Interpretable Structure-Function Coupling via State-Space Models for Brain Connectome Analysis

Amashi Niwarthana[1 [0009−0007−5124−3239]], Hamka Bin Dzulkarnain[1], Xia Jing[2], and Jagath C. Rajapakse[1 [0000−0001−7944−1658]⋆]

[1] Nanyang Technological University, Singapore
{amashini001,hamka001}@e.ntu.edu.sg, asjagath@ntu.edu.sg
[2] Zhejiang University, Zhejiang, China jing_xia@zju.edu.cn

Abstract. Brain structural connectivity (SC) provides the anatomical foundation for functional connectivity (FC), yet existing methods model structure-function coupling as one-to-one correspondence between paired regions. This contradicts neuroscientific evidence showing that functional interactions emerge from distributed structural pathways spanning multiple brain areas. We propose a many-to-many coupling approach using graph state-space models, which encode SC and FC graphs separately, and thereafter sequentially scan their interactions where each position accumulates information from all previous positions. Unlike graph convolutional networks that suffer from over-smoothing or attention mechanisms that compute region-pair interactions independently, we show that our approach captures long-range distributed coupling through state accumulation. We further introduce a hierarchical variant that incorporates modular brain organization, performing within-module fusion before cross-module integration for biologically constrained many-to-many learning. Quantitative analysis using interaction sparsity, distance, and modularity metrics validates the distributed coupling capacity. Experiments across four datasets (HCP, PPMI, ADNI, COBRE) show consistent improvements: r=0.32 for fluid cognition and upto 97.0% accuracy on disease classification. Our method activates ≈49% of possible SC–FC region pairs versus ≈20% for attention, with mean hop distances of $\bar{d}$=0.98 and cross-module ratios exceeding 85%. Interpretability analysis reveals disease-specific patterns including subcortical-motor disruption in Parkinson's disease.

Keywords: Brain connectome · Disease classification · Graph state-space models · Mamba · Multimodal fusion · Structure-function coupling.

1 Introduction

Brain structural connectivity (SC) representing white matter pathway integrity estimated from diffusion tensor imaging (DTI), and functional connectivity (FC)

⋆ Corresponding author.

capturing temporal dependencies between BOLD signals via resting-state fMRI (rs-fMRI) are closely coupled: FC tends to be strongest among structurally connected regions and alterations in their relationship are associated with cognitive change and brain disorders [2, 8, 5, 27, 18]. Yet most functional connections form without a direct structural link [19], and high-order polysynaptic interactions give rise to complex structure-function couplings that go beyond simple region-to-region correspondence [8, 19].

Graph neural networks (GNN) have become the dominant framework for connectome analysis [3, 6], with recent multimodal extensions modelling SC–FC interactions through learned coupling weights or modularity-guided convolution [25, 4, 24]. However, these methods enforce a one-to-one correspondence: structural region i interacts only with functional region i. This contradicts the neuroscientific evidence above and fails to capture distributed polysynaptic influences [19, 8].

We address this by proposing *many-to-many* SC–FC coupling where every structural region can influence all functional regions. We implement this through state-space model, specifically Mamba [9] that process tokens sequentially so that each position accumulates information from all preceding positions via a hidden state. Constructing an interleaved SC–FC sequence allows structural region i to propagate naturally to all subsequent functional regions $j{>}i$ without enumerating $O(N^2)$ pairs. Cross-attention provides all-to-all pairing but treats each pair independently; element-wise operations (e.g., concatenation) enforce one-to-one correspondence by design.While Mamba has been applied to multimodal fusion via coupled parallel state chains [13] and to brain image fusion through interleaved token sequences [15], and has shown promise for unimodal connectome analysis [23], its application to multimodal SC–FC fusion remains unexplored.

