Structural Connectome Analysis using a Graph-based Deep Model for Age and Dementia Prediction

Anees Kazi[1,2], **Jocelyn Mora**[1], **Bruce Fischl**[1,2], **Adrian V. Dalca**[1,2,3], **Iman Aganj**[1,2]

[1] Athinoula A. Martinos Center for Biomedical Imaging, Radiology Department, Massachusetts General Hospital, Boston, USA

[2] Radiology Department, Harvard Medical School, Boston, USA

[3] CSAIL, Massachusetts Institute of Technology, Cambridge, USA

Keywords: Graph neural network (GNN), Graph convolutional network (GCN), Prediction, Deep learning, Age, Dementia

ABSTRACT

We tackle the prediction of age and mini-mental state examination (MMSE) score based on structural brain connectivity derived from diffusion magnetic resonance images. We propose a machine-learning model inspired by graph convolutional networks (GCNs), which takes a brain connectivity input graph and processes the data separately through a parallel GCN mechanism with multiple branches, thereby disentangling the input node and graph features. The novelty of our work lies in the model architecture, especially the connectivity attention module, which learns an embedding representation of brain graphs while providing graph-level attention. We show experiments on publicly available datasets of PREVENT-AD and OASIS3. Through our experiments, we validate our model by comparing it to existing methods and via ablations. This quantifies the degree to which the connectome varies depending on the task, which is important for improving our understanding of health and disease across the population. The proposed model generally demonstrates higher performance especially for age prediction compared to the existing machine-learning algorithms we tested, including classical methods and (graph and non-graph) deep learning. We provide a detailed analysis of each component.

Corresponding author: Anees Kazi, AKAZI1@mgh.harvard.edu

INTRODUCTION

Structural connections between the brain regions constitute complex brain networks, known as the connectome (Hofman, 2015). Structural connectome can be quantified via diffusion-weighted MRI (dMRI), thereby revealing the physical connections between the brain regions, which are essential for understanding the brain's overall organization (Sporns, 2011). Analyzing the connectome provides insight into neural circuitry, connectivity patterns, and their implications for cognition and behavior, contributing to our understanding of brain development, aging, and the impact of neurological conditions on brain wiring (Hagmann et al., 2008; F. Zhang et al., 2022).

Structural connectome graphs have been used to study a wide range of neurological and psychiatric disorders, including Alzheimer's disease (AD) (Aganj, Mora, Frau-Pascual, Fischl, & Initiative, 2023; Amoroso et al., 2017; Frau-Pascual et al., 2021; He, Chen, & Evans, 2008; J. Wang et al., 2018), schizophrenia (Karlsgodt, Sun, & Cannon, 2010; Shi et al., 2012; Y.-m. Wang et al., 2020), autism spectrum disorders (ASD) (Tolan & Isik, 2018), and Parkinson's disease (Arrigo et al., 2019; Jellinger, 2022; X. Zhang et al., 2018; Y. Zhang, Zhan, Cai, Thompson, & Huang, 2019), as well as to understand cognitive decline (Xu et al., 2021) and normal brain development and aging (Coelho et al., 2021; Damoiseaux, 2017; Dennis et al., 2013; Lewis, O'Reilly, Bock, Theilmann, & Townsend, 2022; Neudorf, Shen, & McIntosh, 2024). Recently, (Prescott et al., 2022) also showed that structural connectivity is a crucial factor in identifying early-onset AD risk, as individuals with a genetic predisposition show lower connectivity, especially in the frontoparietal control network, and this reduced connectivity is linked to the estimated time until dementia symptoms emerge. Studying brain aging through the structural connectome is crucial as it provides insight into how brain connectivity and network organization change over time, which can be related to cognitive decline, neurodegenerative diseases, and overall brain health (Meunier, Achard, Morcom, & Bullmore, 2009).

The connectome is a relatively new avenue for studying the brain, which is not yet included in clinical practice but holds promise for discovery. Given the complexity and the vast amount of data involved, deep learning (DL) techniques are likely to be effective in analyzing the connectome. Structural brain connectivity analysis is a data-driven approach that explores the overall structural connections in the brain. It involves studying patterns of anatomical connectivity via fiber tracking (tractography) throughout the brain (Zalesky, Cocchi, Fornito, Murray, & Bullmore, 2012). Structural brain connectivity

analysis using artificial intelligence (AI) is an emerging field that remains notably new in the research landscape (Dubost, 2020; Sjöblom, Westerlund, Neimantaite, & Andersson, 2020). The exploration of AI applications in deciphering dMRI data to study brain connectivity is an evolving domain that needs further investigation (Faiyaz, Doyley, Schifitto, & Uddin, 2023). In this paper, we attempt to analyze dMRI-derived connectivity using DL for age and dementia prediction.

The task of predicting clinical and demographic data using the structural connectome involves analyzing high-dimensional and complex data and can be challenging for conventional DL methods. Graph neural networks (GNNs), particularly the graph convolutional network (GCN) architecture (S. Zhang, Tong, Xu, & Maciejewski, 2019), provide a powerful and flexible framework for analyzing brain connectivity data. GCNs can learn from the complex interrelationships between nodes and edges, capturing both local and global patterns in the graph structure. This makes GCNs well suited for the prediction task and for identifying the most predictive regions and connections of the brain. Furthermore, GCNs can be trained on large datasets, thus increasing their generalizability and applicability to different populations and contexts.

GCNs can leverage the rich structural information in the connectome to make accurate predictions. They do so by performing iterative message-passing between neighboring nodes in the graph, using learnable functions to aggregate and transform information from neighboring nodes, and updating the features of each node based on the aggregated information. This allows GCNs to capture the complex relationships between brain regions and their connections, and make predictions based on this information. A GCN-based model is promising for analyzing structural connectome and has shown great potential for improving our understanding of neurological and psychiatric disorders, as well as normal brain development and aging.

