

ClinGRAD: Clinically-Guided Genomics and Radiomics Interpretable GNN for Dementia Diagnosis

Salma Hassan¹^[0009-0006-5498-8319], Mostafa Salem^{1,2}^[0000-0001-9915-8390],
Vijay Ram Kumar Papineni³^[0000-0002-6162-3290], Ayman
Elsayed³^[0000-0003-3381-066X], and Mohammad Yaqub¹^[0000-0001-6896-1105]

¹ Mohamed Bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE
{Salma.Hassan, Mostafa.Salem, Mohammad.Yaqub}@mbzuai.ac.ae

² Department of Computer Science, Faculty of Computers and Information, Assiut
University, Assiut, Egypt

³ Sheikh Shakhboub Medical City, Abu Dhabi, UAE
rkpvijay@gmail.com, ayaahmed@ssmc.ae

Abstract. Alzheimer’s Disease (AD) remains a major diagnostic challenge due to the complex interplay of genomic, radiomic, and structural factors in disease progression. While deep learning methods can classify AD, current approaches fail to effectively combine multimodal data with clinical knowledge, compromising both accuracy and interpretability. We present ClinGRAD, a clinically-guided heterogeneous graph neural network that combines genomic and radiomic data using connections based on diffusion-weighted imaging (DWI) maps and gene co-expression networks. ClinGRAD’s contributions include: (1) a multimodal fusion architecture that integrates validated structural and genetic connectivity patterns for consistent biological feature analysis; (2) a multi-scale graph framework capturing both local brain structure and global genomic pathway relationships; (3) an attention mechanism that provides clinically relevant explanations of gene-structure interactions; and (4) pathway-based gene clustering that reveals underlying biological mechanisms and their clinical implications. ClinGRAD outperforms existing models, achieving an accuracy of 93.15%, distinguishing AD from control, mild cognitive impaired, and vascular dementia patients while maintaining biological coherence through its clinical guidance framework. The code is available at <https://github.com/BioMedIA-MBZUAI/ClinGRAD>.

Keywords: Graph Neural Network · Dementia · Neuroimaging · Genomics · Multi-omics · Interpretability · DWI

1 Introduction

Alzheimer’s Disease (AD) is one of the most pressing healthcare challenges, with its devastating impact on both individuals and society growing rapidly. Currently, 75% of individuals with dementia worldwide remain undiagnosed, and

cases are projected to surge to 139 million by 2050 [12]. Beyond the immediate health impact, the economic burden is severe, with global costs exceeding \$1.3 trillion annually and expected to increase nine-fold by 2050 [14].

Understanding AD’s progression requires integrating multiple biological perspectives. Clinical research has shown that brain connectivity maps from diffusion-weighted imaging (DWI), a specialized MRI technique that tracks water molecule movement along neural pathways, provide crucial insights into disease progression [9]. At the molecular level, gene co-expression networks, which capture how genes are activated together across different brain regions, have revealed key disease mechanisms and potential therapeutic targets [11]. These clinically validated approaches suggest that effective AD diagnosis must consider both structural brain changes and genetic interactions simultaneously.

Traditional machine learning approaches have attempted to address AD diagnosis through various modalities: genomic analysis to identify risk-associated variants [11], structural MRI to detect anatomical changes [27], and clinical assessments to track cognitive decline [7]. However, these single-modality approaches fail to capture AD’s complex pathophysiology, which spans multiple biological scales. While deep learning has advanced multimodal integration [3], existing fusion techniques often overlook critical cross-modal interactions.

Graph Neural Networks (GNNs) have emerged as powerful tools for medical imaging analysis, particularly in neurological applications, where they excel at capturing spatial and relational patterns in brain data [8]. While advanced graph architectures have demonstrated success in neurodegenerative disease diagnosis [13,19], the integration of multimodal data remains challenging due to fundamental differences in data structure [25,18]. Recent innovations in multimodal GNNs and message-passing techniques have improved predictive performance [26], yet three critical limitations persist: (1) insufficient incorporation of established clinical knowledge about brain connectivity patterns, (2) inadequate modeling of cross-scale biological interactions, and (3) limited clinical interpretability of predictions. Despite the wealth of insights generated from anatomical parcellation studies [22] and connectivity mapping research, current deep learning models rarely leverage this clinical expertise. Even with efforts to standardize evaluation through benchmarks like NeuroGraph [17], bridging the gap between model predictions and clinical interpretability remains a significant challenge.