This paper outlines the following contributions: (i) We propose an state-space many-to-many fusion (SSFuse) framework using interleaved SC–FC sequences for distributed structure-function coupling; (ii) We introduce a hierarchical variant (GM-SSFuse-H) that constrains within-module fusion before cross-module integration, consistent with the modular organisation of the brain [26]; (iii) We define three quantitative interaction metrics; active pair fraction, mean hop distance, and cross-module ratio, to validate distributed coupling empirically; and (iv) Our methods evaluated on four datasets (HCP, PPMI, ADNI, COBRE) outperform six baselines, with gradient-based analysis recovering neuroanatomically interpretable SC–FC patterns consistent with known disease mechanisms [27, 18].

2 Materials and Methodology

2.1 Datasets and Image Preprocessing

For fluid cognition prediction, we used 838 participants from the Human Connectome Project (HCP) [21]. For disease classification, we used three datasets: PPMI [14] (75 controls, 72 Parkinson's disease (PD) patients), ADNI [10] (105 controls, 100 Alzheimer's disease (AD) patients), and COBRE [1] (72 controls,

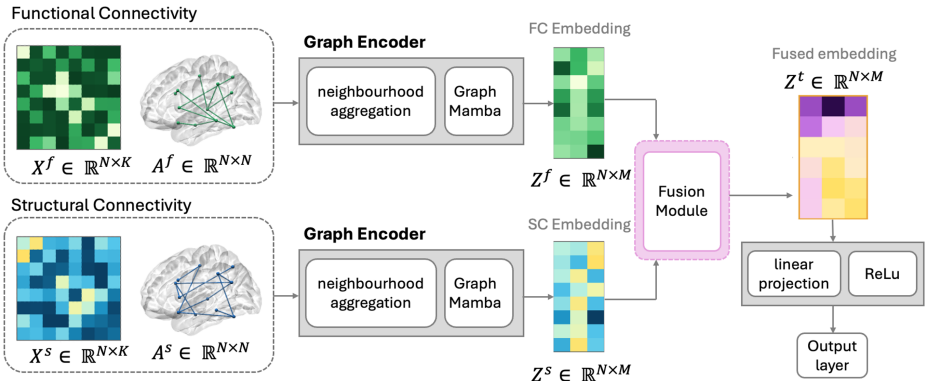

Fig. 1. Overview of proposed GM-SSFuse. FC and SC graphs are encoded independently via modality-specific Graph Mamba (GM), and then fused through a state-space many-to-many fusion (SSFuse) module. The fused embedding is passed to a readout and output layer for regression (fluid cognition) or classification (neurological diseases).

60 schizophrenia (SZ) patients). All datasets include T1w MRI, rs-fMRI, and DTI images. For HCP, rs-fMRI and DTI images were preprocessed with the DPARSF 5.1 toolkit and the brain diffusion toolkit within FSL [11], respectively. For PPMI, ADNI, and COBRE, DTI images were preprocessed via Clinica [17], including eddy current correction, with structural connectivity matrices obtained as streamline counts between ROI pairs via probabilistic tractography using FSL following the AAL atlas. fMRI images were preprocessed with fMRIPrep [7], including motion correction, slice timing correction, spatial normalisation to MNI152NLin2009cAsym space, and nuisance regression. FC and SC were constructed with 116 ROIs following the AAL protocol [20], covering both cortical and subcortical structures. Fluid cognition scores were extracted from CogFluidComp in the HCP phenotype list, ranging from 87 to 147. FC between paired ROIs was computed using the Pearson correlation coefficient. A modified Yeo 7-network parcellation [26] was adopted by introducing an additional subcortical module, yielding $P = 8$ modules.

2.2 Proposed Method

As shown in Fig. 1, the proposed framework (GM-SSFuse) consists of (1) modality-specific Graph Mamba (GM) encoders and (2) a State-Space Fusion (SSFuse) module. Further, we propose a hierarchical variant (GM-SSFuse-H) incorporating modular brain organization, and introduce interaction analysis metrics.