Several state-of-the-art methods have shown the application of GCNs in brain networks analysis. Most recent ones involve functional brain network analysis (X. Li et al., 2021), psychiatric disorder diagnosis (Zheng, Yu, Chen, Dang, & Chen, 2024), optimization of GNN architectures for schizophrenia spectrum disorder prediction (S. Wang et al., 2024), classification of ASD versus hyper complex brain networks (Hu, Cao, Li, Dong, & Li, 2021), and sex classification (Ktena et al., 2018). Many papers have shown work using structural networks for causal inference (Wein et al., 2021), brain age estimation (Moon et al., 2024), early diagnosis of AD (Y. Zhang, He, Chan, Teng, & Rajapakse, 2023), and schizophrenia

diagnosis (Sebenius, Campbell, Morgan, Bullmore, & Liò, 2021), achieving high performance compared to existing machine-learning (ML) and DL methods. GCNs have also been combined with recurrent neural networks to predict sex on temporal fMRI brain graphs (Kazi et al., 2022). Spectral GCN has been employed for region-of-interest (ROI) identification in functional connectivity graphs and for sex classification (Arslan, Ktena, Glocker, & Rueckert, 2018).

Methodologically, most GCN methods use the Laplacian-filter based spectral implementation (Kipf & Welling, 2016) and graph attention networks (GATs) (Veličković et al., 2017) incorporate attention mechanisms for edge weighting. Dynamic Graph CNNs (DGCNNs) (Phan, Le Nguyen, Nguyen, & Bui, 2018) are used for temporal graph analysis of multi-level graph structures, and hierarchical learning and spatiotemporal GNNs are used for integration of spatial and temporal data. GCNs cannot generalize to new graphs due to dependency on the fixed graph Laplacian (Q. Li, Han, & Wu, 2018; **?**). Further, GATs are inefficient in capturing long-range dependencies due to over-smoothing (Dasoulas, Scaman, & Virmaux, 2021). Both GCN and GAT have high computational costs for large (Liu, Park, & Yoo, 2020) and dense (Wu et al., 2020) graphs, respectively. DGCNNs are computationally inefficient and unstable due to repeated graph construction at each layer, leading to high memory usage and potential noise amplification.

In this paper, we propose a computationally efficient novel network architecture that uses 1) residual connections inspired by ResGCN (Pei, Huang, Van Ipenburg, & Pechenizkiy, 2022), 2) linear layers, and 3) a newly proposed connectivity attention module (CAM). The proposed model captures feature embeddings at both the subject and node levels, effectively spanning from granular details to high-level information. We show through experiments on two public databases that the proposed model outperforms non-graph and graph-based DL methods and, in many cases, conventional ML methods. The rest of the paper includes a description of the proposed method, experiments, discussion, and conclusion.

METHOD

Let us say a dataset has S number of subjects. Any subject comes with a brain graph $\mathbf{G} \in \mathbb{R}^{N \times N}$, with N the number of brain regions (graph nodes) and a feature matrix $\mathbf{X} \in \mathbb{R}^{N \times M}$, with M the number of features per node, and a label $y \in \mathbb{R}$ (i.e., value of the clinical or demographic variable to be predicted).

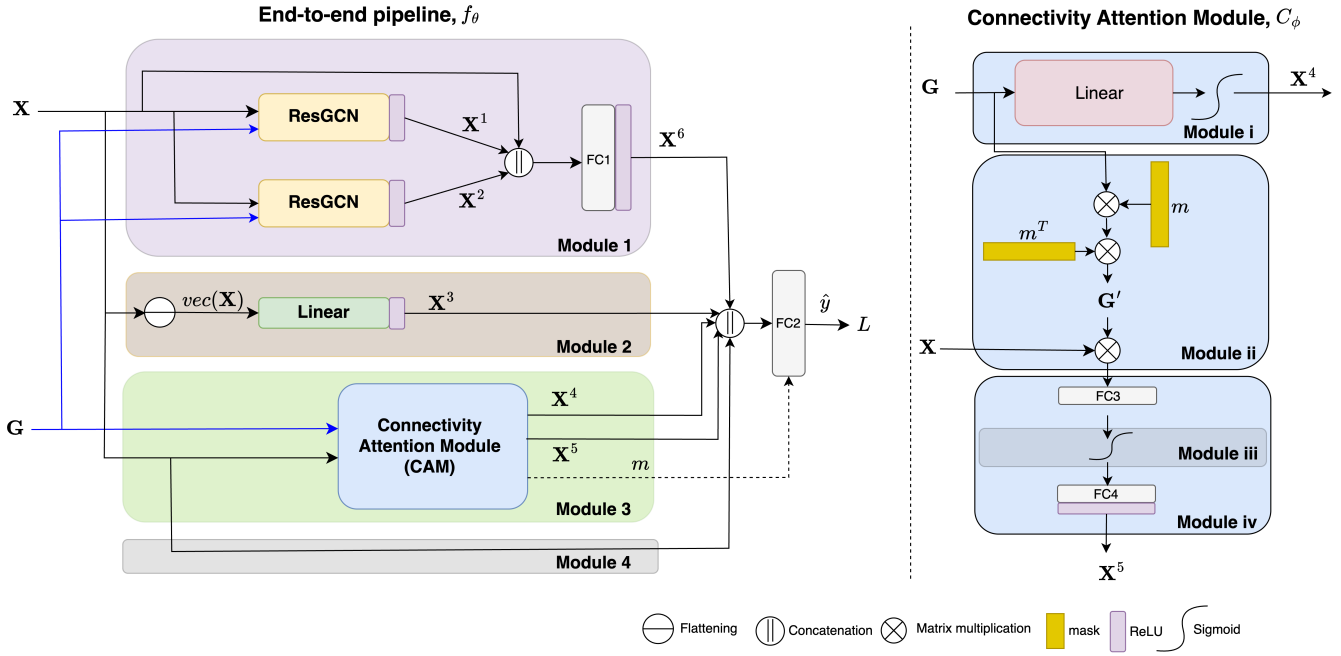

Figure 1. End-to-end pipeline of the proposed model (left) and Connectivity Attention Module (right).