We introduce ClinGRAD, a novel clinically-guided heterogeneous GNN that addresses these challenges by integrating multimodal data with multi-scale graph representations to capture comprehensive AD pathology. Our contributions are:

1. **Clinically-Guided Multimodal Fusion:** ClinGRAD proposes a heterogeneous GNN that fuses MRI and genetic data through validated brain connectivity maps and gene co-expression, enforcing anatomical constraints during cross-modal learning.
2. **Multi-Scale Graph Representations:** ClinGRAD implements a hierarchical graph framework that models AD through coupled biological scales: molecular pathways, brain structures, and clinical manifestations.

3. **Clinical Interpretability:** ClinGRAD identifies key genes and brain regions influencing each diagnosis, providing physicians with clear evidence.
4. **Pathway-Based Gene Clustering:** ClinGRAD implements clinically-guided supernodes to cluster genes in an unsupervised manner based on known biological pathways and disease mechanisms.

2 Methodology

Feature Extraction and Pre-processing. To extract structural features, we segment brain MRI scans into 32 regions using SynthSeg [1], which outperforms traditional tools like FSL and FreeSurfer in handling elderly brain scans with atrophy [16]. From these segmentations, we extract 107 radiomic features per structure using PyRadiomics [23]. For genomic features, we process data for 75 AD-associated genes from [6]. Crucially, we integrate established structural connectivity patterns between brain regions derived from DWI studies [20]. These DWI-derived connectivity maps serve as biologically validated priors in our graph, ensuring message passing follows anatomical pathways. The connectivity weights between regions are derived from pairwise Pearson correlation coefficients of white matter fiber densities across subjects, as shown in Fig. 1C.

Heterogeneous Graph Structure. ClinGRAD employs a heterogeneous graph structure $G = (V, E)$ that models patients, genes, and brain structures as distinct node types within a clinically informed architecture. The node representa-




Fig. 1: Architecture for heterogeneous ClinGRAD model. The process involves three main steps: (A) segmenting MRI scans using a 3D UNet model and extracting radiomics features, (B) selecting relevant genetic data, and (C) DWI connection map. The combined data feeds into ClinGRAD (D) featuring gene-to-gene, structure-to-structure, and patient to structure and gene interactions.

tions capture key clinical markers through 1) patient nodes (V_p) with features $\mathbf{X}_p \in \mathbb{R}^{n_p \times d_p}$ including cognitive scores, 2) gene nodes (V_g) encoding expression values in $\mathbf{X}_g \in \mathbb{R}^{n_g \times 1}$, and 3) structure nodes (V_s) representing radiomic features in $\mathbf{X}_s \in \mathbb{R}^{n_s \times d_s}$, where p is the number of patients, g is the number of selected genes, and s is the number of structures. These nodes are connected through four biologically-informed edge types: patient-gene edges (E_{pg}) linking patients to expression profiles, patient-structure edges (E_{ps}) connecting imaging features, weighted gene edges (E_{gg}) by co-expression scores from GeneMANIA [28], and structure-structure edges (E_{ss}). For E_{ss} , we learn edge weights through a combination of anatomical and connectivity information: $w_{ij} = \text{MLP}([\alpha \cdot d_{ij}, \beta \cdot c_{ij}])$, where d_{ij} represents the scaled 3D Euclidean distance between regions i and j , c_{ij} denotes the corresponding DWI connectivity strength, and α, β are learnable parameters. The MLP aims to learn optimal edge representations that balance spatial proximity and structural connectivity. The graph, formally represented as $G = (V_p, V_g, V_s), (E_{pg}, E_{ps}, E_{gg}, E_{ss})$, uniquely integrates molecular and structural biomarkers and preserves clinically relevant relationships across modalities.