Graph Representation The brain connectome is modeled as a graph with ROIs as nodes and connectivity strengths as edges. For each participant, a FC graph $G^f = (V^f, A^f, X^f)$ is constructed from rs-fMRI, where $A^f \in \mathbb{R}^{N\times N}$ is a thresholded FC matrix preserving the top 10% positive edges per node [24], and

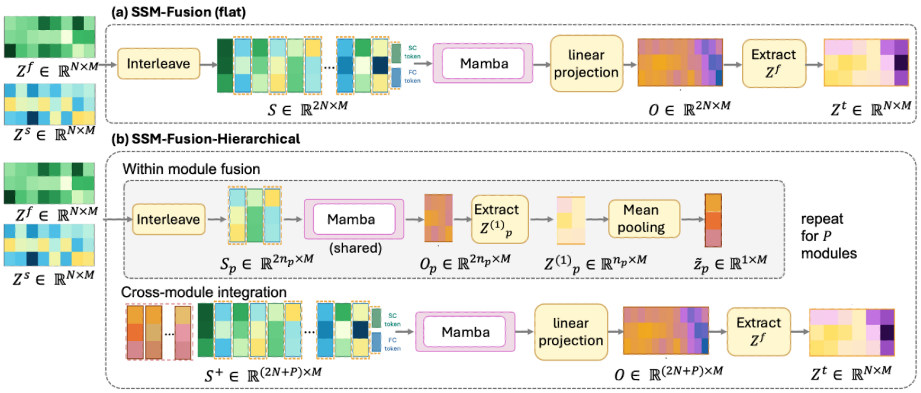

Fig. 2. Detailed view of SSFuse variants. *(a)* SSFuse: SC and FC node embeddings are interleaved into a single sequence and processed by a bidirectional Mamba scan, enabling many-to-many SC-FC coupling across all regions. *(b)* Hierarchical fusion (SSFuse-H): within-module Mamba scans are applied independently per brain network, constraining local SC-FC coupling to biologically coherent groups; mean-pooled module tokens are then integrated by a cross-module Mamba scan.

$X^f \in \mathbb{R}^{N\times K}$ contains per-node FC profiles. Similarly, a structural connectivity graph $G^s = (V^s, A^s, X^s)$ is constructed from DTI, where $A^s \in \mathbb{R}^{N\times N}$ is derived from fiber tractography and $X^s \in \mathbb{R}^{N\times K}$ contains per-node SC profiles. V^f and V^s share the same N ROIs.

Graph Mamba (GM) Encoder Each graph is encoded independently using GM [22]. The core idea is to enrich each node's representation with local structural context before processing the node sequence with a Mamba selective state-space model [9], so that the hidden state propagates both topology-aware and task-relevant information along the scan.

Graph-augmented input: Given node features $X \in \mathbb{R}^{N\times K}$ and adjacency matrix A, a linear projection first produces token embeddings $H = XW_{\text{in}} \in \mathbb{R}^{N\times M}$. For each node i, a one-hop neighborhood aggregate $\tilde{h}_i = \sum_j A_{ij} h_j$ is computed, and the node input to the Mamba scan is formed by concatenation: $u_i = [h_i \| \tilde{h}_i] \in \mathbb{R}^{2M}$. This augmentation injects local structural context into every position of the sequence before the SS model begins propagating information globally. The Mamba block then processes $U \in \mathbb{R}^{N\times 2M}$ sequentially: at each position t,

$$h_t = \bar{A}_t h_{t-1} + \bar{B}_t u_t, \qquad o_t = C_t h_t, \tag{1}$$

where $\bar{A}_t$, $\bar{B}_t$, C_t are input-dependent selection matrices, with $\bar{A}_t = \exp(\Delta_t A)$ and $\bar{B}_t = \Delta_t B_t$ as zero-order-hold discretizations. A linear projection maps the output back to $\mathbb{R}^M$.

