Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Anatomical Interfaces as a Complementary Substrate for Structural MRI Representation in Alzheimer's Disease

Yongha Gi
Korea University
Hyeongjin Lim
Korea University
Saemin Park
Korea University
Jeongwon Lee
Korea University
Sangyoon Park
Korea University
A Hyun Jung
Korea University
Seol-Hee Baek
Korea University
Jong Hyun Kim
FieldCure Co
Byung-Jo Kim
Korea University
Myonggeun Yoon
radioyoon@korea.ac.kr

Korea University

Research Article

Keywords: Structural MRI, Alzheimer's disease, Image representation, Inter-regional interfaces, Brain morphometry, Graph representation

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11067124/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Background

Structural MRI analysis of Alzheimer’s disease (AD) commonly relies on region-of-interest (ROI) summaries or voxel-level deep networks, neither of which explicitly quantifies the physical interfaces shared by adjacent brain regions. We investigated whether these anatomical interfaces provide structural information complementary to conventional regional measures.

Methods

T1-weighted MRI was segmented into 79 anatomical regions using a pretrained nnU-Net. Each region was represented as a graph node with six morpho-textural descriptors, while edges were defined only between regions sharing voxel-level face contact and characterized by Structural Contact and Boundary Appearance features.

Results

Correlation analyses showed generally limited redundancy between node- and interface-level descriptors, supporting partially distinct representational information. Node-controlled analyses further showed that selected interface features remained significantly associated with disease stage after adjustment for the morpho-textural characteristics of both connected regions, indicating that interface-level variation was not fully reducible to regional abnormalities. White matter hypointensity interfaces repeatedly emerged among the retained stage-associated effects. Selected Structural Contact features also showed convergent associations with clinical severity across CDR-SB, ADAS-Cog13, and MMSE. Representation ablation demonstrated that disrupting anatomical adjacency or removing interface attributes reduced classification performance, supporting the contribution of anatomically organized relational information. The full representation also improved multi-class classification relative to conventional ROI-vector and voxel-level baselines, with this advantage preserved on the independent NACC cohort without fine-tuning.

Conclusions

These findings indicate that anatomical interfaces provide measurable and complementary structural information beyond conventional region-wise morphometry, offering an interpretable framework for representing inter-regional structural relationships in neurodegenerative disease.

1. Introduction

Structural T1-weighted MRI provides a non-invasive means of characterizing the anatomical alterations associated with Alzheimer's disease (AD), including regional atrophy, ventricular enlargement, and changes in cortical and subcortical structures [1–5]. Accordingly, a wide range of computational approaches have been developed to represent these structural abnormalities and distinguish cognitively normal (CN), mild cognitive impairment (MCI), and AD groups. A fundamental challenge, however, lies not only in the choice of predictive model but also in how brain anatomy is represented from MRI. Different representations preserve different aspects of structural organization, and clinically meaningful information may be lost when complex anatomical relationships are reduced to either isolated regional summaries or unconstrained voxel-level patterns.

Traditional structural MRI pipelines commonly segment the brain into anatomically defined regions of interest (ROIs) and characterize each region using quantitative measures such as volume or cortical thickness [6, 7]. Additional morphological and textural descriptors can further capture regional shape, complexity, and intensity heterogeneity. Such ROI-based representations are compact, anatomically interpretable, and readily related to established patterns of neurodegeneration. However, they typically treat each ROI as an independent unit. Compressing a region into a small number of summary measures consequently provides limited information about how neighboring anatomical structures are spatially arranged with respect to one another, including the geometry and appearance of the physical boundaries they share [3, 5, 8, 9]. Thus, alterations expressed at inter-regional interfaces may remain unrepresented even when the individual regions themselves are extensively characterized.

Voxel-level deep learning approaches offer an alternative by learning spatial representations directly from full MRI volumes or local image patches [10, 11]. Convolutional neural networks and vision transformers can capture complex spatial patterns without restricting the representation to predefined regional summaries. However, this flexibility comes at the cost of substantially higher dimensionality and reduced anatomical interpretability. Full-volume and patch-based models may require considerable computational capacity and can be susceptible to overfitting in the sample-size regime typical of neuroimaging studies. Furthermore, anatomical relationships are learned implicitly rather than represented explicitly, making it difficult to determine whether predictive information arises from regional abnormalities, relationships between structures, or other spatial patterns. Consequently, voxel-level models do not directly resolve the representational gap between interpretable ROI-wise measurements and the spatial organization of neighboring anatomical structures [12, 13].

Graph-based representations provide a natural framework for bridging these levels by modeling brain regions as nodes and relationships between them as edges [14–17]. Existing brain graph approaches have successfully incorporated structural, functional, covariance-based, or similarity-based relationships between regions. Nevertheless, ROI-based structural graphs often encode local anatomical relationships using binary adjacency or abstract statistical associations [18–20]. Even when two regions are known to be anatomically adjacent, the physical interface through which they meet is generally not quantified as a structural object in its own right. Consequently, the extent of contact between adjacent regions and the characteristics of their shared boundary remain largely absent from conventional ROI-graph

representations. This distinction is important because anatomical change may alter not only the size or internal properties of individual regions but also the geometry of the interfaces connecting neighboring structures.

To address this representational gap, we formulate structural brain anatomy as an ROI-level graph in which anatomical interfaces are explicitly quantified rather than represented as binary adjacency alone. Each ROI is encoded using complementary morpho-textural node attributes, whereas edges are restricted to regions sharing voxel-level face contact. For each such interface, we derive a scale-normalized Structural Contact measure reflecting the extent of physical contact between the connected regions and a Boundary Appearance descriptor reflecting local intensity-gradient characteristics. The resulting representation therefore jointly describes what each anatomical region is like and how adjacent regions physically relate to one another.

The central question of this study is whether these interface-level attributes provide structural information that is complementary to conventional regional measures. We evaluate this question from several perspectives: by examining the dependence between node- and interface-level descriptors, testing whether interface associations with disease stage persist after controlling for the characteristics of both connected ROIs, examining their anatomical and clinical relevance, assessing the contribution of anatomical topology and interface attributes through representation ablation, and finally evaluating classification performance and external generalization. Through these analyses, we investigate anatomical interfaces as a potentially distinct level of structural MRI representation for characterizing AD-related anatomical variation.

2. Materials and Methods

2.1 Dataset

For the development of the proposed model and its internal evaluation, we utilized T1-weighted MRI scans from the Alzheimer's Disease Neuroimaging Initiative phase 1 (ADNI1) cohort [21]. Data used in this study were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). ADNI1 is a large, multi-center longitudinal study that includes CN, MCI, and AD participants, providing representation across all three diagnostic groups for training and cross-validated performance assessment. All scans used in this study were acquired using standardized 3D Magnetization-Prepared Rapid Gradient Echo (MPRAGE) imaging protocols on 1.5T scanners across participating sites. Although acquisition parameters differed slightly between centers, voxel sizes typically approximated $1.1 \times 1.1 \times 1.2$ mm^3.

To evaluate generalization on an independent dataset, we further employed the National Alzheimer's Coordinating Center (NACC) dataset [22]. In contrast to ADNI1, the NACC cohort comprises data collected across multiple Alzheimer's Disease Research Centers (ADRCs), resulting in greater heterogeneity in scanner types, imaging protocols, and acquisition parameters. This inter-site variability

introduces a realistic cross-dataset domain shift, thereby providing a more stringent evaluation of model robustness and generalizability. T1-weighted MRI scans in NACC were acquired using site-specific protocols, with voxel resolutions generally comparable to those of ADNI1 but not fully standardized across centers.

For both cohorts, all subjects with available T1-weighted MRI and a corresponding CN, MCI, or AD diagnosis were included, without additional demographic or clinical screening. All MRI scans underwent N4 bias-field correction [23], intensity normalization, and resampling to an isotropic resolution of 1 mm^3, after which they were conformed to a standardized 256 × 256 × 256 voxel grid. A summary of diagnostic categories and demographic characteristics for both datasets is presented in Table 1.

Table 1
Summary of demographic characteristics and diagnostic group distribution for the ADNI1 and NACC cohorts. The table lists both the number and proportion of subjects by diagnosis and sex. Age values are presented as the mean ± standard deviation (SD), along with the minimum–maximum range.

Dataset	Subjects	Women, n(%)	Age (Range)	CN, n(%)	MCI, n(%)	AD, n(%)
ADNI1	639	267 (42)	75.3 ± 6.7 (55–90)	194 (30)	312 (49)	133 (21)
NACC	540	364 (67)	70.4 ± 8.2 (54–95)	293 (54)	191 (35)	56 (10)

2.2 Full Pipeline Overview

The overall pipeline is illustrated in Fig. 1. Structural T1-weighted MRI scans are first segmented into anatomical ROIs using a pretrained nnU-Net model [24]. From the resulting segmentation masks and corresponding MRI intensities, we extract ROI-level node features consisting of morphological and textural descriptors. In parallel, an anatomically structured topology is constructed by defining edges between ROIs that share voxel-level face contact under a strict 6-neighborhood connectivity criterion. For each anatomically adjacent ROI pair, edge features are then derived by jointly leveraging the segmentation labels and MRI intensities to quantify structural contact and boundary appearance at the shared anatomical interface. By integrating the node features, edge features, and anatomically structured topology, we construct an anatomically informed ROI-level graph. This graph serves as the input to an edge-aware graph neural network (GNN) [25], which is subsequently trained to classify subjects into CN, MCI, and AD groups.