Given $\mathbf{G}$ and $\mathbf{X}$, the task is to predict y, for which we define a model f_θ as:

$$y = f_\theta(\mathbf{G}, \mathbf{X}), \tag{1}$$

where θ is the set of learnable parameters. The model we propose comprises four sub-parts, modules 1 to 4, as shown in Figure 1. Each module is designed to process the combination of $\mathbf{X}$ and $\mathbf{G}$ to produce latent embedding. The last part combines all former outputs to produce the predicted label $\hat{y}$. We explain all four modules separately and explain the whole end-to-end model.

Module 1: Graph Convolution

This module comprises two branches, each with a graph convolution with a different embedding size, and a skip connection as shown in Figure 1. We employ a combination of ResGCN layers (Bresson & Laurent, 2017), which captures low-level features of the graph. Each ResGCN (with a different embedding size) transforms the same input data into a different space, enabling the model to learn diverse representations of the graph data. One layer focuses on capturing low-dimensional, essential features, while the other learns higher-dimensional, more nuanced relationships. The skip connection is then

combined with these learned representations, allowing the model to leverage both the basic, essential features and the more intricate relationships within the graph, potentially leading to a richer and more informative representation. The outputs of these layers are fed to a fully connected (FC) layer to produce the output of the module, X^6.

Module 2: Linear Layer

In this layer, we remove any structure that is present in the feature vector sequence $\mathbf{X}$ and try to simply fetch information from the raw data. The linear layer can be mathematically defined as: $\mathbf{X}^3 = \rho\left(MLP(vec(\mathbf{X}))\right)$, where $vec(\mathbf{X}) \in \mathbb{R}^{NM\times 1}$ is the flattened matrix $\mathbf{X}$, MLP is multi-layer perceptron, and ρ is the non-linearity (ReLU). A linear layer processing node features in parallel with GCN layers offers several potential benefits. It allows the model to learn an independent representation of the node features, capturing information that might not be directly reflected in the graph structure. This is particularly advantageous when the node features themselves hold significant information for the task at hand. In our case, we have dMRI-derived measures such as structural connectivity of the node ROI to the rest of the ROIs and mean fractional anisotropy (FA) and mean apparent diffusion coefficient (ADC) in the node ROI, as well as volume of the node ROI and the degree of each node. In predicting age and AD, diffusion MRI analysis can benefit from a combination of node features. FA and ADC, reflecting white matter organization and water diffusion, respectively, capture microstructural changes due to age and AD, whereas mean segmentation volume captures brain region size that is potentially reduced in age-related atrophy or AD. Finally, node degree (total connectivity to a brain region) reveals alterations in brain network connectivity patterns.

Module 3: Connectivity Attention Module

Attention mechanisms in DL enhance model performance by dynamically focusing on the most relevant features in input data, improving interpretability and efficiency in tasks like natural language processing and computer vision (Bahdanau, 2014; Vaswani, 2017). Attention mechanisms have significantly advanced medical image analysis, particularly in brain imaging applications, by enabling models to focus on critical regions, thereby enhancing diagnostic accuracy and interpretability (Jiao et al., 2023; Ranjbarzadeh et al., 2021; Y. Zhang, Teng, et al., 2023).

We propose a customized connectivity attention module (CAM) C_ϕ so the model learns an embedding representation of a brain graph, as well as to provide a graph-level attention mechanism. The CAM, depicted in Figure 1 (right), is defined as

$$\mathbf{X^4}, \mathbf{X^5}, \hat{\mathbf{m}} = C_\phi(\mathbf{X}, \mathbf{G}, \mathbf{m}), \tag{2}$$

where X^4 is a scalar embedding of each subject's brain connectivity, which is also treated as a factor of importance of each subject with respect to the population. X^5 is the lower dimensional representation that is sculpted out of the corresponding G and X from the model. $\hat{\mathbf{m}}$, the external trainable parameter, is the attention mask assigned to the nodes, i.e. ROIs in the brain. The CAM $C_\phi = \{C_p, C_e\}$. All are defined as:

$$C_\phi = \begin{cases} C_p: & \mathbf{X}^4 = \sigma\left[\langle \mathbf{G}, \mathbf{E} \rangle_F\right] \\ C_e: & \mathbf{X}^5 = \rho[f_4(\sigma[f_3(\mathbf{G}\hat{\mathbf{m}}\hat{\mathbf{m}}^T\mathbf{X})])] \\ \hat{\mathbf{m}}: & \hat{\mathbf{m}} \in \mathbb{R}^{1\times N} \end{cases} \tag{3}$$

$\mathbf{X}^4$ is the Frobenius inner product of the $\mathbf{G}$ matrix by a weight matrix $\mathbf{E}$ learned by the model, followed by sigmoid non-linearity, σ. For $\mathbf{X}^5$, we effectively project $\mathbf{G}$ onto a single one-dimensional orientation in the N-dimensional space, from which we create a rank-one matrix $\mathbf{G}' := \mathbf{G}\hat{\mathbf{m}}\hat{\mathbf{m}}^T$, with the projection weights $\hat{\mathbf{m}}$ learned by the model. f_3 and f_4 are the fully connected layers FC_3 and FC_4 as shown in Figure 1, both followed by non-linearities.

