.; and Lei, B. 2022. Multicenter and multichannel pooling GCN for early AD diagnosis based on dual-modality fused brain network. *IEEE Transactions on Medical Imaging*, 42(2): 354–367.
- Weiner, M. W.; Veitch, D. P.; Aisen, P. S.; Beckett, L. A.; Cairns, N. J.; Green, R. C.; Harvey, D.; Jack Jr, C. R.; Jagust, W.; Morris, J. C.; et al. 2017. The Alzheimer’s disease neuroimaging initiative 3: Continued innovation for clinical trial improvement. *Alzheimer’s & Dementia*, 13(5): 561–571.
- Xu, J.; Bian, Q.; Li, X.; Zhang, A.; Ke, Y.; Qiao, M.; Zhang, W.; Sim, W. K. J.; and Gulyás, B. 2024. Contrastive graph pooling for explainable classification of brain networks. *IEEE Transactions on Medical Imaging*, 43(9): 3292–3305.
- Xu, M.; Sanz, D. L.; Garces, P.; Maestu, F.; Li, Q.; and Pantazis, D. 2021. A graph Gaussian embedding method for predicting Alzheimer’s disease progression with MEG brain networks. *IEEE Transactions on Biomedical Engineering*, 68(5): 1579–1588.
- Yang, M.; Huang, X.; Huang, L.; and Cai, G. 2023. Diagnosis of Parkinson’s disease based on 3D ResNet: The frontal lobe is crucial. *Biomedical Signal Processing and Control*, 85: 104904.
- Ying, C.; Cai, T.; Luo, S.; Zheng, S.; Ke, G.; He, D.; Shen, Y.; and Liu, T.-Y. 2021. Do transformers really perform badly for graph representation? *Advances in Neural Information Processing Systems*, 34: 28877–28888.
- Yu, S.; Jin, S.; Li, M.; Sarwar, T.; and Xia, F. 2024. Long-range brain graph transformer. *Advances in Neural Information Processing Systems*, 37: 24472–24495.
- Yue, G.; Wei, P.; Zhou, T.; Song, Y.; Zhao, C.; Wang, T.; and Lei, B. 2024. Specificity-aware federated learning with dynamic feature fusion network for imbalanced medical image classification. *IEEE Journal of Biomedical and Health Informatics*, 28(11): 6373–6383.
- Zhang, H.; Song, R.; Wang, L.; Zhang, L.; Wang, D.; Wang, C.; and Zhang, W. 2022. Classification of brain disorders in rs-fMRI via local-to-global graph neural networks. *IEEE Transactions on Medical Imaging*, 42(2): 444–455.
- Zhang, S.; Chen, X.; Shen, X.; Ren, B.; Yu, Z.; Yang, H.; Jiang, X.; Shen, D.; Zhou, Y.; and Zhang, X.-Y. 2023a. A-GCL: Adversarial graph contrastive learning for fMRI analysis to diagnose neurodevelopmental disorders. *Medical Image Analysis*, 90: 102932.
- Zhang, Y.; Cao, G.; Xu, X.; and Kang, G. 2023b. Graph-based Alzheimer’s disease diagnosis with contrastive learning and graph transformer. In *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2368–2373.
- Zhou, J.; Jie, B.; Wang, Z.; Zhang, Z.; Du, T.; Bian, W.; Yang, Y.; and Jia, J. 2024. LCGNet: Local sequential feature coupling global representation learning for functional connectivity network analysis with fMRI. *IEEE Transactions on Medical Imaging*, 43(12): 4319–4330.