2.3 ROI Segmentation and Node feature extraction

Whole-brain segmentation was performed using a pretrained nnU-Net model trained on 2,378 scans from the OASIS-3 dataset [26], with labels derived from the Desikan–Killiany atlas. Following the strategy employed in FastSurfer [27], ROIs known to exhibit high noise or instability were merged bilaterally, resulting in a total of 79 anatomical regions spanning cortical, subcortical, and CSF compartments. The names of the 79 anatomical regions are provided in Supplementary Table 1. The overall segmentation

performance of the nnU-Net model, along with the performance for key ROIs, is reported in Supplementary Table 2.

Based on the resulting segmentation and the corresponding MR images, a set of six quantitative descriptors was extracted for each ROI to capture complementary morphological and textural characteristics. These six measurements were concatenated to form a 6-dimensional feature vector for each ROI node in the graph.

2.3.1 Morpho Features

Morphological features—volume, thickness, and fractal dimension—were computed for each ROI using the segmentation generated by the pretrained nnU-Net model.

1. Volume: Voxel count of each ROI.
2. Thickness: A generalized surface-based distance measure was applied to all ROIs. Boundary vertices of each ROI were matched with nearby boundary points located on the opposite side using KD-tree search, and the distance along the local surface normal direction was used to estimate ROI-wise thickness [28, 29].
3. Fractal Dimension: The fractal dimension of each ROI was computed using a 3D box-counting method. For multiple box sizes s, the number of occupied boxes $N(s)$within the ROI's bounding volume was estimated, and the slope of the linear fit between $\log N(s)$and $\log(1/s)$was taken as the fractal dimension [30].

2.3.2 Textural Features

Textural features were computed using the input T1-weighted MRI and the ROI segmentation generated by the pretrained nnU-Net model.

1. Intensity: Mean T1 intensity within each ROI.

1. Gradient: A 3D gradient-magnitude image was computed using Sobel-based edge filtering. For each ROI, gradient values were sampled from a thin boundary band obtained by dilating and eroding the segmentation mask. The sampled values were fitted with a beta distribution, and the shape parameter $b$was used as the gradient feature for that ROI [31].
2. Entropy: For each ROI, voxel intensities were discretized into a fixed number of bins to construct a histogram, from which the probability mass function (PMF) was estimated. Shannon entropy was then computed from the PMF to quantify the degree of intensity heterogeneity within each region [32].

2.4 ROI Adjacency Graph Construction with Structural and Boundary-Based Features

Using the same segmentation generated by the pretrained nnU-Net model, we constructed an anatomically informed ROI adjacency graph. Each of the 79 ROIs was represented as a node, and two ROIs were defined as adjacent only when they shared at least one voxel-level face contact, following a strict 6-neighborhood connectivity criterion in the 3D segmentation mask. For every anatomically adjacent ROI pair, we computed two edge descriptors capturing inter-regional contact strength and boundary appearance. These quantities were stored as a 2-dimensional edge feature vector for each undirected edge. The resulting graph typically contained approximately 400–450 edges, depending on inter-ROI adjacency.

2.4.1 Structural Contact Feature

For each adjacent ROI pair (i, j), we first computed the number of face-contacting voxel pairs, denoted as shared boundary voxels:

$$SB(i,j) = \#\left\{\text{face-adjacent}\backslash\text{:voxel}\backslash\text{:pairs}\backslash\text{:between}\backslash\text{:ROIs}\backslash\text{:}i\backslash\text{:and}\backslash\text{:}j\right\}.$$

We further defined an ROI-wise “ROI–ROI surface exposure” term by accumulating all boundary contacts incident to each ROI:

$$S\left(i\right) = \sum\nolimits_{k\in\,\mathcal{N}(i)} SB(i,k),$$

where $\mathcal{N}\left(i\right)$is the neighbor set of ROI i.

To obtain a scale-normalized contact strength, we applied a square-root normalization to compute the contact weight:

$$w_{\text{contact}}(i,j) = \frac{SB(i,j)}{\sqrt{S\left(i\right)\,\backslash\text{hspace}\{0.17\text{em}\}S\left(j\right)}}.$$

This weighting reduces the dominance of large regions and emphasizes relatively strong inter-ROI contacts given their boundary exposure [33].

2.4.2 Boundary Appearance Feature

When the corresponding T1-weighted MRI was available, we computed additional boundary-based appearance descriptors to capture contrast and sharpness at the interface between adjacent ROIs. The MRI intensity was first robustly normalized by clipping values above the 99.5th percentile and linearly rescaling the intensities to the range [0,1]. A 3D gradient-magnitude image was then computed using finite differences.

For each contacting ROI pair (i, j), boundary voxels contributing to $SB(i, j)$were used to estimate the boundary gradient strength:

$$g_{\text{boundary}}(i,j) = \frac{1}{2}\left(g_i^{'} + g_j^{'}\right),$$

where $g_i^{'}$ and $g_j^{'}$ denote mean gradient magnitudes on each side of the contacting boundary [34].

2.5 Graph Representation

For downstream learning, the structural brain graph was encoded as a PyTorch Geometric (PyG) Data object. The node feature matrix $X \in \mathbb{R}^{79\times 6}$ consisted of six ROI-level descriptors: volume, thickness, fractal dimension, intensity, gradient, and entropy.

Anatomical adjacency between ROIs was represented as a bidirectional adjacency list, $\text{edge_index} \in \mathbb{N}^{2\times 2E}$, where E denotes the number of unique unordered ROI pairs defined by the strict 6-neighborhood connectivity criterion [35].

The edge attribute matrix $\mathbf{E} \in \mathbb{R}^{2E\times 2}$ encoded a two-dimensional feature vector for each edge, composed of the scale-normalized Structural Contact weight and the Boundary Appearance feature described in Section 2.4.

By integrating morphological node descriptors with anatomically grounded edge features, the resulting ROI graph provides a compact yet interpretable representation of structural brain organization. This graph was used as input to an edge-aware graph neural network that jointly leverages node- and edge-level information. Through message passing across anatomically adjacent ROIs, the model learns spatially contextualized structural patterns relevant to AD classification.

2.6 Statistical Analysis of Node- and Edge-Level Features

To assess redundancy between node- and edge-derived descriptors, Spearman's rank correlations were calculated among the pooled features across ROIs and subjects in the ADNI1 cohort, and absolute correlation coefficients were used for visualization.

Disease-stage differences in node- and edge-level features were evaluated across the CN, MCI, and AD groups using the Kruskal–Wallis H test. For each feature, effect size was quantified using eta squared (η^2) derived from the Kruskal–Wallis statistic:

$$\eta^2 = \frac{H - k + 1}{n - k},$$

where H denotes the Kruskal–Wallis statistic, k is the number of diagnostic groups, and n is the number of observations included in the corresponding analysis.

For edge-level analyses, only anatomical interfaces present in at least 70% of subjects were included to reduce instability arising from sparsely observed connections. Statistical significance was corrected for multiple comparisons using the Benjamini–Hochberg false discovery rate (FDR) procedure, with

$q < 0.05$ considered significant. Effect sizes of η² ≥ 0.06 and η² ≥ 0.14 were considered moderate and large, respectively, for visualization.

To characterize the direction of disease-stage–related change, diagnostic stage was ordinally coded as CN = 0, MCI = 1, and AD = 2, and Spearman's rank correlation coefficient was calculated between each feature and the ordinal stage variable. Positive coefficients therefore indicate increasing feature values across the CN–MCI–AD continuum, whereas negative coefficients indicate decreasing values.

2.7 Node-controlled edge analysis

To determine whether stage-associated edge-level effects merely reflected regional (node-level) alterations, we performed a node-controlled edge analysis for each anatomically adjacent ROI pair [36].

For each edge (i, j), the edge feature value was modeled as:

$$\mathrm{EdgeFeature}_{ij} = \beta_0 + \beta_{\mathrm{stage}} \cdot \mathrm{stage} + \gamma^{\top} x_i + \delta^{\top} x_j + \epsilon,$$

where x_i and x_j denote the six-dimensional node feature vectors (volume, thickness, fractal dimension, intensity, gradient, entropy) of the incident ROIs. The disease stage (CN→MCI→AD) was encoded as an ordinal variable. This framework isolates the unique contribution of disease stage to each edge feature while adjusting for regional morpho-textural characteristics.

For each edge, the effect size of the stage term was quantified using partial η² [37]. Statistical significance was assessed across all tested edges with false discovery rate (FDR) correction (q < 0.05). Edges exceeding predefined effect-size thresholds were visualized in the corresponding network figures.

2.8 Clinical association analysis

To assess the clinical relevance of interface-level features, we examined associations between Structural Contact and three clinical measures: CDR-SB, ADAS-Cog13, and MMSE. For each anatomical interface, a node-controlled regression model was fitted with the clinical score as the dependent variable and Structural Contact as the predictor of interest. Age, sex, diagnosis, and the six morpho-textural features of each of the two connected ROIs were included as covariates, with diagnosis entered as a categorical variable.