Here, instead of applying attention to each element of $\mathbf{G}$, we leverage the matrix properties of $\mathbf{G}$ by focusing on eigenvalues. Our goal is to retain a single representative rank-one matrix $\mathbf{G}'$ that best helps the prediction task at hand. This operation simplifies the representation of the matrix while preserving its most significant features.

Module 4: Skip Connection

We add an overall skip connection (Module 4 in Figure 1) to help mitigate the vanishing gradient problem and allow effective information propagation across layers by directly connecting $vec(\mathbf{X})$ to the final layer. This helps to preserve modality-specific features, preventing over smoothing.

Data Fusion

Table 1. Description of dataset size (number of available scans), distribution across the classes, and partitioning. Due to missing demographic data, nine subjects were removed from the OASIS3 dataset. The female ratio is the portion of scans from female subjects.

Name	Subjects	Total samples	Samples-10%	10%	Female ratio
PREVENT–AD	347	789	710	79	72%
OASIS3	771	1294	1164	121	56%

The outputs of all four modules are concatenated and fed to a FC layer to produce the final prediction $\hat{y}$. Mathematically, $\hat{y} = MLP([X^3, X^4, X^5, X^6, vec(\mathbf{X})])$.

Loss Function and Optimization

For the regression task, we use the Huber loss, $L_\delta(y - \hat{y})$, which is a piecewise function that is quadratic for small errors and linear for large errors, defined as:

$$L_\delta(e) = \begin{cases} \frac{1}{2}e^2 & \text{if } |e| \le \delta, \\ \delta \cdot (|e| - \frac{1}{2}\delta) & \text{otherwise,} \end{cases}$$

where δ is a threshold that distinguishes between quadratic and linear loss behavior. This function applies a quadratic loss when the absolute error is less than or equal to δ and a linear loss when the error exceeds this threshold, effectively combining the differentiability of the quadratic loss at small errors with the robustness of the linear loss to outliers. We also regularize $\mathbf{m}$ by adding a second term $L^R = \sum_i^N |m_i| - \sum_i^N |m_i| \log |m_i|$. The overall loss for the model is therefore $L = L_\delta + \alpha L^R$, where $\alpha > 0$ is a scaling factor.

Datasets

Pre-symptomatic Evaluation of Experimental or Novel Treatments for Alzheimer's Disease (PREVENT-AD) (Leoutsakos, Gross, Jones, Albert, & Breitner, 2016) is a publicly available dataset that aims to provide a comprehensive set of data on individuals who are at risk for developing AD (https://prevent-alzheimer.net). The database contains neuroimaging studies such as MRI (including dMRI) and PET scans, a range of demographic, clinical, cognitive, and genetic data, as well as data on lifestyle factors such as diet and exercise. The dataset comprises 347 subjects, some with multiple (longitudinal) dMRI scans, totaling 789 dMRI scans.

Open Access Series of Imaging Studies, whose third release (OASIS3) we used here (LaMontagne et al., 2019), is a longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and AD, provided freely to researchers (http://www.oasis-brains.org). The OASIS3 dataset contains MRI scans (including dMRI), cognitive assessments, demographic information, and clinical diagnoses for subjects, including healthy controls, individuals with MCI, and AD patients. We used 1294 brain scans from 771 subjects.

Pre-processing

We used FreeSurfer (Fischl, 2012) to process the databases (additionally applying the longitudinal processing pipeline (Reuter, Schmansky, Rosas, & Fischl, 2012) for PREVENT-AD). We then ran the FreeSurfer diffusion processing pipeline and propagated the 85 automatically segmented cortical and subcortical regions from the structural to the diffusion space. These 85 regions act as the nodes in our graph setup. Next, we used our public toolbox (http://www.nitrc.org/projects/csaodf-hough) to reconstruct the diffusion orientation distribution function in constant solid angle (Aganj et al., 2010), run Hough-transform global probabilistic tractography (Aganj et al., 2011) to generate 10,000 fibers per subject, compute symmetric structural connectivity matrices, and augment the matrices with indirect connections (Aganj et al., 2014). More details on the pipeline can be found in our previous publication (Aganj et al., 2023). Once we had all the graphs G_i, we performed a population-level normalization on edge weights. For node features, we used the volume, ADC, FA, and the degree of each node in the affinity matrix obtained for each ROI, as well as the row in G_i representing connectivity to the rest of the brain. Therefore, for each subject we obtained $G_i \in \mathbb{R}^{85\times 85}$ and corresponding $X_i \in \mathbb{R}^{85\times 89}$ (i.e., $N = 85$ and $M = 89$).

Implementation details. All the experiments were run via 10-fold cross validation with the same folds across methods and experiments. For model robustness, we added zero-mean Gaussian noise with a standard deviation of 0.0001 to the training samples. All the experiments were run on a Linux machine with 512 GB of RAM, an Intel (R) Xeon (R) Gold 6256 CPU @ 3.60 GHz, and an NVIDIA RTX A6000 (48 GB) graphics processing unit. The total number of parameters used in the proposed model was 5073, which was comparable to GCNConv (2453), DGCNN (4653), Graphconv (4653), ResGatedGraphConv (RGGC) (9128), and GINConv (2683). In our experiments, the values of d_1, d_2, d_3, and d_4 were 25, 20, 5,

and 2, respectively. We kept 10% of the data aside from each dataset so as not to heuristically fit the model to the entire data, and then tested the model at the end on the unseen data, the results of which are reported in the next section.