Topology-aware node ordering: The Mamba hidden state at position t accumulates information from all preceding positions, so the order in which

nodes are presented to the scan determines which neighborhood relationships are captured implicitly through state propagation. We traverse the graphs from the highest-degree node and evaluate four orderings: *degree* (hub nodes first, maximizing early-state coverage), breadth-first search *BFS* (topologically proximal nodes consecutive, encouraging local neighborhood coherence in the hidden state), depth-first search *DFS* (long structural paths preserved sequentially), and *spectral* (nodes sorted by the Fiedler vector of the normalized Laplacian $L = I - D^{-\frac{1}{2}}AD^{-\frac{1}{2}}$, placing same-community nodes adjacently and aligning the scan with the structure of the connectome). After encoding, the output is inverse-permuted to restore the original node ordering. We find BFS optimal for FC and degree ordering for SC, and use these throughout.

Modality-specific encoders: The above procedure is applied independently to each modality, yielding $Z^f \in \mathbb{R}^{N\times M}$ and $Z^s \in \mathbb{R}^{N\times M}$:

$$Z^f = \mathrm{GM}(A^f, X^f), \qquad Z^s = \mathrm{GM}(A^s, X^s). \tag{2}$$

State-space many-to-many Fusion (SSFuse) To overcome one-to-one coupling of corresponding brain regions, we reframe fusion of FC and SC as a sequential state accumulation problem. Given Z^s and Z^f, we construct an interleaved sequence of length $2N$ by alternating structural and functional embeddings:

$$S = \left[z_1^s,\, z_1^f,\, z_2^s,\, z_2^f,\, \ldots,\, z_N^s,\, z_N^f\right] \in \mathbb{R}^{2N\times M}. \tag{3}$$

When the SS model processes z_i^f, its hidden state already encodes all preceding structural nodes $z_1^s, \ldots, z_i^s$ through sequential state updates. Structural information at position i propagates forward to influence all subsequent functional positions $i{+}1, \ldots, N$ through the hidden state, realizing genuine many-to-many coupling without explicitly enumerating $O(N^2)$ region pairs. A Mamba block processes S via Eq. (1), producing $O \in \mathbb{R}^{2N\times M}$. Fused embeddings are extracted at functional positions as $Z_t = O[1{::}2] \in \mathbb{R}^{N\times M}$. The final unified embedding Z_t is then passed through a linear output projection.

Hierarchical Fusion Variant (GM-SSFuse-H) Flat fusion over all N nodes does not exploit the modular organization of the brain and may conflate signals across functionally dissimilar regions. We introduce a GM-SSFuse-H that operates at two levels.

Within-module fusion: for each module $\mathcal{M}_p$ of size n_p, a local interleaved sequence $S_p = [z_i^s,\, z_i^f]_{i\in\mathcal{M}_p} \in \mathbb{R}^{2n_p\times M}$ is processed by a shared within-module Mamba block, producing local fused embeddings $Z_p^{(1)} \in \mathbb{R}^{n_p\times M}$.

Cross-module integration: per-module embeddings are mean-pooled into module-level tokens $\tilde{z}_p = \frac{1}{n_p}\sum_{i\in\mathcal{M}_p} Z_p^{(1)}[i]$, forming $\tilde{Z} \in \mathbb{R}^{P\times M}$, which is prepended to the flat interleaved sequence to form $S^+ = [\tilde{z}_1, \ldots, \tilde{z}_P,\, z_1^s, z_1^f, \ldots, z_N^s, z_N^f] \in \mathbb{R}^{(2N+P)\times M}$. A forward Mamba scan over S^+ primes the hidden state with all module-level context before processing any node pair; outputs at the prepended positions are discarded. The final node-level representation is obtained via a learned linear projection enabling inter-network interactions.

Output Layer and Loss The fused embedding Z_t is processed by a readout layer using the flattening method [24], followed by an output layer producing O:

$$O = \mathrm{Lin}_2\big(\mathrm{Dropout}\big(\mathrm{ReLU}\big(\mathrm{Lin}_1\big(\mathrm{Readout}(Z_t)\big)\big)\big)\big) \tag{4}$$

For classification, a softmax follows Lin_2. The loss function $\mathcal{L}$ is RMSE between predicted and true fluid cognition scores for the regression task, and cross-entropy between predicted and true labels for all disease classification tasks.