Standardized regression coefficients were calculated to facilitate comparison across clinical outcomes. Because higher CDR-SB and ADAS-Cog13 scores indicate greater impairment whereas lower MMSE scores indicate greater impairment, the sign of the standardized MMSE coefficient was reversed to align all effect directions with clinical severity. These severity-aligned standardized coefficients are denoted as β*. Accordingly, a positive β* indicates that greater Structural Contact is associated with greater clinical severity, whereas a negative β* indicates that reduced Structural Contact is associated with greater clinical severity.

The unique contribution of each interface beyond the covariates was quantified using partial R^2. Statistical significance was corrected for multiple comparisons using false discovery rate (FDR) correction, with $q < 0.05$ considered significant. Interfaces that showed significant associations across all three clinical measures and achieved partial $R^2 \geq 0.01$ for each outcome were retained as convergent clinical associations.

2.9 Primary Classifier Architectures for Comparative Evaluation

To evaluate the effectiveness of the proposed anatomically informed ROI-graph representation, we conducted comparative experiments against two primary baseline models under a consistent evaluation framework.

2.9.1 ROI-Vector Baseline (Non-Graph MLP)

To evaluate the contribution of explicit anatomical adjacency modeling, we constructed a non-graph baseline in which all ROI-wise node features (79 regions × 6 features) were concatenated into a 474-dimensional subject-level feature vector. This vector was input to a multilayer perceptron (MLP) classifier with two fully connected layers and ReLU activation [38].

This baseline assesses classification performance when regional morpho-textural information is used without incorporating inter-regional structural relationships.

2.9.2 3D DenseNet Baseline (Voxel-Level Model)

As an image-based baseline, we trained a compact 3D DenseNet model directly on the preprocessed T1-weighted MRI volumes [39]. This model operates at the voxel level and learns spatial patterns implicitly from the full 3D image, without relying on ROI segmentation or explicit anatomical adjacency.

The DenseNet architecture consisted of an initial convolutional embedding layer followed by densely connected convolutional blocks and a global pooling layer, culminating in a fully connected classification head for diagnostic prediction.

2.9.3 Proposed Anatomically Informed ROI-Level Graph Classifier

For the proposed method, each subject was represented as an anatomically informed ROI-level graph as described in Sections 2.3–2.5. Classification was performed using an edge-aware GNN implemented in a message-passing framework [25]. In this architecture, node embeddings were iteratively updated by aggregating information from anatomically adjacent ROIs together with their corresponding edge attributes, enabling joint modeling of regional morpho-textural features and inter-regional structural relationships.

After two edge-aware graph convolution layers, global pooling was applied to obtain a subject-level graph representation. This representation was then passed through a fully connected MLP classifier to predict diagnostic labels.

2.10 Ablation Settings and Evaluation Protocol

To determine which components of the proposed representation contribute to classification performance, we conducted a series of targeted ablation experiments. The analyses were designed to distinguish the contribution of (1) relational message passing, (2) anatomically meaningful topology, and (3) quantitative interface attributes. Except for the non-graph tabular baseline, all graph-based variants used the same GNN architecture, training procedure, and hyperparameter settings as the proposed model. Thus, differences among the graph variants primarily reflect changes in the relational representation supplied to the network rather than differences in model capacity. The full anatomically informed graph described in Section 2.9.3 served as the reference condition against which the following ablation settings were compared.

The following ablation settings were evaluated.

2.10.1 Tabular MLP (Node + Edge Summary)

To evaluate whether the proposed graph representation provides information beyond the availability of node- and edge-derived features, the graph features were converted into a fixed-length subject-level vector and classified using an MLP. This setting removes explicit graph topology and message passing while retaining the feature content in a non-relational form. It therefore serves as a non-graph reference for assessing whether organizing the same anatomical information within a relational structure provides additional predictive value.

2.10.2 Self-Loop Graph

1. To evaluate the contribution of inter-regional message passing, a graph was constructed in which each ROI was connected only to itself. Node features were preserved, but information exchange between different anatomical regions was eliminated. Because the same GNN architecture was retained, this setting tests whether message passing across distinct ROIs contributes beyond node-wise processing within the graph architecture itself.

2.10.4 Complete Graph

To determine whether message passing alone is sufficient in the absence of anatomically meaningful topology, all ROI pairs were connected to form a complete graph. This configuration permits unrestricted inter-regional information exchange but removes the anatomical constraint that limits communication to physically adjacent regions. The comparison between the complete and anatomically structured graphs therefore assesses whether biologically grounded adjacency provides value beyond generic relational connectivity.

2.10.5 Permuted Adjacency Graph

To assess whether performance depends specifically on anatomically meaningful topology rather than graph connectivity itself, the anatomical adjacency structure was randomly permuted while preserving the overall graph density. This procedure disrupts the correspondence between graph connections and true anatomical neighborhood relationships while maintaining a comparable amount of relational connectivity.

2.10.6 Anatomical Graph Without Interface Attributes

To isolate the contribution of quantitative interface attributes, the anatomically defined adjacency structure was retained while Structural Contact and Boundary Appearance features were removed. Message passing was therefore performed according to binary anatomical adjacency using only ROI-level node features. This setting evaluates whether the quantitative properties of shared anatomical interfaces provide additional information beyond knowledge of which regions are anatomically adjacent.

2.11 Implementation Details

For the whole-brain segmentation process described in Section 2.3, the nnU-Net model was trained on 2,378 scans from the OASIS-3 dataset using its standard 3D full-resolution pipeline. The network was optimized with stochastic gradient descent with Nesterov momentum (0.99), an initial learning rate of 0.01 following a polynomial decay schedule, and a weight decay of 3×10^{-5}. The objective function combined soft Dice loss and cross-entropy loss, with deep supervision applied at multiple resolution levels. Patch size and batch size were automatically determined by the nnU-Net planner based on available GPU memory.

To maintain hemispheric asymmetry and ensure stable anatomical alignment, mirror-based data augmentation was disabled during training, and all MR images were spatially aligned according to their affine matrices prior to preprocessing. This configuration improved the consistency of medial temporal structures compared to the default augmentation settings. The trained model was subsequently applied to the ADNI1 and NACC datasets without additional fine-tuning.

For the graph-based classification model, all node and edge features were processed using PyG. The graph-based classifier described in Section 2.9 was implemented in PyG. Models were trained using the Adam optimizer (learning rate: 1×10^{-3}, weight decay: 1×10^{-5}), with class-weighted cross-entropy applied to mitigate class imbalance in the ADNI1 cohort.

Model evaluation was performed using stratified five-fold cross-validation at the subject level. For each outer fold, one fold was held out exclusively for final evaluation, while the remaining subjects constituted the training portion. A separate validation subset was split from the training portion and used only for early stopping based on validation loss. The held-out test fold was not used for early stopping, hyperparameter selection, or any other model-selection procedure. Feature standardization parameters were estimated using only the training data of each fold and were subsequently applied to the

corresponding validation and test data. The same outer-fold assignments were used for all compared models to ensure paired evaluation across methods.

For the non-graph baseline, the six ROI-wise features (79 regions × 6 features = 474 dimensions) were concatenated into a single subject-level feature vector. The MLP classifier consisted of two fully connected layers with ReLU activation, with a dropout rate of 0.2 applied after the hidden layer. The model was trained using the same optimizer, learning rate, weight decay, and early stopping strategy as the GNN to ensure a fair comparison.

As an image-based baseline, the 3D DenseNet model described in Section 2.9.2 was trained directly on the preprocessed T1-weighted MRI volumes. The model was optimized using AdamW (learning rate: 5×10^{-4}) with cosine decay scheduling and a batch size of 8.

All experiments were implemented in Python (v3.8.19), PyTorch (v2.0.1), and the PyTorch Geometric (v2.3). Computations were performed on an NVIDIA RTX 4090 GPU.

3 Results

1.1 Complementarity of Node and Edge Feature Representations

Figure 2 presents the absolute Spearman correlation heatmaps computed to evaluate redundancy and complementarity between node- and edge-derived features. Correlations were calculated using the ADNI1 cohort in a pooled manner across all ROIs and subjects.

At the feature-family level (Fig. 2a), morphological, textural, and edge-derived features showed generally low inter-group correlations. The association between morphological and textural features was modest (~ 0.3), indicating limited overlap between structural size-related and intensity-based descriptors. Edge features exhibited weak correlations with both morphological and textural families, particularly with textural features. Although edge–morphological correlation was slightly higher, the overall magnitude remained low. This pattern suggests that edge-derived descriptors encode information not directly captured by conventional ROI-wise morphology or texture measures.

At the individual feature level (Fig. 2b), correlations were generally weak across most node features, with only limited associations observed among certain morphological measures. Volume and fractal dimension showed modest correlation, reflecting partially shared geometric properties. Structural Contact demonstrated a comparatively stronger association with thickness (~ 0.6) than with other node features, suggesting a plausible coupling between regional morphology and inter-regional interface strength. However, beyond this thickness-related association, correlations between edge features and the remaining node features were generally low. Structural Contact and Boundary Appearance also exhibited weak mutual correlation (~ 0.2).