RESULTS AND DISCUSSION

Baselines and Comparative Methods

Table 2 presents a comparative evaluation of various conventional ML and advanced DL models for the age prediction task across the two datasets of PREVENT-AD (789 subjects) and OASIS3 (1294 subjects), as well as for MMSE prediction in the OASIS3 dataset. Performance was assessed using metrics such as the root mean square error (RMSE) and mean absolute error (MAE) of the prediction, as well as Pearson's correlation (PC) and Spearman's correlation (SC) coefficients between the predicted and ground-truth values. For each of the three tasks, we also computed baseline values for our performance metrics as RMSE and MAE of a naive model that always predicts, respectively, the mean and median of the training-set values (PC and SC are zero for such a fixed-value predictor). The conventional ML models that we tested include linear regression, support vector regression (SVR), decision tree, regression tree, ensemble tree, and neural networks (all implemented in using sklearn package (Komer, Bergstra, & Eliasmith, 2019)). We further compared our method with DL approaches such as MLP, and especially more advanced GNN models such as GCN, GIN, GraphConv, and ResGCN. For both datasets, the proposed method outperformed all other models in age prediction in 7 out of 8 metrics, achieving the lowest prediction error (except it finished second in MAE for PREVENT-AD) while also demonstrating the highest correlation scores, indicating superior predictive accuracy. For MMSE prediction (in OASIS3), the proposed method still outperformed the other DL methods, but not the conventional ML methods of ensemble tree and (mostly) SVR. Note that the narrow range of MMSE in the OASIS3 dataset (as reflected in the Baseline row) resulted in performance metric values that were very close to each other among models. Classical ML approaches of linear regression (not shown in the table), decision tree, and regression tree generally exhibited poorer performance, indicating limited effectiveness for the prediction task.

We additionally performed these experiments on data from the second phase of the Alzheimer's Disease Neuroimaging Initiative, with a dataset size of 200 samples. While the proposed method still

outperformed the rest of the methods, its prediction error was not substantially below the baseline, presumably due to the small dataset size (PC/SC were still 0.50/0.55 and 0.34/0.36 for age and MMSE, respectively).

Ablation Tests

Table 3 presents an ablation study on the prediction tasks on both datasets, where different components were removed from the model architecture to assess their individual contributions. The same evaluation metrics as in the previous subsection were used.

For PREVENT-AD, when removing specific components such as the linear block, CAM block, skip connection, or GCN block, performance degrades across all datasets, confirming the effectiveness of the model's complete architecture. Omitting the linear component leads to an increase in RMSE and a drop in correlation scores. Similarly, removing the GCN block results in performance loss. The trend remains generally consistent in OASIS3 (age); however, the original model without the linear component maintains the best scores, whereas other ablations cause notable declines in performance, indicating weaker predictive capability.

For OASIS3 (MMSE), the original model without the skip connection outperforms the other ablated versions in terms of correlations, though the performance drop is less pronounced. Removing any component results in only a marginal increase in RMSE and a decrease in correlation metrics, possibly due to the narrow range of MMSE values in OASIS3.

The ablation results generally emphasize that each architectural component plays a crucial role in maintaining predictive accuracy, with graph-based learning contributing significantly to overall model performance. The GCN block, in particular, appears essential for capturing complex dependencies, as its removal leads to the most significant decline across all datasets.

CONCLUSION

In this paper, we proposed a simple yet effective model capable of capturing complementary information from structural brain connectivity graphs, which we evaluated in the context of age and MMSE prediction. The configuration of input data, the initialization of neighborhood information as node

Table 2. Comparative methods for age and MMSE prediction tasks

	PREVENT-AD (age) (789 subjects)				OASIS3 (age) (1294 subjects)				OASIS3 (MMSE) (1294 subjects)			
Method	RMSE	MAE	PC	SC	RMSE	MAE	PC	SC	RMSE	MAE	PC	SC
Baseline	6.17	5.65	0	0	9.46	7.21	0	0	3.02	1.71	0	0
SVR	6.25	5.19	0.41	0.42	7.56	5.55	0.67	0.67	3.00	**1.63**	0.35	**0.35**
Decision Tree	7.50	6.15	0.20	0.17	10.86	8.37	0.29	0.28	3.40	1.92	0.20	0.19
Regression Tree	7.11	5.94	0.22	0.27	10.21	7.90	0.32	0.29	3.33	1.86	0.19	0.20
Ensemble Tree	5.54	4.69	0.37	0.39	7.35	5.57	0.64	0.62	**2.76**	1.64	**0.41**	0.31
Neural Network	5.39	4.60	0.40	0.43	7.47	5.91	0.67	0.67	3.48	2.56	0.25	0.22
MLP	5.80	4.81	0.43	0.52	10.12	7.14	0.26	0.35	3.31	2.02	0.17	0.15
GCN	5.34	4.46	0.56	0.57	6.72	5.41	0.66	0.68	3.08	1.79	0.16	0.23
GIN	5.37	4.46	0.47	0.48	9.51	6.36	0.45	0.60	2.96	2.11	0.27	0.24
GraphConv	5.49	4.45	0.48	0.51	8.92	7.02	0.25	0.24	3.48	2.02	0.15	0.25
ResGCN	5.28	**4.39**	0.52	0.55	8.50	5.50	0.54	0.68	3.07	1.94	0.22	0.23
Proposed	**5.27**	4.42	**0.62**	**0.66**	**6.04**	**4.70**	**0.75**	**0.77**	2.89	1.76	0.34	0.27

Bold and underlined indicate the **best** and second-best, respectively, in each column.

features, and multiple operations from all four modules helped to learn better representations of each subject's graph. We have shown that our model often outperforms competing techniques on two publicly available datasets, while also ablating all the components of the model. Future work includes the addition of interpretability to the models to find the brain subnetworks that are informative for the prediction task. Further steps would be trying different graph convolution mechanisms, such as gated attention graph convolutions.

ACKNOWLEDGMENTS

Table 3. Ablations on the prediction tasks. Each column represents the respective block removed from the pipeline.