Message Passing and Aggregation. ClinGRAD implements a clinically informed message-passing framework using Graph Attention Networks (GAT) [24] to adaptively aggregate information across heterogeneous nodes while preserving clinical relationships. For each node type $\tau(i)$, we compute attention coefficients that reflect the importance of different node interactions:

$$\alpha_{ij}^\phi = \frac{\exp(\text{LeakyReLU}(\mathbf{a}_{\tau(i),\phi}^T [\mathbf{W}_{\tau(i),\phi} \mathbf{h}_i^{(t)} \parallel \mathbf{W}_{\tau(i),\phi} \mathbf{h}_j^{(t)}]))}{\sum_{k \in \mathcal{N}_\phi(i)} \exp(\text{LeakyReLU}(\mathbf{a}_{\tau(i),\phi}^T [\mathbf{W}_{\tau(i),\phi} \mathbf{h}_i^{(t)} \parallel \mathbf{W}_{\tau(i),\phi} \mathbf{h}_k^{(t)}]))}. \quad (1)$$

The node features are then updated through multi-head attention:


$$\mathbf{h}_i^{(t+1)} = \sigma \left(\frac{1}{K} \sum_{k=1}^K \left(\mathbf{w}_{\tau(i)}^k \mathbf{h}_i^{(t)} + \sum_{\phi \in \Phi} \sum_{j \in \mathcal{N}_\phi(i)} \alpha_{ij}^{\phi,k} \mathbf{w}_{\tau(i),\phi}^k \mathbf{h}_j^{(t)} \right) \right), \quad (2)$$

Where K is the number of attention heads, $\mathbf{W}_{\tau(i)}^k$ and $\mathbf{W}_{\tau(i),\phi}^k$ are learnable weight matrices for self and inter-type connections respectively, $\mathbf{a}_{\tau(i),\phi}$ is the attention vector for relation type ϕ , α_{ij}^ϕ is the attention coefficient between nodes i and j , $\mathbf{h}_i^{(t)}$ is node i 's feature vector at layer t , $\tau(i)$ is node i 's type, and $\mathcal{N}_\phi(i)$ represents neighbors of node i connected by relation ϕ . This attention mechanism dynamically weighs interactions based on learned patterns and clinical priors.

AD Classification. The final patient node embeddings are passed through a classification layer with softmax activation: $\hat{y}_i = \text{softmax}(\mathbf{W}_c \mathbf{h}_i^{(T)} + \mathbf{b}_c)$, where $\mathbf{W}_c, \mathbf{b}_c$ are learnable parameters and $\hat{y}_i$ represents the predicted AD status probability distribution. The class with the highest probability is the final prediction.



Fig. 2: Network visualization depicting gene clustering based on their functional associations within AD pathways. The supernodes (stars) represent distinct biological clusters. Genes (blue nodes) are connected within and across clusters.



(a) Each node represents a gene, with connections indicating significant interactions and color representing gene influence level.

(b) Each sphere represents the center of structures, and color represents the level of influence.

Fig. 3: Interpretability of ClinGRAD (a) Network visualization of gene interactions highlighting influential genes in AD pathogenesis. (b) Corresponding radiomics-based analysis of brain regions affected in AD.

3 Experimental Details

Datasets. We utilize ANMerge [2], the **only** public dataset that combines longitudinal MRI scans with comprehensive genomic data for dementia analysis. While ADNI lacks the necessary genomic depth and multi-class labels for our heterogeneous graph approach, ANMerge provides rich molecular profiles alongside structural MRI data for four distinct classes: AD, Vascular Dementia (VaD), Mild Cognitive Impairment (MCI), and Control (CTL). While we derived our initial structural connectivity priors from [20], we validated and refined these connection maps using DWI scans from the ethnically diverse BrainLat [15] dataset for both AD and CTL cohorts, ensuring robust connectivity estimates.

Experimental Setup. ClinGRAD is trained using a stratified 3-fold cross-




Fig. 4: Left: A chord diagram visualizing the high-level relationships between supernodes biological clusters and their degree of interplay. Right: Gene-level connections highlighting detailed interactions among genes.

validation approach to ensure a balanced representation of the classes. For training, the Adam optimizer ($\text{lr}=1e^{-3}$, $\text{weight decay}=5e^{-4}$) was used. Full hyperparameters are available in the code.

Ablation Study. We conduct comprehensive comparisons against both unimodal and multimodal approaches. For unimodal MRI analysis, we evaluate 3D CNN and 3D ViT architectures. In multimodal settings, we compare against recent methods, including MINDSETS [5], Flex-MOE [29], HSGO [21], and FT-Transformer [4]. To isolate ClinGRAD’s key components, we perform ablation studies on multimodal integration, graph edges, and message-passing mechanisms.

Interpretability. ClinGRAD provides interpretable insights through three key mechanisms: (1) GNNExplainer identifies critical subgraphs and features driving predictions, (2) Gene supernode clustering reveals functional groups based on co-expression patterns (Fig. 2), and (3) Cross-modal analysis maps relationships between genetic pathways and structural changes in AD-critical regions (Fig. 3). These complementary approaches enable clinicians to trace model decisions back to specific biological mechanisms, offering transparent reasoning for diagnosis.