Overall, node- and edge-derived features occupy partially distinct representational subspaces with limited redundancy. The observed complementarity supports the proposed anatomically informed graph formulation, in which intra-regional attributes and inter-regional structural relationships contribute to disease classification performance.

1.2 Edge-Level Structural Associations with Disease Stage

Figure 3 presents edge-level networks derived from the proposed ROI graph, restricted to large effect sizes ($\eta^2 \geq 0.14$) after FDR correction ($q < 0.05$). This threshold isolates the most robust disease-stage–associated structural interfaces. Edge networks at the moderate threshold ($\eta^2 \geq 0.06$) are shown in Supplementary Fig. 1. Notably, no Boundary Appearance edges reached the large-effect threshold under the raw model, indicating that strong group differences at this magnitude were specific to Structural Contact features.

As shown in Fig. 3a, large-effect Structural Contact edges form a sparse but anatomically structured network rather than a diffuse pattern. The surviving connections are concentrated among regions strongly implicated in AD pathology, including the hippocampus (left: ROI 13; right: ROI 28), amygdala (left: ROI 14; right: ROI 29), cortical white matter (left: ROI 1; right: ROI 19), inferior lateral ventricles (left: ROI 3; right: ROI 21), and external CSF (ROI 79). Negative associations are observed between medial temporal structures—hippocampus and amygdala—and cortical white matter symmetrically across hemispheres, indicating reduced structural contact strength with advancing disease stage. In contrast, the adjacency axis linking the inferior lateral ventricles and external CSF shows positive associations with advancing disease stage. This configuration is consistent with established pathological processes in AD, including medial temporal atrophy, ventricular enlargement, and remodeling of surrounding white matter. The spatial concentration of large-effect edges along this ventricular–limbic–white matter axis suggests that strong Structural Contact alterations are selectively aligned with a characteristic anatomical pattern of the disease.

After applying the node-controlled edge analysis framework described in Section 2.7—where morpho-textural features of both incident ROIs were included as covariates—several large-effect Structural Contact edges remained statistically significant (Fig. 3b). Here, partial η^2 quantifies the association between edge features and disease stage after accounting for regional morpho-textural properties.

Within the node-controlled Structural Contact network, several large-effect interfaces remained significant. White matter hypointensities (ROI 33) were involved in multiple retained connections, including interfaces with the right putamen (ROI 26), pars triangularis (ROI 51), and left paracentral cortex (ROI 48). Additional large-effect interfaces were observed between the right ventral diencephalon (ROI 31) and insula (ROI 64), the left pericalcarine cortex (ROI 52) and right lingual cortex (ROI 69), and the superior temporal cortex (ROI 61) and right precentral cortex (ROI 76).

Figure 3c further shows that Boundary Appearance also retains large-effect edges within the node-controlled network. Similar to Structural Contact, white matter hypointensities (ROI 33) again appear

centrally involved, showing significant associations with the putamen (left: ROI 8), paracentral cortex (right: ROI 72), and lateral occipital cortex (ROI 42). Additional large-effect connections are observed between the superior frontal cortex (left: ROI 59) and posterior cingulate cortex (right: ROI 75), as well as between the accumbens (left: ROI 16) and lateral orbitofrontal cortex (left: ROI 43).

The persistence of these edges after explicit adjustment for regional morpho-textural characteristics indicates that the edge-level changes observed across disease stages cannot be fully explained by atrophy or morpho-textural alterations within individual ROIs alone. Rather than representing a simple byproduct of regional degeneration, these findings suggest that stage-related differences include additional interface-level structural alterations between adjacent regions. The repeated emergence of white matter hypointensities (ROI 33) across both edge-feature types further indicates that interfaces surrounding white matter abnormality–related regions constitute a consistent locus of stage-associated structural coupling alteration along the CN–MCI–AD continuum.

1.3 Anatomical Visualization of Node- and Interface-Level Effects

To facilitate anatomical interpretation of node- and interface-level structural variation, we projected node- and edge-level effect sizes onto anatomical MRI, as shown in Fig. 4. Effect sizes were quantified using η^2 values derived from group-wise statistical analysis across the CN–MCI–AD continuum. Only effects exceeding the moderate threshold ($\eta^2 \geq 0.06$) were visualized.

Figure 4a illustrates node-derived structural effect-size maps, in which ROI-wise η^2 values were projected onto the corresponding anatomical volumes. Prominent node-level effects were observed in medial temporal and subcortical regions, including the hippocampus and amygdala, as well as in periventricular and CSF-related structures. In particular, external CSF exhibited relatively strong effect sizes, consistent with global atrophic expansion across the AD continuum. The spatial distribution of node-level effects aligns with established pathological trajectories of AD, reflecting region-specific structural alterations.

Figure 4b presents edge-derived structural effect-size maps, where η^2 values of anatomically adjacent ROI pairs were projected onto their voxel-level contact surfaces. In contrast to node-level maps that emphasize entire regional volumes, the edge-derived visualization localizes effects to anatomical interfaces between regions. Notably, many of the highlighted interfaces were situated along the medial temporal–periventricular regions, overlapping partially with regions exhibiting strong node-level effects. However, the edge maps reveal a more localized and boundary-specific pattern, particularly at white matter–ventricular and limbic interfaces. These findings suggest that disease-related structural information is expressed not only in region-wise degeneration but also along anatomical interfaces between adjacent regions.

Taken together, the node- and edge-derived visualizations demonstrate partially overlapping yet complementary anatomical patterns. While node-level maps primarily capture region-specific structural alterations—including pronounced CSF expansion—edge-level maps emphasize interaction-specific

changes at anatomical boundaries. This distinction supports the premise that organizing MRI-derived features within an anatomically informed graph enables simultaneous modeling of intra-regional degeneration and inter-regional structural relationships, providing a richer and more interpretable representation of disease-related structural remodeling.

Figure 5 visualizes edge-level partial η^2 effects derived from the node-controlled edge analysis framework described in Section 2.7. In this analysis, morpho-textural features of both incident ROIs were included as covariates in the regression model prior to effect-size estimation. The resulting partial η^2 values therefore quantify the association between each edge feature and disease stage after accounting for regional morpho-textural properties. Effect sizes were computed across the CN–MCI–AD continuum, and only effects exceeding the moderate threshold (partial $\eta^2 \geq 0.06$) were retained for visualization.

In the Structural Contact map (Fig. 5a), significant partial η^2 effects appear spatially sparse and highly localized, forming small island-like clusters rather than continuous boundary patterns. Although several effects are situated near ventricular–white matter and medial temporal interfaces, they do not extend uniformly along these boundaries. Compared with the broader raw edge patterns shown in Fig. 4b, the node-controlled map reveals a markedly reduced and spatially fragmented subset of interfaces. This pattern indicates that only specific portions of geometric contact surfaces exhibit stage-related variation beyond what can be explained by nodal atrophy or intensity changes.

Similarly, the Boundary Appearance map (Fig. 5b) shows a sparse distribution of effects, again manifesting as discrete focal clusters along anatomical interfaces. While fewer edges survive node-level adjustment, the retained effects remain predominantly located near periventricular and deep white matter regions. The focal and discontinuous nature of these clusters suggests that intensity-gradient–based interface differentiation reflects localized interface remodeling rather than diffuse boundary-wide changes.

Taken together, Fig. 5 demonstrates that only a restricted subset of interface regions retains significant stage associations after node-feature control. Whereas Fig. 4 displays more distributed edge effects, Fig. 5 isolates small, spatially confined coupling alterations that cannot be reduced to intra-regional degeneration alone. This localized pattern supports a more selective interpretation of interface-level structural alteration across the AD continuum.

1.4 Clinical associations with disease severity

To evaluate the clinical relevance of the interface-level features identified in the graph-based analysis, we examined their associations with three complementary measures of disease severity: CDR-SB, ADAS-Cog13, and MMSE. For each anatomical interface, node-controlled regression models were used to assess whether Structural Contact remained associated with clinical severity after accounting for age, sex, diagnosis, and the six morpho-textural features of both incident ROIs. To facilitate comparison across clinical measures, regression coefficients were expressed as severity-aligned standardized

coefficients (β*); the direction of the MMSE coefficient was reversed because lower MMSE scores indicate greater cognitive impairment.

Eight Structural Contact interfaces showed convergent associations across all three clinical measures, satisfying FDR significance ($q < 0.05$) and a partial $R^2 \geq 0.01$ for each outcome (Fig. 6 and Table 2). All eight interfaces exhibited concordant directions of association across CDR-SB, ADAS-Cog13, and MMSE, indicating that the observed relationships were not specific to a single clinical scale.

Among these interfaces, the Fusiform–left Parahippocampal connection showed a consistent negative association with clinical severity, with severity-aligned β* values of − 0.106 for CDR-SB, − 0.185 for ADAS-Cog13, and − 0.138 for MMSE. The corresponding partial R^2 values were 0.013, 0.034, and 0.017, respectively, indicating that reduced Structural Contact at this interface was associated with greater clinical impairment even after controlling for regional morpho-textural characteristics.