	Model-level					CAM layer				
Metric	Proposed	Mod 1	Mod 2	Mod 3	Mod 4	Mod i	Mod ii		Mod iii	Mod iv
		GCN block	Linear	CAM	skip	Linear	$G\hat{\mathbf{m}}\hat{\mathbf{m}}^T$	G in $G\hat{\mathbf{m}}\hat{\mathbf{m}}^T$	sigmoid	MLPs
	PREVENT-AD (age)									
RMSE	**5.27**	5.76	5.30	5.65	5.62	6.47	6.38	6.47	6.02	6.44
MAE	**4.42**	4.68	4.50	4.7	4.68	5.14	5.09	5.16	4.87	5.14
PC	**0.62**	0.55	0.54	0.54	0.52	0.50	0.51	0.50	0.53	0.50
SC	**0.66**	0.57	0.58	0.57	0.56	0.51	0.51	0.51	0.58	0.51
	OASIS3 (age)									
RMSE	6.04	6.04	**5.78**	6.44	6.72	7.04	7.01	7.02	7.01	6.90
MAE	4.70	4.63	**4.47**	5.29	5.10	5.21	5.06	5.09	5.20	5.08
PC	0.75	0.73	**0.76**	0.71	0.68	0.64	0.64	0.64	0.65	0.65
SC	0.77	0.75	**0.77**	0.70	0.70	0.67	0.67	0.68	0.67	0.68
	OASIS3 (MMSE)									
RMSE	2.89	2.85	2.82	2.84	2.94	2.84	2.87	2.85	2.80	**2.79**
MAE	1.76	**0.83**	1.75	1.91	1.65	1.73	1.68	1.70	1.77	1.79
PC	0.34	0.35	0.35	0.36	**0.39**	0.34	0.35	0.34	0.38	0.38
SC	0.27	0.27	0.25	0.27	**0.33**	0.28	0.27	0.25	0.31	0.26

Bold and underlined indicate the **best** and second-best, respectively, in each row. Mod stands for module.

Support for this research was provided by the National Institutes of Health (NIH), specifically the National Institute on Aging (RF1AG068261, R01AG068261). Additional support was provided in part by the BRAIN Initiative Cell Census Network grant U01MH117023, the National Institute for Biomedical Imaging and Bioengineering (P41EB015896, R01EB023281, R01EB006758, R21EB018907, R01EB019956, P41EB030006), the National Institute on Aging (R56AG064027,

R01AG064027, R01AG008122, R01AG016495, R01AG070988), the National Institute of Mental Health (R01MH121885, RF1MH123195), the National Institute for Neurological Disorders and Stroke (R01NS0525851, R21NS072652, R01NS070963, R01NS083534, U01NS086625, U24NS10059103, R01NS105820), the NIH Blueprint for Neuroscience Research (U01MH093765), part of the multi-institutional Human Connectome Project, and the Michael J. Fox Foundation for Parkinson's Research (MJFF-021226). Computational resources were provided through the Massachusetts Life Sciences Center.

B. Fischl is an advisor to DeepHealth, a company whose medical pursuits focus on medical imaging and measurement technologies. His interests were reviewed and are managed by Massachusetts General Hospital and Mass General Brigham in accordance with their conflict-of-interest policies.

REFERENCES

Aganj, I., Lenglet, C., Jahanshad, N., Yacoub, E., Harel, N., Thompson, P. M., & Sapiro, G. (2011). A hough transform global probabilistic approach to multiple-subject diffusion mri tractography. *Medical image analysis*, *15*(4), 414–425.

Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., & Harel, N. (2010). Reconstruction of the orientation distribution function in single-and multiple-shell q-ball imaging within constant solid angle. *Magnetic resonance in medicine*, *64*(2).

Aganj, I., Mora, J., Frau-Pascual, A., Fischl, B., & Initiative, A. D. N. (2023). Exploratory correlation of the human structural connectome with non-mri variables in alzheimer's disease. *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring*, *15*(4), e12511.

Aganj, I., Prasad, G., Srinivasan, P., Yendiki, A., Thompson, P. M., & Fischl, B. (2014). Structural brain network augmentation via kirchhoff's laws. In *Joint annual meeting of ismrm-esmrmb, http://nmr. mgh. harvard. edu/~ iman/conductancemodel_ismrm14_iman. pdf* (Vol. 22, p. 2665).

Amoroso, N., La Rocca, M., Bruno, S., Maggipinto, T., Monaco, A., Bellotti, R., & Tangaro, S. (2017). Brain structural connectivity atrophy in alzheimer's disease. *arXiv preprint arXiv:1709.02369*.

Arrigo, A., Calamuneri, A., Milardi, D., Mormina, E., Gaeta, M., Corallo, F., ... others (2019). Claustral structural connectivity and cognitive impairment in drug naïve parkinson's disease. *Brain Imaging and Behavior*, *13*, 933–944.

Arslan, S., Ktena, S. I., Glocker, B., & Rueckert, D. (2018). Graph saliency maps through spectral convolutional networks: Application to sex classification with brain connectivity. In *Graphs in biomedical image analysis and integrating*

medical imaging and non-imaging modalities: Second international workshop, grail 2018 and first international workshop, beyond mic 2018, held in conjunction with miccai 2018, granada, spain, september 20, 2018, proceedings 2 (pp. 3–13).

Bahdanau, D. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Bresson, X., & Laurent, T. (2017). Residual gated graph convnets. *arXiv preprint arXiv:1711.07553*.

Coelho, A., Fernandes, H. M., Magalhães, R., Moreira, P. S., Marques, P., Soares, J. M., … others (2021). Reorganization of brain structural networks in aging: A longitudinal study. *Journal of Neuroscience Research*, *99*(5), 1354–1376.

Damoiseaux, J. S. (2017). Effects of aging on functional and structural brain connectivity. *Neuroimage*, *160*, 32–40.