4 Results and Discussion

Comparative Performance Analysis. Our extensive experimentation, detailed in Tables 1 and 2, demonstrates ClinGRAD’s superior performance across multiple AD-related classification tasks. In the AD vs CTL classification, ClinGRAD achieves 98.75% accuracy, outperforming both traditional single-modality approaches (3D CNN: 84.27%, 3D ViT: 87.12%) and recent multimodal methods (MINDSETS [5]: 97.92%). The consistent performance across these diverse

Table 1: Performance comparison of different models across various metrics and classification tasks across ANMerge datasets with varying data types. Hyphens indicate unreported metrics. The standard deviation for [10,21] was not reported.

Model	Data	Metric	AD Vs CTL	AD Vs MCI	MCI Vs CTL	AD Vs VaD
3D CNN	MRI	Accuracy	84.27 $\pm$ 5.32	72.01 $\pm$ 5.68	77.73 $\pm$ 4.52	72.83 $\pm$ 5.62
		F1-Score	80.50 $\pm$ 4.85	70.96 $\pm$ 5.12	77.89 $\pm$ 3.98	69.91 $\pm$ 4.51
		Recall	79.67 $\pm$ 5.02	70.08 $\pm$ 4.21	77.77 $\pm$ 4.09	69.56 $\pm$ 3.84
		Precision	81.34 $\pm$ 4.85	71.87 $\pm$ 5.20	78.02 $\pm$ 4.65	70.26 $\pm$ 4.76
3D ViT	MRI	Accuracy	87.12 $\pm$ 3.12	75.43 $\pm$ 3.45	79.95 $\pm$ 3.87	75.62 $\pm$ 4.23
		F1-Score	83.48 $\pm$ 4.09	73.88 $\pm$ 4.12	79.12 $\pm$ 4.05	72.45 $\pm$ 4.18
		Recall	81.67 $\pm$ 3.76	72.91 $\pm$ 3.98	78.89 $\pm$ 3.92	71.98 $\pm$ 3.95
		Precision	83.31 $\pm$ 4.87	74.85 $\pm$ 4.21	79.35 $\pm$ 4.15	72.92 $\pm$ 4.08
Maddalena et al. [10]	MRI + Gen	Accuracy	87.40	72.20	77.80	-
		F1-Score	87.20	71.00	78.60	-
		Recall	87.20	72.00	79.00	-
		Precision	87.90	71.00	76.50	-
Flex-MOE [29]	MRI + Gen	Accuracy	92.92 $\pm$ 1.85	88.75 $\pm$ 1.82	84.95 $\pm$ 2.15	84.85 $\pm$ 2.05
		F1-Score	92.85 $\pm$ 1.92	88.42 $\pm$ 1.85	83.88 $\pm$ 2.22	83.92 $\pm$ 2.12
		Recall	92.76 $\pm$ 1.88	88.28 $\pm$ 1.87	84.25 $\pm$ 2.18	83.45 $\pm$ 2.08
		Precision	92.94 $\pm$ 1.90	88.56 $\pm$ 1.84	83.52 $\pm$ 2.20	84.39 $\pm$ 2.15