Several interfaces involving external CSF showed the opposite pattern, with greater Structural Contact associated with greater clinical severity. The left Parahippocampal–external CSF interface demonstrated positive associations across ADAS-Cog13 (β* = 0.147, partial R^2 = 0.035), CDR-SB (β* = 0.091, partial R^2 = 0.016), and MMSE (β* = 0.107, partial R^2 = 0.016). The right Ventral DC–external CSF interface showed the largest unique contribution of any interface for ADAS-Cog13 (β* = 0.142, partial R^2 = 0.038). Additional convergent positive associations were observed for the Inferior Temporal–external CSF, right Parahippocampal–external CSF, right Lateral Orbitofrontal–external CSF, and right cortical white matter–right Caudate interfaces.

This directional split warrants caution. Because cortical retreat mechanically increases contact with surrounding CSF, associations at cortical–CSF interfaces may partly reflect regional atrophy that the volume of a single global external-CSF node cannot fully control for. The three interfaces not involving external CSF — Fusiform–Parahippocampal, the interhemispheric Posterior Cingulate connection, and right cortical white matter–right Caudate — are not subject to this specific concern. In particular, the interhemispheric Posterior Cingulate connection, which spans the midline between hemispheres, is difficult to attribute solely to volumetric reduction in either region.

Table 2
Convergent node-controlled associations of Structural Contact with clinical severity. β* denotes the severity-aligned standardized regression coefficient; MMSE coefficient signs were reversed so that positive β* consistently indicates greater clinical severity.

Anatomical interface	ADAS β*	Partial R^2	CDR β*	Partial R^2	MMSE β*	Partial R^2
Fusiform – Parahippocampal (lh)	−0.185	0.034	−0.106	0.013	−0.138	0.017
Cortical white matter (rh) – Caudate (rh)	0.107	0.013	0.087	0.010	0.120	0.014
Posterior Cingulate (lh) – Posterior Cingulate (rh)	−0.099	0.020	−0.066	0.011	−0.079	0.011
Parahippocampal (lh) – external CSF	0.147	0.035	0.091	0.016	0.107	0.016
Ventral DC (rh) – external CSF	0.142	0.038	0.075	0.012	0.082	0.011
Inferior Temporal – external CSF	0.137	0.027	0.091	0.013	0.122	0.017
Parahippocampal (rh) – external CSF	0.115	0.022	0.082	0.013	0.088	0.011
Lateral Orbitofrontal (rh) – external CSF	0.078	0.010	0.082	0.013	0.085	0.010

To place these effect sizes in context, we computed partial R^2 for established volumetric markers in the same regression framework (Supplementary Table 3). Classical volumetric markers remained strong correlates of clinical severity in this cohort: bilateral hippocampal volume reached partial $R^2 = 0.124$ for ADAS-Cog13, exceeding every interface effect. Interface attributes are therefore not proposed as a substitute for volumetric measurement. The relevant comparison is nonetheless region-specific rather than global, and it divides the eight interfaces into two groups.

For three interfaces, the interface exceeded the volumes of both regions it connects. Most strikingly, right ventral diencephalon volume explained essentially no variance (partial $R^2 = 0.002$) while its interface with external CSF yielded the largest effect in Table 2 (0.038). The interhemispheric posterior cingulate interface (0.020) likewise exceeded the volume of either posterior cingulate region (0.012 for both), as did the right cortical white matter–caudate interface (0.013 vs. 0.008 and 0.006). These are precisely the regions in which volume itself is a weak marker.

For the remaining five interfaces, regional volume explained more variance than the interface — for example, inferior temporal volume (0.095) versus its interface with external CSF (0.027). These regions lie along the medial temporal atrophy axis, where volumetric measurement is already well established.

This division suggests that interface attributes do not duplicate volumetric information but extend it to regions where volume is comparatively uninformative. Notably, the comparison is asymmetric in favour

of volume: interface values were obtained under stricter adjustment, including the volumes of both incident regions as covariates, whereas the volumetric markers were adjusted for age, sex, and diagnosis only.

1.5 Representation ablation

To isolate what the proposed representation contributes beyond feature content and model capacity, we compared five alternative representation settings with the full model on the multi-class task (CN vs. MCI vs. AD). The alternatives included a non-graph flattened representation, three graph configurations that removed anatomically meaningful inter-regional structure through self-loops, complete connectivity, or permuted adjacency, and an anatomical graph without interface attributes. The five GNN-based configurations—including the full model—shared an identical architecture, hyperparameter setting, and parameter count (53,244), differing only in the edge index and edge attributes supplied to the network. Differences among these configurations therefore reflect the representation rather than architectural capacity. Results are summarized in Table 3.

Table 3
Representation ablation for multi-class disease classification (CN vs. MCI vs. AD).

Topology	Message passing	Anatomical	Interface attr.	Evaluation metric			Params
				AUC	Macro-F1	BACC	
Flattened (MLP)	–	–	–	0.689 ± 0.050	0.512 ± 0.037	0.536 ± 0.039	2.01M
Self-loops only	□	–	–	0.669 ± 0.041	0.470 ± 0.028	0.475 ± 0.027	53k
Complete	✓	□	–	0.697 ± 0.047	0.468 ± 0.037	0.473 ± 0.038	53k
Permuted	✓	□	–	0.681 ± 0.049	0.475 ± 0.045	0.476 ± 0.044	53k
Anatomical	✓	✓	□	0.698 ± 0.038	0.486 ± 0.032	0.492 ± 0.028	53k
Anatomical	✓	✓	✓	**0.745 ± 0.026**	**0.562 ± 0.035**	**0.600 ± 0.044**	53k

The three structure-free settings converged on a similar performance level despite removing structure through different mechanisms (macro-F1 0.468–0.475, BACC 0.473–0.476). Disabling message passing entirely, supplying a complete graph in which every region pair is connected uniformly, and permuting anatomical adjacency yielded closely similar performance within the observed fold-level variability. This convergence indicates that the graph neural network provided no clear advantage in the absence of anatomically meaningful structure: when the topology carries no anatomical information, the model performs at a level comparable to a network without meaningful relational modeling.

Introducing anatomical adjacency as binary connections increased macro-F1 to 0.486 and BACC to 0.492, while attaching quantitative interface attributes to the same topology further increased them to 0.562 and 0.600, respectively. The full representation therefore exceeded the structure-free settings by approximately 0.09 in macro-F1 and 0.13 in BACC, with a larger proportion of the gain attributable to the addition of interface attributes than to topology alone.

A dissociation was observed between ranking performance and class-wise balance. Anatomical adjacency alone left AUC essentially unchanged relative to the complete graph (0.698 vs. 0.697) while improving macro-F1 and BACC, suggesting that anatomical topology primarily improves the balance of predictions across diagnostic classes rather than the model's ability to rank subjects. The gain in discriminative performance emerged only when interface attributes were incorporated into the anatomical topology (AUC 0.745). These two components therefore appear to contribute through distinct mechanisms: topology organizes the relational structure, whereas interface attributes provide additional discriminative information.

Finally, the non-graph baseline received the same node- and edge-derived feature content in flattened form and required approximately 38 times more parameters, yet it did not exceed the proposed representation on any metric (macro-F1 0.512 vs. 0.562; BACC 0.536 vs. 0.600). Feature availability alone therefore does not account for the observed performance; rather, the anatomical organization of those features within a relational representation contributes to the performance gain.

Taken together, these ablations indicate that the advantage of the proposed formulation arises from organizing morpho-textural features within anatomically faithful relational structure, rather than from the expressive capacity of the graph neural network alone or from the mere availability of additional edge-derived features.

3.6 Overall Classification Performance

We evaluated the proposed anatomically informed ROI-level graph model against the two baseline architectures described in Section 2.9 on the ADNI1 dataset. Performance was assessed across four diagnostic tasks (CN vs. AD, CN vs. MCI, MCI vs. AD, and CN vs. MCI vs. AD) using 5-fold cross-validation. Area under the receiver operating characteristic curve (AUC), F1 score, and accuracy-based metrics were used for evaluation. For binary classification tasks, F1 denotes the positive-class F1 score and ACC denotes overall accuracy. For the three-class task, AUC denotes the macro-averaged one-vs-rest area under the receiver operating characteristic curve, F1 denotes macro-F1, and BACC denotes balanced accuracy, calculated as the unweighted mean of class-wise recall.

As summarized in Table 4, the proposed anatomically informed ROI-level graph model demonstrated consistently strong performance across all diagnostic tasks. In the CN vs. AD binary classification task, the proposed method achieved the highest AUC among all models, indicating superior discriminative capability. Although its F1 score and ACC were slightly lower than those of the ROI-Vector baseline, the differences were marginal, suggesting comparable overall performance.

For the CN vs. MCI and MCI vs. AD tasks, which involve more subtle inter-group distinctions, the proposed model consistently improved AUC while maintaining competitive F1 and ACC. These findings suggest that incorporating anatomically structured inter-regional relationships enhances sensitivity to transitional disease stages.