Dasoulas, G., Scaman, K., & Virmaux, A. (2021). Lipschitz normalization for self-attention layers with application to graph neural networks. In *International conference on machine learning* (pp. 2456–2466).

Dennis, E. L., Jahanshad, N., McMahon, K. L., de Zubicaray, G. I., Martin, N. G., Hickie, I. B., … Thompson, P. M. (2013). Development of brain structural connectivity between ages 12 and 30: a 4-tesla diffusion imaging study in 439 adolescents and adults. *Neuroimage*, *64*, 671–684.

Dubost, F. (2020). Artificial intelligence with light supervision: Application to neuroimaging.

Faiyaz, A., Doyley, M. M., Schifitto, G., & Uddin, M. N. (2023). Artificial intelligence for diffusion mri-based tissue microstructure estimation in the human brain: an overview. *Frontiers in Neurology*, *14*, 1168833.

Fischl, B. (2012). Freesurfer. *Neuroimage*, *62*(2), 774–781.

Frau-Pascual, A., Augustinack, J., Varadarajan, D., Yendiki, A., Salat, D. H., Fischl, B., … Initiative, A. D. N. (2021). Conductance-based structural brain connectivity in aging and dementia. *Brain connectivity*, *11*(7), 566–583.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., & Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS biology*, *6*(7), e159.

He, Y., Chen, Z., & Evans, A. (2008). Structural insights into aberrant topological patterns of large-scale cortical networks in alzheimer's disease. *Journal of Neuroscience*, *28*(18), 4756–4766.

Hofman, M. A. (2015). *Evolution of the human brain: from matter to mind*. Springer.

Hu, J., Cao, L., Li, T., Dong, S., & Li, P. (2021). Gat-li: a graph attention network based learning and interpreting method for functional brain network classification. *BMC bioinformatics*, *22*, 1–20.

Jellinger, K. A. (2022). Morphological basis of parkinson disease-associated cognitive impairment: an update. *Journal of Neural Transmission*, *129*(8), 977–999.

Jiao, T., Li, F., Cui, Y., Wang, X., Li, B., Shi, F., … Zeng, Q. (2023). Deep learning with an attention mechanism for differentiating the origin of brain metastasis using mr images. *Journal of Magnetic Resonance Imaging*, *58*(5),

1624–1635.

Karlsgodt, K. H., Sun, D., & Cannon, T. D. (2010). Structural and functional brain abnormalities in schizophrenia. *Current directions in psychological science*, *19*(4).

Kazi, A., Markova, V., Kondamadugula, P. R., Liu, B., Adly, A., Faghihroohi, S., & Navab, N. (2022). Dg-gru: dynamic graph based gated recurrent unit for age and gender prediction using brain imaging. In *Medical imaging 2022: Computer-aided diagnosis* (Vol. 12033, pp. 277–281).

Kipf, T. N., & Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.

Komer, B., Bergstra, J., & Eliasmith, C. (2019). Hyperopt-sklearn. *Automated machine learning: methods, systems, challenges*, 97–111.

Ktena, S. I., Parisot, S., Ferrante, E., Rajchl, M., Lee, M., Glocker, B., & Rueckert, D. (2018). Metric learning with spectral graph convolutions on brain connectivity networks. *NeuroImage*, *169*, 431–442.

LaMontagne, P. J., Benzinger, T. L., Morris, J. C., Keefe, S., Hornbeck, R., Xiong, C., … others (2019). Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *MedRxiv*, 2019–12.

Leoutsakos, J.-M., Gross, A., Jones, R., Albert, M., & Breitner, J. (2016). 'alzheimer's progression score': Development of a biomarker summary outcome for ad prevention trials. *The journal of prevention of Alzheimer's disease*, *3*(4), 229.

Lewis, J. D., O'Reilly, C., Bock, E., Theilmann, R. J., & Townsend, J. (2022). Aging-related differences in structural and functional interhemispheric connectivity. *Cerebral Cortex*, *32*(7), 1379–1389.

Li, Q., Han, Z., & Wu, X.-M. (2018). Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 32).

Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, J., … Duncan, J. S. (2021). Braingnn: Interpretable brain graph neural network for fmri analysis. *Medical Image Analysis*, *74*, 102233.

Liu, S., Park, J. H., & Yoo, S. (2020). Efficient and effective graph convolution networks. In *Proceedings of the 2020 siam international conference on data mining* (pp. 388–396).

Meunier, D., Achard, S., Morcom, A., & Bullmore, E. (2009). Age-related changes in modular organization of human brain functional networks. *Neuroimage*, *44*(3), 715–723.

Moon, H. S., Mahzarnia, A., Stout, J., Anderson, R. J., Han, Z. Y., Tremblay, J. T., … Badea, A. (2024). Feature attention graph neural network for estimating brain age and identifying important neural connections in mouse models of genetic risk for alzheimer's disease. *Imaging Neuroscience*, *2*, 1–22.

Neudorf, J., Shen, K., & McIntosh, A. R. (2024). Reorganization of structural connectivity in the brain supports preservation

of cognitive ability in healthy aging. *Network Neuroscience*, *8*(3), 837–859.

Pei, Y., Huang, T., Van Ipenburg, W., & Pechenizkiy, M. (2022). Resgcn: attention-based deep residual modeling for anomaly detection on attributed networks. *Machine Learning*, *111*(2), 519–541.

Phan, A. V., Le Nguyen, M., Nguyen, Y. L. H., & Bui, L. T. (2018). Dgcnn: A convolutional neural network over large-scale labeled graphs. *Neural Networks*, *108*, 533–543.