HSGO [21]	Rad + Gen	Accuracy	87.66	-	71.53	-
		F1-Score	80.36	-	63.96	-
		Recall	85.73	-	63.48	-
		Precision	83.65	-	69.46	-
SVM	Rad + Gen	Accuracy	80.46 $\pm$ 3.70	61.47 $\pm$ 3.90	69.90 $\pm$ 3.95	81.71 $\pm$ 4.29
		F1-Score	80.89 $\pm$ 3.10	52.55 $\pm$ 3.85	69.85 $\pm$ 3.81	85.31 $\pm$ 5.26
		Recall	83.72 $\pm$ 4.30	52.99 $\pm$ 7.24	69.89 $\pm$ 5.21	91.72 $\pm$ 4.32
		Precision	78.26 $\pm$ 3.70	53.11 $\pm$ 3.36	69.83 $\pm$ 3.75	79.76 $\pm$ 3.85
MINDSETS [5]	Rad + Gen	Accuracy	97.92 $\pm$ 2.15	90.69 $\pm$ 1.65	86.45 $\pm$ 2.87	89.25 $\pm$ 5.62
		F1-Score	97.92 $\pm$ 3.06	90.68 $\pm$ 1.87	84.23 $\pm$ 3.84	81.13 $\pm$ 4.58
		Recall	97.91 $\pm$ 2.54	90.69 $\pm$ 1.85	82.63 $\pm$ 4.02	83.61 $\pm$ 3.89
		Precision	98.92 $\pm$ 3.20	90.85 $\pm$ 1.87	87.24 $\pm$ 4.11	79.21 $\pm$ 3.88
FT-Transformer [4]	Rad + Gen	Accuracy	88.45 $\pm$ 2.25	82.75 $\pm$ 1.98	79.92 $\pm$ 2.75	80.85 $\pm$ 2.45
		F1-Score	88.12 $\pm$ 2.31	82.42 $\pm$ 2.05	78.85 $\pm$ 2.92	79.92 $\pm$ 2.52
		Recall	87.98 $\pm$ 2.28	82.28 $\pm$ 2.12	78.25 $\pm$ 2.88	79.45 $\pm$ 2.48
		Precision	88.26 $\pm$ 2.35	82.56 $\pm$ 2.08	79.45 $\pm$ 2.82	80.39 $\pm$ 2.55
GCN GNN	Rad + Gen	Accuracy	91.79 $\pm$ 3.62	89.27 $\pm$ 2.65	84.46 $\pm$ 2.36	85.31 $\pm$ 3.48
		F1-Score	91.90 $\pm$ 3.42	90.44 $\pm$ 2.66	<u>84.38 $\pm$ 2.38</u>	<u>86.22 $\pm$ 3.48</u>
		Recall	92.02 $\pm$ 2.78	<u>91.64 $\pm$ 2.92</u>	<u>84.63 $\pm$ 2.63</u>	87.15 $\pm$ 3.77
		Precision	91.79 $\pm$ 2.98	89.27 $\pm$ 3.03	84.14 $\pm$ 2.81	<u>85.31 $\pm$ 3.48</u>
ClinGRAD (Ours)	Rad + Gen	Accuracy	98.75 $\pm$ 1.54	94.25 $\pm$ 1.55	89.66 $\pm$ 2.42	89.45 $\pm$ 1.96
		F1-Score	98.71 $\pm$ 1.64	92.43 $\pm$ 1.73	87.93 $\pm$ 1.91	88.54 $\pm$ 2.03
		Recall	98.65 $\pm$ 1.60	92.55 $\pm$ 1.72	89.43 $\pm$ 1.85	87.82 $\pm$ 1.91
		Precision	<u>98.78 $\pm$ 1.47</u>	93.25 $\pm$ 1.79	<u>86.42 $\pm$ 1.85</u>	89.73 $\pm$ 1.92