In the most challenging multi-class task (CN vs. MCI vs. AD), the proposed model achieved the best overall performance across all evaluation metrics. Compared with the ROI-Vector baseline, the proposed approach improved AUC, macro-F1, and BACC simultaneously. The gains were even more pronounced relative to the 3D DenseNet baseline, particularly in macro-F1, indicating improved class-wise balance under the graph-based formulation. Together, these results highlight the advantage of anatomically informed graph representations in modeling heterogeneous structural patterns across diagnostic stages.

Table 4
Classification performance comparison across diagnostic tasks (5-fold cross-validation, mean ± SD). ACC denotes overall accuracy for binary tasks, whereas BACC denotes balanced accuracy for the three-class task.

Task	Method	Evaluation metric		
		AUC	F1	ACC/BACC
CN vs AD	ROI-Vector Baseline	0.919 ± 0.048	**0.840 ± 0.064**	**0.873 ± 0.044**
	3D DenseNet Baseline	0.886 ± 0.051	0.757 ± 0.073	0.788 ± 0.062
	Proposed	**0.928 ± 0.033**	0.839 ± 0.033	0.868 ± 0.027
CN vs MCI	ROI-Vector Baseline	0.755 ± 0.062	0.734 ± 0.052	0.696 ± 0.052
	3D DenseNet Baseline	0.709 ± 0.058	0.611 ± 0.081	0.616 ± 0.058
	Proposed	**0.774 ± 0.062**	**0.736 ± 0.057**	**0.709 ± 0.049**
MCI vs AD	ROI-Vector Baseline	0.673 ± 0.048	0.496 ± 0.035	0.646 ± 0.051
	3D DenseNet Baseline	0.586 ± 0.078	0.462 ± 0.004	0.311 ± 0.025
	Proposed	**0.689 ± 0.079**	**0.501 ± 0.077**	**0.652 ± 0.042**
CN vs MCI vs AD	ROI-Vector Baseline	0.708 ± 0.051	0.535 ± 0.051	0.549 ± 0.052
	3D DenseNet Baseline	0.697 ± 0.039	0.361 ± 0.041	0.531 ± 0.058
	Proposed	**0.745 ± 0.026**	**0.562 ± 0.035**	**0.600 ± 0.044**

To statistically assess model differences in the multi-class setting, a Friedman test was performed (Table 5). The results indicated significant differences among the three models across evaluation metrics. Overall, these findings suggest that organizing ROI-level morphological and textural features within anatomically defined adjacency structures provides a robust and balanced representation for multi-stage AD classification.

Table 5
Friedman Test Results for Multi-Class Classification (CN vs. MCI vs. AD)

Task	Evaluation metric	Statistical test results	
		χ^2 (df = 2)	p-value
CN vs MCI vs AD	AUC	7.20	**0.027**
	Macro-F1	7.60	**0.022**
	BACC	6.40	**0.041**

3.7 External Validation on Independent Dataset

To further evaluate generalizability, we conducted external validation on the independent NACC dataset under a three-class classification setting (CN vs. MCI vs. AD). Each model was trained on ADNI1 using the protocol described in Section 2.11 and directly applied to the NACC cohort without additional fine-tuning, thereby assessing robustness to cross-dataset domain shift.

As shown in Table 6, the proposed ROI-level graph-based model achieved the best overall performance across all evaluation metrics. Compared to the ROI-vector baseline, the proposed method improved AUC by + 0.016, macro-F1 by + 0.088, and BACC by + 0.056. Notably, the relative improvement in macro-F1 (approximately 22% over the ROI-vector baseline) was substantially larger than the gain in AUC, indicating that the proposed framework primarily enhanced class-wise prediction balance rather than merely improving ranking performance. The proposed model also outperformed the 3D DenseNet baseline across all metrics, supporting the advantage of ROI-level relational modeling over voxel-level CNN representations in the external multi-class setting.

Overall, these results demonstrate that the proposed graph-based representation maintains both discriminative performance and class balance under cross-dataset domain shift, enabling more reliable multi-class generalization to an independent cohort.

Table 6
External validation results on the NACC dataset. Values are reported as mean ± SD across the five models trained on the ADNI1 cross-validation folds, each evaluated on the full NACC cohort.

Task	Method	Evaluation metric		
		AUC	Macro-F1	BACC
CN vs MCI vs AD	ROI-Vector Baseline	0.703 ± 0.024	0.392 ± 0.062	0.429 ± 0.059
	3D DenseNet Baseline	0.655 ± 0.045	0.331 ± 0.043	0.420 ± 0.072
	Proposed	**0.719 ± 0.015**	**0.480 ± 0.029**	**0.485 ± 0.030**

Figure 7 presents the class-wise ROC and precision–recall (PR) curves for the three-class classification setting (CN vs. MCI vs. AD) on the independent external dataset, comparing the ROI-Vector baseline, the

3D DenseNet baseline, and the proposed graph-based model. Performance varied across diagnostic categories.

For the CN class (Fig. 7a, 7b), the proposed model demonstrated the strongest discriminative performance in both ROC and PR analyses, achieving an AUC improvement of approximately + 0.03 over the ROI-Vector baseline and showing improved PR balance across the recall range.

For the MCI class (Fig. 7c, 7d), performance differences were more pronounced. The proposed graph-based model improved AUC by approximately + 0.05 compared with the ROI-Vector baseline and showed consistent gains in PR space, indicating enhanced sensitivity to the intermediate disease stage within the multi-class framework.

In contrast, for the AD class (Fig. 7e, 7f), the ROI-Vector baseline achieved the highest discriminative performance. The proposed model showed a modest decrease in AUC (approximately − 0.03 relative to ROI-Vector) but remained competitive and continued to outperform the 3D DenseNet baseline.

Overall, these findings indicate that model performance varied across diagnostic categories. The proposed graph-based framework provided consistent improvements for CN and MCI classification, whereas the ROI-Vector baseline achieved the strongest performance in AD detection.

4. Discussion

In this study, we introduced an anatomically informed ROI-level graph representation for structural MRI analysis of AD. Rather than treating ROIs as isolated summary vectors or relying solely on voxel-level deep networks, our framework represents each region as a morpho-textural node embedded within a geometry-aware anatomical graph. The principal contribution of this work lies in the representation itself: regional MRI-derived attributes are organized together with quantitative descriptors of the physical interfaces shared by adjacent structures, enabling explicit characterization of inter-regional structural relationships while preserving anatomical interpretability.

The proposed representation integrates enriched regional descriptors with geometry-aware interface information. Each ROI was characterized using morphological features—volume, thickness, and fractal dimension—and textural descriptors—intensity, gradient, and entropy—thereby capturing multiple aspects of regional structural variation. Importantly, the graph formulation extends beyond binary adjacency by incorporating graded interface attributes, including Structural Contact and Boundary Appearance. These descriptors quantify properties of the anatomical interface itself rather than only the characteristics of the individual regions it connects.

An important question is whether these interface-level descriptors provide information that is distinct from conventional regional measurements. The feature complementarity analysis indicated that node- and edge-derived descriptors occupy partially distinct representational spaces. Although Structural Contact showed a comparatively stronger association with thickness than with other node features,

correlations between interface attributes and most regional descriptors remained limited, and Structural Contact and Boundary Appearance themselves showed only weak mutual correlation. These findings support the view that interface-level descriptors capture aspects of anatomical organization that are not fully represented by conventional ROI-wise morphology or texture.

The interface-level patterns observed in this study further suggest that AD-related structural variation is not confined to changes occurring within individual regions. In contrast to node-derived effect maps, edge-based analyses revealed spatially organized alterations localized to anatomical contact interfaces between adjacent structures. In particular, the Structural Contact network showed coherent patterns involving medial temporal, ventricular, and white matter structures, consistent with established anatomical changes across the AD continuum. These observations suggest that regional tissue loss may be accompanied by systematic modification of the geometry through which neighboring structures contact one another. Structural Contact should therefore not be interpreted simply as another measure of regional size, but as a geometric descriptor that may capture aspects of coordinated tissue remodeling expressed at the interface between adjacent anatomical regions.

A further question is whether these interface alterations merely arise as secondary consequences of regional morphological or textural abnormalities. Node-controlled edge analyses demonstrated that several interface-level associations with disease stage remained significant after adjustment for the six morpho-textural characteristics of both incident ROIs. Although such adjustment cannot establish complete independence from all aspects of regional anatomy, it indicates that the observed interface effects are not fully reducible to the measured intra-regional characteristics alone. The corresponding effect maps further showed that the retained associations were spatially localized rather than uniformly distributed across anatomical boundaries, supporting a selective rather than diffuse pattern of interface-level structural alteration.

The clinical association analysis provides an additional line of evidence for the relevance of interface-level information. A subset of Structural Contact features showed convergent associations with CDR-SB, ADAS-Cog13, and MMSE after adjustment for age, sex, diagnosis, and the morpho-textural characteristics of both connected ROIs. The direction of these associations was consistent across the three clinical measures, suggesting that the identified interfaces were related to a broader dimension of clinical severity rather than to variation specific to a single cognitive or functional scale. These findings extend the disease-stage analysis by showing that selected anatomical interfaces also carry clinically meaningful information beyond diagnostic category and measured regional morphology.