Prescott, J. W., Doraiswamy, P. M., Gamberger, D., Benzinger, T., Petrella, J. R., & Network, D. I. A. (2022). Diffusion tensor mri structural connectivity and pet amyloid burden in preclinical autosomal dominant alzheimer disease: The dian cohort. *Radiology*, *302*(1), 143–150.

Ranjbarzadeh, R., Bagherian Kasgari, A., Jafarzadeh Ghoushchi, S., Anari, S., Naseri, M., & Bendechache, M. (2021). Brain tumor segmentation based on deep learning and an attention mechanism using mri multi-modalities brain images. *Scientific Reports*, *11*(1), 1–17.

Reuter, M., Schmansky, N. J., Rosas, H. D., & Fischl, B. (2012). Within-subject template estimation for unbiased longitudinal image analysis. *Neuroimage*, *61*(4), 1402–1418.

Sebenius, I., Campbell, A., Morgan, S. E., Bullmore, E. T., & Liò, P. (2021). Multimodal graph coarsening for interpretable, mri-based brain graph neural network. In *2021 ieee 31st international workshop on machine learning for signal processing (mlsp)* (pp. 1–6).

Shi, F., Yap, P.-T., Gao, W., Lin, W., Gilmore, J. H., & Shen, D. (2012). Altered structural connectivity in neonates at genetic risk for schizophrenia: a combined study using morphological and white matter networks. *Neuroimage*, *62*(3), 1622–1633.

Sjöblom, L., Westerlund, D., Neimantaite, A. D., & Andersson, J. (2020). Braph for python: graph theory and graph neural networks for brain connectivity analysis. In *Emerging topics in artificial intelligence 2020* (Vol. 11469, p. 114691Y).

Sporns, O. (2011). The human connectome: a complex network. *Annals of the new York Academy of Sciences*, *1224*(1), 109–125.

Tolan, E., & Isik, Z. (2018). Graph theory based classification of brain connectivity network for autism spectrum disorder. In *Bioinformatics and biomedical engineering: 6th international work-conference, iwbbio 2018, granada, spain, april 25–27, 2018, proceedings, part i 6* (pp. 520–530).

Vaswani, A. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2017). Graph attention networks. *arXiv preprint arXiv:1710.10903*.

Wang, J., Khosrowabadi, R., Ng, K. K., Hong, Z., Chong, J. S. X., Wang, Y., … others (2018). Alterations in brain network

topology and structural-functional connectome coupling relate to cognitive impairment. *Frontiers in Aging Neuroscience*, *10*, 404.

Wang, S., Tang, H., Himeno, R., Solé-Casals, J., Caiafa, C. F., Han, S., … Sun, Z. (2024). Optimizing graph neural network architectures for schizophrenia spectrum disorder prediction using evolutionary algorithms. *Computer Methods and Programs in Biomedicine*, *257*, 108419.

Wang, Y.-m., Zhang, Y.-j., Cai, X.-l., Yang, H.-x., Cheung, E. F., Chan, R. C., et al. (2020). Altered grey matter volume and white matter integrity in individuals with high schizo-obsessive traits, high schizotypal traits and obsessive-compulsive symptoms. *Asian Journal of Psychiatry*, *52*, 102096.

Wein, S., Malloni, W. M., Tomé, A. M., Frank, S. M., Henze, G.-I., Wüst, S., … Lang, E. W. (2021). A graph neural network framework for causal inference in brain networks. *Scientific reports*, *11*(1), 8061.

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Philip, S. Y. (2020). A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, *32*(1), 4–24.

Xu, X., Wang, T., Li, W., Li, H., Xu, B., Zhang, M., … Xiao, S. (2021). Morphological, structural, and functional networks highlight the role of the cortical-subcortical circuit in individuals with subjective cognitive decline. *Frontiers in Aging Neuroscience*, *13*, 688113.

Zalesky, A., Cocchi, L., Fornito, A., Murray, M. M., & Bullmore, E. (2012). Connectivity differences in brain networks. *Neuroimage*, *60*(2), 1055–1062.

Zhang, F., Daducci, A., He, Y., Schiavi, S., Seguin, C., Smith, R. E., … O'Donnell, L. J. (2022). Quantitative mapping of the brain's structural connectivity using diffusion mri tractography: A review. *Neuroimage*, *249*, 118870.

Zhang, S., Tong, H., Xu, J., & Maciejewski, R. (2019). Graph convolutional networks: a comprehensive review. *Computational Social Networks*, *6*(1), 1–23.

Zhang, X., He, L., Chen, K., Luo, Y., Zhou, J., & Wang, F. (2018). Multi-view graph convolutional network and its applications on neuroimage analysis for parkinson's disease. In *Amia annual symposium proceedings* (Vol. 2018, p. 1147).

Zhang, Y., He, X., Chan, Y. H., Teng, Q., & Rajapakse, J. C. (2023). Multi-modal graph neural network for early diagnosis of alzheimer's disease from smri and pet scans. *Computers in Biology and Medicine*, *164*, 107328.

Zhang, Y., Teng, Q., He, X., Niu, T., Zhang, L., Liu, Y., & Ren, C. (2023). Attention-based 3d cnn with multi-layer features for alzheimer's disease diagnosis using brain images. In *2023 45th annual international conference of the ieee engineering in medicine & biology society (embc)* (pp. 1–4).

Zhang, Y., Zhan, L., Cai, W., Thompson, P., & Huang, H. (2019). Integrating heterogeneous brain networks for predicting

brain disease conditions. In *Medical image computing and computer assisted intervention–miccai 2019: 22nd international conference, shenzhen, china, october 13–17, 2019, proceedings, part iv 22* (pp. 214–222).

Zheng, K., Yu, S., Chen, L., Dang, L., & Chen, B. (2024). Bpi-gnn: Interpretable brain network-based psychiatric diagnosis and subtyping. *NeuroImage*, *292*, 120594.