tasks, coupled with low standard deviations, suggests that ClinGRAD effectively captures subtle disease markers that differentiate these conditions.

Multimodal Integration Benefits. The ablation studies in Table 2 reveal crucial insights into the effectiveness of our multimodal architecture. Using radiomics features alone achieves moderate accuracy, while genomic features independently show stronger discriminative power (89.52%). The integration of both modalities, even without structural connections, yields a substantial improvement to 90.80% accuracy, representing a significant gain over single-modality approaches. This improvement validates our hypothesis that radiomics and genomics capture complementary disease signatures. The inclusion of DWI connections further enhances performance, pushing accuracy from 92.05% to 93.15%.

Impact of Edge Connections. Our systematic evaluation of different edge

Table 2: Breakdown of ClinGRAD results to evaluate the impact of data types, edge connections, and corresponding performance metrics. The filled circle indicates inclusion, and the empty circle indicates exclusion.

Task	Data		Edge Connections			Classes				Metrics			
	Gen	Rad	Struct	DWI	Co-Exp	CTL	MCI	AD	VaD	Accuracy (%)	F1-Score (%)	Precision (%)	Recall (%)
Binary	○	●	○	○	○	●	○	●	○	66.21 ± 4.18	67.31 ± 4.20	68.45 ± 5.88	66.21 ± 4.18
	○	●	●	○	○	○	○	●	○	77.13 ± 2.14	78.37 ± 1.68	79.65 ± 3.97	77.13 ± 2.14
	○	●	●	●	○	○	○	●	○	79.45 ± 2.15	80.42 ± 1.70	81.42 ± 3.95	79.45 ± 2.15
	●	○	○	○	○	●	○	●	○	85.49 ± 3.49	86.44 ± 3.58	87.42 ± 3.82	85.49 ± 3.49
	●	○	○	○	○	○	○	●	○	89.52 ± 3.65	90.66 ± 3.70	91.85 ± 2.88	89.52 ± 3.65
	●	○	○	○	○	○	○	●	○	90.80 ± 4.80	90.42 ± 4.85	90.05 ± 5.08	90.80 ± 4.80
	●	○	○	○	○	○	○	●	○	93.95 ± 1.78	94.05 ± 1.83	94.15 ± 2.02	93.95 ± 1.78
	●	○	○	○	○	○	○	●	○	96.15 ± 1.80	96.25 ± 1.85	96.35 ± 2.04	96.15 ± 1.80
	●	○	○	○	○	○	○	●	○	98.25 ± 1.58	97.53 ± 1.62	97.12 ± 1.45	97.95 ± 1.58
	●	○	○	○	○	○	○	●	○	98.75 ± 1.54	98.71 ± 1.64	98.78 ± 1.47	98.65 ± 1.60
Multiclass	○	●	○	○	○	●	○	●	○	55.62 ± 2.08	44.95 ± 3.35	37.75 ± 8.12	55.62 ± 2.08
	○	●	●	○	○	○	○	●	○	57.32 ± 5.20	53.73 ± 5.58	50.45 ± 4.48	57.32 ± 5.20
	○	●	●	○	○	○	○	●	○	59.85 ± 5.22	56.05 ± 5.60	52.68 ± 4.50	59.85 ± 5.22
	○	○	○	○	○	○	○	○	○	80.58 ± 3.02	81.35 ± 2.72	82.15 ± 1.88	80.58 ± 3.02
	○	○	○	○	○	○	○	○	○	82.85 ± 2.18	81.77 ± 3.25	80.72 ± 4.45	82.85 ± 2.18
	○	○	○	○	○	○	○	○	○	85.62 ± 2.15	85.09 ± 0.80	84.58 ± 4.42	85.62 ± 2.15
	○	○	○	○	○	○	○	○	○	84.82 ± 1.62	81.77 ± 1.72	78.92 ± 1.88	84.82 ± 1.62
	○	○	○	○	○	○	○	○	○	86.95 ± 1.64	83.91 ± 1.74	81.05 ± 1.90	86.95 ± 1.64
	○	○	○	○	○	○	○	○	○	92.05 ± 3.60	92.15 ± 3.55	92.25 ± 3.48	92.05 ± 3.60
	○	○	○	○	○	○	○	○	○	93.15 ± 3.62	93.25 ± 3.57	93.35 ± 3.50	93.15 ± 3.62

connection types reveals the crucial role of each connectivity component in ClinGRAD’s architecture. The addition of basic structural connections improves accuracy by 10.92% in radiomics-only models, highlighting the importance of modeling spatial relationships in neuroimaging analysis. The incorporation of DWI-guided connections provides further enhancement, improving accuracy by 2.32% in radiomics-only models and by 2.20% in the full model. This improvement is particularly noteworthy given the high baseline performance, indicating that DWI-guided connections help capture subtle but clinically relevant structural patterns that might otherwise be missed. The integration of genetic co-expression relationships provides an additional performance boost, particularly evident in multi-class scenarios where accuracy improves from 86.95% to 93.15%. This substantial gain demonstrates the importance of modeling gene-gene interactions for distinguishing between different neurodegenerative conditions.

Interpretability. ClinGRAD highlights novel interactions between pathways, particularly between mitochondrial dysfunction and cell signaling networks. ClinGRAD’s gene clustering revealed seven distinct functional modules: neuroinflammation, APP/Tau pathology, vascular integrity, synaptic function, proteostasis, mitochondrial function, and cell signaling clusters (in Fig. 4).

Technical Contributions. ClinGRAD’s superior performance is driven by several technical innovations. First, the inclusion of DWI-guided connectivity patterns provides biologically informed constraints during feature learning, where data fusion is driven by clinically validated structural connectivity values rather than learned arbitrarily. Second, our multi-scale hierarchical framework enables simultaneous modeling of molecular mechanisms (gene co-expression), microscale structural alterations (radiomic features), and macroscale brain connectivity patterns (DWI edges). Third, ClinGRAD’s attention-based message-passing mechanism effectively models complex cross-modal interactions.