Among these interfaces, reduced Structural Contact between the fusiform and left parahippocampal regions was particularly notable. This finding is anatomically and biologically plausible given the vulnerability of the ventral medial temporal system in AD. Recent neuropathological evidence has demonstrated inter-regional associations between these regions, with axonal damage in the fusiform gyrus associated with synaptic loss in the parahippocampal gyrus, and vice versa [40]. In the present study, lower Fusiform–Parahippocampal Structural Contact was consistently associated with greater

impairment across CDR-SB, ADAS-Cog13, and MMSE. Importantly, this association remained significant after controlling for the morpho-textural characteristics of both connected ROIs. This observation is consistent with previous structural imaging evidence that inter-regional alterations in AD may contain information that is not fully explained by regional atrophy alone [41, 42]. Accordingly, reduced contact at this interface may reflect coordinated geometric alteration across adjacent vulnerable temporal regions rather than isolated changes within either region alone. This interpretation extends conventional region-based observations by suggesting that AD-related structural abnormalities may also be expressed through alterations in the anatomical relationship between neighboring structures.

White matter hypointensities (ROI 33) also repeatedly emerged in the node-controlled interface analyses. White matter hypointensities appear as low-signal abnormalities on T1-weighted MRI and have been associated with chronic small vessel disease, demyelination, and ischemic injury [43, 44]. Although they are not a defining pathological hallmark of classical AD, they are strongly associated with aging and vascular pathology and have been linked to cognitive decline [45–47]. In the present analysis, multiple interfaces involving ROI 33 retained significant disease-stage associations after adjustment for node-level morpho-textural features and were observed in both Structural Contact and Boundary Appearance analyses. This pattern suggests that the findings cannot be attributed solely to the local size or signal characteristics of the white matter hypointensity region itself.

The corresponding node-controlled effect maps further showed that these associations were spatially localized rather than diffusely distributed throughout white matter. Retained effects appeared as small, discontinuous clusters concentrated near periventricular and deep white matter–adjacent interfaces. Such spatial restriction is more consistent with localized interface alteration than with generalized white matter atrophy or global boundary degradation. This interpretation is compatible with recent evidence that AD-related pathological processes may interact with structurally connected and vascularly vulnerable white matter regions. Cortical tau deposition, for example, has been associated with degeneration along connected white matter pathways [48], while periventricular and deep white matter regions are recognized as vulnerable to cerebral small vessel disease and blood–brain barrier dysfunction [49]. White matter lesion burden has also been associated with Alzheimer's pathology and cognitive decline [50], and vascular and clearance-related mechanisms have been proposed to contribute to spatially heterogeneous network abnormalities [51, 52]. Taken together, these observations provide biological context for the localized interface alterations involving white matter abnormalities observed in the present study, although the underlying mechanisms cannot be determined from the current structural MRI analysis.

From a clinical perspective, the present findings suggest that geometry-aware interface attributes may complement conventional regional volumetric assessments rather than replace them. Regional volumes remain strong markers of AD-related structural abnormality, particularly within established medial temporal regions. However, the present analyses showed that interface features can provide additional information in anatomical relationships for which regional size alone is comparatively less informative. By characterizing how adjacent structures are geometrically arranged with respect to one another,

Structural Contact and Boundary Appearance provide a different level of description from conventional ROI-wise measures. The observed interface alterations may therefore occur alongside macroscopic atrophy while capturing aspects of anatomical organization that are not fully represented by regional size. This complementary role is consistent with the central premise of the proposed representation: AD-related structural information can be expressed both within individual regions and at the physical interfaces connecting them.

The representation ablation analyses provided further support for this interpretation. Disrupting anatomically structured topology or removing interface attributes consistently reduced classification performance. In particular, structure-free graph configurations converged on similar performance despite differing in how relational structure was disrupted, whereas restoring anatomical adjacency improved class-wise balance and incorporating quantitative interface attributes produced an additional gain in discriminative performance. Moreover, the performance advantage of the full representation could not be explained simply by increased feature availability or model capacity, because reorganizing the same information without anatomically meaningful relational structure also degraded performance. Together, these findings indicate that the organization of features within anatomically faithful relationships contributes to classification beyond the regional feature content alone.

Consistent with the ablation findings, the proposed representation demonstrated its strongest and most consistent predictive advantage in the multi-class setting (CN vs. MCI vs. AD), both in internal cross-validation and external validation. Although the overall gains in AUC were moderate, improvements in macro-F1 together with higher balanced accuracy indicated better class-wise balance across diagnostic groups. Class-wise ROC and PR analyses further revealed a nuanced pattern: the graph-based model showed stronger discrimination for CN and MCI, whereas the ROI-vector baseline performed slightly better for AD. This pattern suggests that relational information may be particularly useful for distinguishing early and intermediate diagnostic stages, whereas advanced AD may already be characterized by sufficiently pronounced regional abnormalities to be captured by conventional ROI-wise summaries. Importantly, the preservation of the overall advantage on the independent NACC cohort without additional fine-tuning suggests that the representation is not limited to the acquisition characteristics of the development dataset.

Several limitations should be acknowledged. First, representing each ROI as a single node assumes relative intra-regional homogeneity, and highly focal pathological changes may therefore be attenuated through regional aggregation. Increasing graph granularity through subregional partitioning or surface-based sampling may improve sensitivity while preserving anatomical coherence. Second, the current graph topology is restricted to local anatomical adjacency and does not incorporate long-range structural or functional connectivity. Future integration of multimodal connectivity information may enable characterization of network-level abnormalities extending beyond physically adjacent structures [53]. Third, although node-controlled analyses adjusted for multiple regional morpho-textural characteristics, residual dependence on unmeasured aspects of regional geometry cannot be completely excluded. Finally, the present analysis is cross-sectional. Longitudinal studies will be required

to determine whether interface alterations change over time, whether they precede or follow regional atrophy, and whether they provide prognostic information regarding subsequent clinical decline or treatment response.

The anatomically grounded ROI graph paradigm may also be applicable to neurological disorders beyond AD in which structural abnormalities follow spatially organized anatomical patterns. More fundamentally, the present study demonstrates that the organization of neuroimaging features may be as important as the predictive architecture applied to them. By explicitly encoding physical anatomical interfaces and geometry-constrained relationships, the proposed framework extends structural MRI representation from isolated regional summaries toward a joint description of regional and inter-regional anatomy. The convergent evidence from feature complementarity, node-controlled analyses, clinical associations, representation ablation, and external validation supports the view that anatomical interfaces constitute a measurable and complementary substrate of AD-related structural variation rather than merely an architectural component of a graph classifier.

5. Conclusion

Rather than treating ROIs independently or relying solely on voxel-level representations, we proposed an anatomically informed ROI-level graph framework that jointly models regional morpho-textural properties and the physical interfaces between adjacent brain regions in structural T1-weighted MRI. By incorporating geometry-aware interface attributes, the framework extends conventional region-wise representation toward an anatomically structured description of both intra-regional and inter-regional information.

In multi-stage classification and independent external validation, the proposed representation demonstrated improved performance, particularly in class-wise balance. Beyond predictive performance, node-controlled analyses showed that selected interface-level features remained significantly associated with disease stage after accounting for the morpho-textural characteristics of both connected ROIs, and selected Structural Contact features were also consistently associated with clinical severity across multiple cognitive and functional measures.

Collectively, these findings indicate that anatomical interfaces provide measurable and complementary structural information that is not fully captured by conventional regional summaries. The proposed framework therefore offers an interpretable and extensible representation for characterizing AD-related structural variation at both regional and inter-regional levels.

Declarations

Competing Interests

The authors have no conflicts of interest to disclose.

Ethics Approval

This study used only de-identified data from the Alzheimer's Disease Neuroimaging Initiative (ADNI) and the National Alzheimer's Coordinating Center (NACC). Data collection procedures for the original studies were approved by the relevant institutional review boards. The present study involved secondary analysis of de-identified data and no direct interaction with participants; therefore, no additional ethical approval was required.

Consent to Participate

Written informed consent was obtained from all participants as part of the original ADNI and NACC studies.

Consent to Publish

Not applicable. The present study used de-identified data and contains no individually identifiable participant information.

Funding

Declaration

This work was supported by the Korea Medical Device Development Fund grants funded by the Korea government (the Ministry of Science and ICT, the Ministry of Trade, Industry and Energy, the Ministry of Health & Welfare, and the Ministry of Food and Drug Safety) (Project Number: 2410018859, RS-2026-25544586; Project Number: 2410018850, RS-2026-25545662).

Author Contribution

Yongha Gi contributed to the study conception and design, methodology development, software implementation, data analysis, visualization, and preparation of the original manuscript. Hyeongjin Lim, Saemin Park, Jeongwon Lee, Sangyoon Park, A Hyun Jung, Seol-Hee Baek, and Jong Hyun Kim contributed to data analysis, interpretation of the results, and critical review of the manuscript. Byung-Jo Kim contributed to clinical interpretation, study supervision, and critical revision of the manuscript. Myonggeun Yoon contributed to the study conception and design, supervision, project administration, funding acquisition, and critical revision of the manuscript. All authors read and approved the final manuscript.