Interaction Analysis Metrics To validate that SSM fusion captures distributed many-to-many coupling, we introduce three complementary metrics derived from input-gradient attribution of learned coupling weights w_{ij} between structural region i and functional region j. (1) Interaction sparsity $\rho = \frac{1}{N^2}\sum_{i,j} 1[|w_{ij}| > \epsilon]$ measures the fraction of active region pairs, with higher values indicating broader distributed coupling. (2) Interaction distance $\bar{d} = \frac{1}{|\mathcal{A}|}\sum_{(i,j)\in\mathcal{A}} d_{ij}$ is the mean graph hop distance over active pairs $\mathcal{A}$, measuring how far structural influences propagate. (3) Modularity index $Q = |\{(i,j) \in \mathcal{A} : m(i) \neq m(j)\}|/|\mathcal{A}|$ reflects the proportion of cross-module interactions, assessing whether the model captures inter-network dependencies consistent with known brain organisation.

Coupling weights are estimated via gradient attribution. For each subject, we compute the gradient of the model output with respect to the fused structural and functional embeddings Z^s and Z^f, and define the coupling weight between structural region i and functional region j as $w_{ij} = \|\nabla_{z_i^s}\mathcal{L}\|_2 \cdot \|\nabla_{z_j^f}\mathcal{L}\|_2$. Weights are averaged across subjects, normalised to $[0,1]$ before computing ρ, $\bar{d}$, and Q.

3 Experiments and Results

The framework was implemented in PyTorch 2.0.1 with the `mamba_ssm` library (cuda 11.7) on an NVIDIA A100 GPU. The Adam optimizer was used for training with learning rates of 2×10^{-3} for regression and 3–5×10^{-4} for classification, with 20 and 30 training epochs, respectively. We conducted 5-fold cross-validation over 10 runs and report mean and standard deviation. Overfitting was mitigated via dropout (rate 0.5) after the readout layer and between fully connected layers, and ℓ_2 regularization (10^{-3}) with early stopping.

Comparison with the state of the art: Table 1 compares GM-SSFuse and GM-SSFuse-H against six baselines spanning CNN-based (BrainNetCNN [12]), GCN-based (M-GCN, MV-GCN [6, 28]), transformer-based (GraphGPS [16]), and multi-modal fusion approaches (Cross-GNN, MS-Inter-GCN [25, 24]). On fluid cognition, GM-SSFuse-H achieves the highest Pearson r of .32±.02 and lowest RMSE of 10.97±.69, outperforming the strongest baseline MS-Inter-GCN by Δr=.03 and ΔRMSE=0.81. On disease classification, both variants consistently outperform all baselines. GM-SSFuse-H leads on PPMI (.97±.03) and ADNI (.91±.06), while GM-SSFuse leads on COBRE (.92±.05), suggesting that flat

Table 1. Comparison of performances with existing methods on fluid cognition prediction and disease classification: Pearson correlation (r), RMSE, accuracy, and sensitivity.

	HCP		PPMI		ADNI		COBRE	
Method	r↑	RMSE↓	Acc	Sen	Acc	Sen	Acc	Sen
BrainNetCNN	.25±.02	12.86±1.12	.80±.04	.82±.07	.81±.05	.80±.03	.79±.02	.78±.02
M-GCN	.26±.02	12.63±.68	.90±.05	.91±.06	.81±.05	.80±.05	.82±.03	.80±.03
MV-GCN	.27±.03	12.18±.86	.92±.06	.92±.05	.83±.04	.81±.06	.82±.03	.81±.06
GraphGPS	.28±.01	12.25±.67	.93±.08	.93±.06	.85±.07	.84±.06	.84±.08	.84±.06
Cross-GNN	.28±.02	11.95±.82	.94±.07	.94±.08	.89±.06	.86±.07	.85±.06	.84±.04
MS-Inter-GCN	.31±.01	11.32±.77	.96±.06	.96±.07	.90±.07	.88±.05	.89±.06	.89±.05
GM-SSFuse	.30±.02	11.53±.73	.97±.05	.97±.05	.90±.05	.92±.07	**.92±.05**	.90±.06
GM-SSFuse-H	**.32±.02**	**10.97±.69**	**.97±.03**	**.97±.03**	**.91±.06**	**.92±.06**	.89±.05	**.90±.02**