5 Conclusion

ClinGRAD represents a significant advancement in multimodal integration for AD classification. By leveraging heterogeneous GNN to combine genomic and radiomic data, ClinGRAD achieves high accuracy in differentiating dementia subtypes, which demonstrates the power of GNN in capturing complex biological relationships, outperforming individual modalities, and showcasing the synergistic effect of multimodal data integration. ClinGRAD’s interpretability is one of its standout features, offering clear insights into gene-gene interactions and the brain regions most impacted by AD, enhancing both its clinical applicability and our understanding of AD pathogenesis. While promising, future work can focus on validating ClinGRAD on larger, more diverse datasets and integrating additional omics data types to further enhance the model’s robustness.

Acknowledgments. This publication was supported by the Mohamed Bin Zayed University of Artificial Intelligence (MBZUAI) Graduate Student and Post-doctoral Researcher Conference Travel Grant Initiative.

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

References

1. Billot, B., Greve, D.N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A.V., Iglesias, J.E., et al.: Synthseg: Segmentation of brain mri scans of any contrast and resolution without retraining. *Medical image analysis* **86**, 102789 (2023)
2. Birkenbihl, C., Westwood, S., Shi, L., Nevado-Holgado, A., Westman, E., Lovestone, S., Hofmann-Apitius, M., Consortium, A., et al.: ANMerge: a comprehensive and accessible Alzheimer’s disease patient-level dataset. *Journal of Alzheimer’s Disease* **79**(1), 423–431 (2021)
3. Golovanevsky, M., Eickhoff, C., Singh, R.: Multimodal attention-based deep learning for Alzheimer’s disease diagnosis. *Journal of the American Medical Informatics Association* **29**(12), 2014–2022 (2022)
4. Gorishniy, Y., Rubachev, I., Khrulkov, V., Babenko, A.: Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems* **34**, 18932–18943 (2021)
5. Hassan, S., Akaila, D., Arjemandi, M., Papineni, V., Yaqub, M.: Mindsets: Multi-omics integration with neuroimaging for dementia subtyping and effective temporal study. *Scientific Reports* **15**(1), 1–12 (2025)
6. Human Protein Atlas: The Human Protein Atlas. <https://www.proteinatlas.org/> (2024), accessed: 2024-09-13
7. Jutten, R.J., Sikkes, S.A., Amariglio, R.E., Buckley, R.F., Properzi, M.J., Marshall, G.A., Rentz, D.M., Johnson, K.A., Teunissen, C.E., Van Berckel, B.N., et al.: Identifying sensitive measures of cognitive decline at different clinical stages of Alzheimer’s disease. *Journal of the International Neuropsychological Society* **27**(5), 426–438 (2021)