Acknowledgement

Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer’s Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

Data Availability

The data used in this study are available from the Alzheimer's Disease Neuroimaging Initiative (ADNI) and the National Alzheimer's Coordinating Center (NACC), subject to their respective data access policies.

Code Availability

The code used in the current study is available from the corresponding author upon reasonable request.

References

1. Porsteinsson AP, et al. Diagnosis of early Alzheimer's disease: clinical practice in 2021. J Prev Alzheimer's disease. 2021;8(3):371–86.
2. Zhang J, et al. Recent advances in Alzheimer’s disease: Mechanisms, clinical trials and new drug development strategies. Signal Transduct Target therapy. 2024;9(1):211.
3. Miao D, et al. Hippocampal morphological atrophy and distinct patterns of structural covariance network in Alzheimer’s disease and mild cognitive impairment. Front Psychol. 2022;13:980954.

4. Abushakra S et al. Hippocampal Atrophy on Magnetic Resonance Imaging as a Surrogate Marker for Clinical Benefit and Neurodegeneration in Early Symptomatic Alzheimer's Disease: Synthesis of Evidence from Observational and Interventional Trials: S. Abushakra. CNS Drugs (2025): 1–16.
5. Katabathula S, Wang Q, Xu R. Predict Alzheimer's disease using hippocampus MRI data: a lightweight 3D deep convolutional network model with visual and global shape representations. Alzheimers Res Ther. 2021;13(1):104.
6. Arabi EM, Ashraf S, Mohra, Khaled S. Ahmed. High accuracy diagnosis for MRI imaging of Alzheimer's disease using XGBoost. Open Biotechnol J 16.1 (2022).
7. Zheng X, and Alzheimer's Disease Neuroimaging Initiative. Detection of Alzheimer's disease using hybrid meta-ROI of MRI structural images. *Diagnostics* 14.19 (2024): 2203.
8. Tong Y, et al. Seeking optimal region-of-interest (ROI) single-value summary measures for fMRI studies in imaging genetics. PLoS ONE. 2016;11(3):e0151391.
9. Davatzikos C. Why voxel-based morphometric analysis should be used with great caution when characterizing group differences. *Neuroimage* 23.1 (2004): 17–20.
10. Mousavi S, Mohammad K, Moulaei, Leila Ahmadian. Classifying and diagnosing Alzheimer's disease with deep learning using 6735 brain MRI images. Sci Rep. 2025;15(1):22721.
11. Reddy KK. AlzheimerViT: harnessing lightweight vision transformer architecture for proactive Alzheimer's screening. Front Med. 2025;12:1568312.
12. Liu, Wei J, Zhang, Zhao Y. A comparison of deep learning and traditional machine learning approaches in detecting cognitive impairment using MRI scans. 2022 IEEE 46th Annual Computers, Software, and, Conference A. (COMPSAC). IEEE, 2022.
13. Vieira S, Pinaya WHL, Mechelli A. Using deep learning to investigate the neuroimaging correlates of psychiatric and neurological disorders: Methods and applications. Neurosci Biobehavioral Reviews. 2017;74:58–75.
14. Qu Z, et al. A graph convolutional network based on univariate neurodegeneration biomarker for Alzheimer's disease diagnosis. IEEE J Translational Eng Health Med. 2023;11:405–16.
15. Ali S, et al. Graph neural networks in Alzheimer's disease diagnosis: a review of unimodal and multimodal advances. Front NeuroSci. 2025;19:1623141.
16. Zeng X, et al. Knowledge-driven multi-graph convolutional network for brain network analysis and potential biomarker discovery. Med Image Anal. 2025;99:103368.
17. Luo X et al. Graph neural networks for brain graph learning: A survey. *arXiv preprint arXiv:2406.02594* (2024).
18. Mohammadi H, Karwowski W. Graph neural networks in brain connectivity studies: Methods, challenges, and future directions. Brain Sci. 2024;15(1):17.
19. Said A, et al. Neurograph: Benchmarks for graph machine learning in brain connectomics. Adv Neural Inf Process Syst. 2023;36:6509–31.

20. Li X, et al. Braingnn: Interpretable brain graph neural network for fmri analysis. Med Image Anal. 2021;74:102233.
21. Mueller SG, et al. The Alzheimer's disease neuroimaging initiative. Neuroimaging Clin. 2005;15(4):869–77.
22. Beekly DL, et al. The National Alzheimer's Coordinating Center (NACC) database: the uniform data set. Alzheimer Disease Assoc Disorders. 2007;21(3):249–58.
23. Tustison NJ, et al. N4ITK: improved N3 bias correction. IEEE Trans Med Imaging. 2010;29(6):1310–20.
24. Isensee F, et al. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nat Methods. 2021;18(2):203–11.
25. Simonovsky M, Komodakis N. Dynamic edge-conditioned filters in convolutional neural networks on graphs. *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017.
26. LaMontagne PJ et al. OASIS-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer disease. *medrxiv* (2019): 2019-12.
27. Henschel L et al. Fastsurfer-a fast and accurate deep learning based neuroimaging pipeline. *NeuroImage* 219 (2020): 117012.
28. Zhang J et al. Automatic analysis of skull thickness, scalp-to-cortex distance and association with age and sex in cognitively normal elderly. *bioRxiv* (2023).
29. Fischl B. and Anders M. Dale. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proceedings of the National Academy of Sciences* 97.20 (2000): 11050–11055.
30. King RD et al. Fractal dimension analysis of the cortical ribbon in mild Alzheimer's disease. *Neuroimage* 53.2 (2010): 471–479.
31. Lim H et al. Gradient magnitude-based texture feature of the entorhinal cortex as a biomarker for Alzheimer’s disease. Eur J Radiol (2026): 112746.
32. Veluppal A. Differentiation of Alzheimer conditions in brain MR images using bidimensional multiscale entropy-based texture analysis of lateral ventricles. Biomed Signal Process Control. 2022;78:103974.
33. Felzenszwalb PF, Daniel P. Huttenlocher. Efficient graph-based image segmentation. Int J Comput Vision. 2004;59(2):167–81.
34. Zhang Y, Brady M. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. IEEE Trans Med Imaging. 2002;20(1):45–57.
35. Gonzalez RC. Digital image processing. Pearson education india; 2009.
36. Zalesky A, Fornito A, Edward T. Bullmore. Network-based statistic: identifying differences in brain networks. *Neuroimage* 53.4 (2010): 1197–1207.
37. Richardson JTE. Eta squared and partial eta squared as measures of effect size in educational research. Educational Res Rev. 2011;6(2):135–47.
38. Goodfellow I, et al. Deep learning. Vol. 1. No. 2. Cambridge: MIT Press; 2016.

39. Huang G et al. Densely connected convolutional networks. *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017.
40. Bouwman, Maud MA, et al. Distinct regional patterns of synaptic vulnerability across hippocampal and parahippocampal subregions in Alzheimer's disease. Brain Pathol. 2026;36(4):e70081.
41. Andreotti J, et al. Diverging progression of network disruption and atrophy in Alzheimer's disease and semantic dementia. J Alzheimer's Disease. 2016;55(3):981–93.
42. Pajavand A, Mohammad, et al. Structural white matter connectivity differences independent of gray matter loss in mild cognitive impairment with neuropsychiatric symptoms: Early indicators of Alzheimer's disease using network-based statistics. J Alzheimer's Disease. 2024;102(4):1042–56.
43. Kang P, et al. MRI signatures associated with active ischemia and disease severity in cerebral small vessel disease. Neuroimage: Rep. 2025;5(3):100281.
44. Sarbu N et al. White matter diseases with radiologic-pathologic correlation. *Radiographics* 36.5 (2016): 1426–1447.
45. Li T-R, et al. Association of white matter hyperintensities with cognitive decline and neurodegeneration. Front Aging Neurosci. 2024;16:1412735.
46. Markus HS, Frank Erik de Leeuw. Cerebral small vessel disease: recent advances and future directions. Int J Stroke. 2023;18(1):4–14.
47. d'Arbeloff T, et al. White matter hyperintensities are common in midlife and already associated with cognitive decline. Brain Commun. 2019;1(1):fcz041.
48. Pescoller J et al. Cortical tau deposition promotes atrophy in connected white matter regions in Alzheimer's disease. *Brain* 148.12 (2025): 4359–4371.
49. Gu D, et al. Phenotypes of white matter hyperintensities, mechanisms and aetiological differentiation: future targets pre-dementia—a systematic review. Discover Neurosci. 2025;20(1):6.
50. Cha W-J, et al. Association between brain amyloid deposition and longitudinal changes of white matter hyperintensities. Alzheimers Res Ther. 2024;16(1):50.
51. Tian J et al. White matter degeneration pathways associated with tau deposition in Alzheimer disease. *Neurology* 100.22 (2023): e2269-e2278.
52. Tora V, et al. A network-level transport model of tau progression in the Alzheimer's brain. Math Med Biology: J IMA. 2025;42(2):212–38.
53. Pascucci D, et al. Structure supports function: Informing directed and dynamic functional connectivity with anatomical priors. Netw Neurosci. 2022;6(2):401–19.