Table 2. SC–FC interaction analysis. ρ: active region pair fraction (%); $\bar{d}$: mean hop distance; Q: cross-module interaction ratio (%). Diagonal baseline: ρ=0.9, $\bar{d}$=0, Q=0.

	HCP				ADNI				COBRE				PPMI			
Method	ρ↑	$\bar{d}$↑	Q↑	Pairs	ρ↑	$\bar{d}$↑	Q↑	Pairs	ρ↑	$\bar{d}$↑	Q↑	Pairs	ρ↑	$\bar{d}$↑	Q↑	Pairs
Attention	19.6	.96	84.8	2640	19.7	.96	84.8	2646	19.7	.95	84.8	2646	19.7	.96	84.9	2648
GM-SSFuse	43.1	.97	85.1	5800	42.8	.97	85.1	5762	46.2	**.98**	85.0	5814	43.0	.97	85.2	5788
GM-SSFuse-H	**49.3**	**.98**	**85.5**	**6636**	**48.6**	**.98**	**85.5**	**6540**	**49.4**	.97	**85.4**	**6653**	**49.1**	**.98**	**85.5**	**6613**

many-to-many fusion is particularly effective for schizophrenia where dysconnectivity is globally distributed and does not respect module boundaries.

SC–FC interaction analysis. Table 2 quantifies the breadth and distribution of structural-functional coupling. The one-to-one diagonal baseline activates only 116 pairs (ρ=0.9%, $\bar{d}$=0, Q=0% by construction). Attention-based fusion expands this to $\rho\approx$19.7% but remains well below the proposed methods. GM-SSFuse roughly doubles the active pair fraction ($\rho>$42%), and GM-SSFuse-H further increases this to $\rho\approx$49%, engaging over 6500 region pairs. The consistently high cross-module ratios ($Q\approx$85%) and mean hop distances ($\bar{d}\approx$0.97–0.98) confirm that SSM-based fusion captures genuinely distributed, long-range internetwork dependencies rather than reinforcing local homologous connections.

Ablation and qualitative analysis. Fig. 3 (top) ablates five fusion strategies. Concat performs worst across all panels, confirming that element-wise combination of modality embeddings is insufficient. Cross-attention improves substantially but remain below the proposed methods. Bidir-SSM applies bidirectional scanning to a concatenated $[Z_s; Z_f]$ sequence without interleaving, and yields mixed results relative to Cross-Attn, indicating that the alternating structural-functional ordering in GM-SSFuse is a critical factor. GM-SSFuse consistently surpasses Bidir-SSM, and GM-SSFuse-H achieves the best results on HCP and PPMI and ADNI, confirming the benefit of biologically-informed modular fusion. Fig. 3 (bottom) shows the most discriminative SC–FC coupling patterns per task. Fluid cognition is dominated by frontal-parietal and default mode inter-