8. Ktena, S.I., Parisot, S., Ferrante, E., Rajchl, M., Lee, M., Glocker, B., Rueckert, D.: Distance metric learning using graph convolutional networks: Application to functional brain networks. In: Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part I 20. pp. 469–477. Springer (2017)
9. Lella, E., Lombardi, A., Amoroso, N., Diacono, D., Maggipinto, T., Monaco, A., Bellotti, R., Tangaro, S.: Machine learning and DWI brain communicability networks for Alzheimer’s disease detection. *Applied Sciences* **10**(3), 934 (2020)
10. Maddalena, L., Granata, I., Giordano, M., Manzo, M., Guarracino, M.R.: Integrating different data modalities for the classification of Alzheimer’s disease stages. *SN Computer Science* **4**(3), 249 (2023)
11. Mahoney, E.R., Dumitrescu, L., Moore, A.M., Cambronero, F.E., De Jager, P.L., Koran, M.E.I., Petyuk, V.A., Robinson, R.A., Goyal, S., Schneider, J.A., et al.: Brain expression of the vascular endothelial growth factor gene family in cognitive aging and Alzheimer’s disease. *Molecular psychiatry* **26**(3), 888–896 (2021)
12. Nandi, A., Counts, N., Chen, S., Seligman, B., Tortorice, D., Vigo, D., Bloom, D.E.: Global and regional projections of the economic burden of Alzheimer’s disease and related dementias from 2019 to 2050: A value of statistical life approach. *EClinicalMedicine* **51** (2022)
13. 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)
14. Pedroza, P., Miller-Petrie, M.K., Chen, C., Chakrabarti, S., Chapin, A., Hay, S., Tsakalos, G., Wimo, A., Dieleman, J.L.: Global and regional spending on dementia care from 2000–2019 and expected future health spending scenarios from 2020–2050: an economic modelling exercise. *EClinicalMedicine* **45** (2022)
15. Prado, P., Medel, V., Gonzalez-Gomez, R., Sainz-Ballesteros, A., Vidal, V., Santamaría-García, H., Moguilner, S., Mejia, J., Slachevsky, A., Behrens, M.I., et al.: The BrainLat project, a multimodal neuroimaging dataset of neurodegeneration from underrepresented backgrounds. *Scientific Data* **10**(1), 889 (2023)
16. Sadil, P., Lindquist, M.A.: Comparing automated subcortical volume estimation methods; amygdala volumes estimated by FSL and FreeSurfer have poor consistency. Tech. rep., Wiley Online Library (2024)
17. Said, A., Bayrak, R., Derr, T., Shabbir, M., Moyer, D., Chang, C., Koutsoukos, X.: Neurograph: Benchmarks for graph machine learning in brain connectomics. *Advances in Neural Information Processing Systems* **36**, 6509–6531 (2023)
18. Shi, Y., Zu, C., Hong, M., Zhou, L., Wang, L., Wu, X., Zhou, J., Zhang, D., Wang, Y.: ASMFs: Adaptive-similarity-based multi-modality feature selection for classification of Alzheimer’s disease. *Pattern Recognition* **126**, 108566 (2022)
19. Sim, J., Lee, M., Wu, G., Kim, W.H.: Multi-modal Graph Neural Network with Transformer-Guided Adaptive Diffusion for Preclinical Alzheimer Classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 511–521. Springer (2024)
20. Škoch, A., Rehák Bučková, B., Mareš, J., Tintěra, J., Sanda, P., Jajcay, L., Horáček, J., Španiel, F., Hlinka, J.: Human brain structural connectivity matrices—ready for modelling. *Scientific Data* **9**(1), 486 (2022)
21. Song, F., Li, Y., Jiang, M., Li, K., Zhang, J., Zhang, Y., Pang, Z.: HSGO: Harmonized swarm learning with guided optimization for multi-center sMRI classification of Alzheimer’s disease. *IEEE Journal of Biomedical and Health Informatics* (2025)

22. Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M.: Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* **15**(1), 273–289 (2002)
23. Van Griethuysen, J.J., Fedorov, A., Parmar, C., Hosny, A., Aucoin, N., Narayan, V., Beets-Tan, R.G., Fillion-Robin, J.C., Pieper, S., Aerts, H.J.: Computational radiomics system to decode the radiographic phenotype. *Cancer research* **77**(21), e104–e107 (2017)
24. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017)
25. Venugopalan, J., Tong, L., Hassanzadeh, H.R., Wang, M.D.: Multimodal deep learning models for early detection of Alzheimer’s disease stage. *Scientific reports* **11**(1), 3254 (2021)
26. Wang, J., Xiao, T., Zhang, X., Pan, C., Xia, Z., Chen, S.: Multicenter and multi-channel pooling GCN for early AD diagnosis based on dual-modality fused brain network. *IEEE Transactions on Medical Imaging* **41**(6), 1560–1572 (2022)
27. Wang, X., Huang, W., Su, L., Xing, Y., Jessen, F., Sun, Y., Shu, N., Han, Y.: Neuroimaging advances regarding subjective cognitive decline in preclinical Alzheimer’s disease. *Molecular Neurodegeneration* **15**, 1–27 (2020)
28. Warde-Farley, D., Donaldson, I., Comes, O., Zuberi, K., Badrawi, R., Chao, P., Franz, M., Grouios, C., Kazi, F., Lopes, C.T., Maitland, A., Mostafavi, S., Montojo, J., Shao, Q., Wright, G.B., Bader, G.D., Morris, Q.: The GeneMANIA prediction server: biological network integration for gene prioritization and predicting gene function. *Nucleic Acids Research* **38**(suppl_2), W214–W220 (2010). <https://doi.org/10.1093/nar/gkq537>
29. Yun, S., Choi, I., Peng, J., Wu, Y., Bao, J., Zhang, Q., Xin, J., Long, Q., Chen, T.: Flex-moe: Modeling arbitrary modality combination via the flexible mixture-of-experts. *arXiv preprint arXiv:2410.08245* (2024)