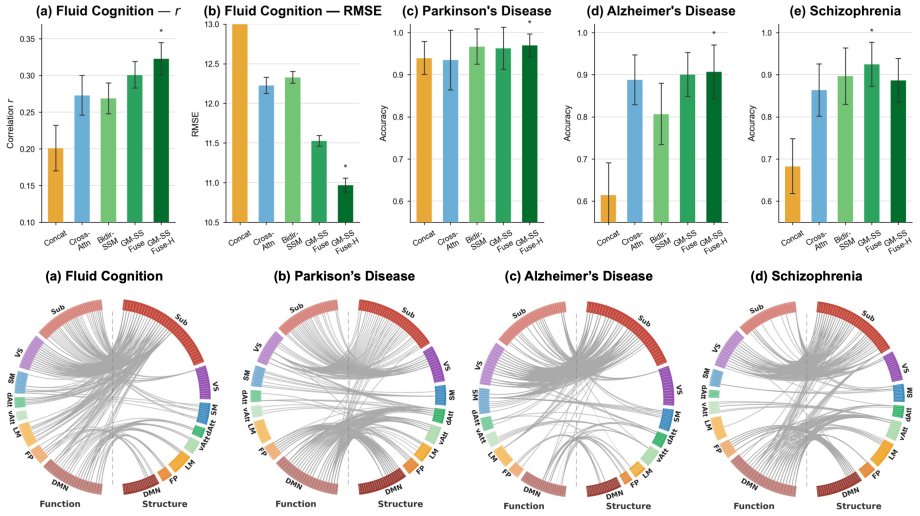

Fig. 3. *Top row:* Ablation of fusion strategies across five evaluation panels: (a) HCP Pearson r; (b) HCP RMSE (lower is better); (c–e) accuracy on PPMI, ADNI, COBRE. Cross-Attn baseline computes dense $N \times N$ attention between SC and FC node embeddings.Stars mark the best variant per panel; error bars show ± 1 std. *Bottom row:* Top 0.5% discriminative SC–FC coupling interactions for (a) fluid cognition, (b) PD, (c) AD, and (d) SZ. Ribbon width reflects coupling strength; arc size is proportional to aggregate subnetwork weight.

actions, consistent with executive and memory demands. PD shows pronounced subcortical-motor decoupling, while AD exhibits hippocampal-medial temporal disruption. SZ is characterised by inter-network dysconnectivity spanning frontal and subcortical systems, in line with the disconnection hypothesis and the observed advantage of flat over hierarchical fusion on that dataset.

4 Conclusion

We presented GM-SSFuse, a framework for multimodal brain connectome analysis that replaces one-to-one SC–FC interaction modelling with many-to-many coupling through sequential state accumulation in state-space models. By constructing an interleaved structural-functional sequence, each structural region naturally influences all functional regions without explicitly enumerating region pairs. A hierarchical variant further incorporates modular brain organisation to constrain within-network fusion while enabling cross-network interactions. Experiments across four datasets demonstrate consistent improvements over six baselines on both cognitive prediction and disease classification. Quantitative interaction analysis confirms that SSM-based fusion engages substantially more region pairs at greater graph distances than attention-based alternatives, and gradient-based interpretation recovers neuroanatomically meaningful coupling

patterns consistent with known disease mechanisms. These results establish sequential state-space models as an effective and interpretable architecture for structure-function coupling in brain network analysis.

Acknowledgement. This research is partially funded by a Tier-1 grant RG15/24 of the Ministry of Education, Singapore.

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

References

1. Aine, C., Bockholt, H.J., Bustillo, J.R., Cañive, J.M., Caprihan, A., Gasparovic, C., Hanlon, F.M., Houck, J.M., Jung, R.E., Lauriello, J., et al.: Multimodal neuroimaging in schizophrenia: description and dissemination. Neuroinformatics **15**(4), 343–364 (2017)
2. Baum, G.L., Cui, Z., Roalf, D.R., Ciric, R., Betzel, R.F., Larsen, B., Cieslak, M., Cook, P.A., Xia, C.H., Moore, T.M., et al.: Development of structure–function coupling in human brain networks during youth. Proceedings of the National Academy of Sciences **117**(1), 771–778 (2020)
3. Chan, Y.H., Girish, D., Gupta, S., Xia, J., Kasi, C., He, Y., Wang, C., Rajapakse, J.C.: Discovering robust biomarkers of psychiatric disorders from resting-state functional mri via graph neural networks: A systematic review. NeuroImage p. 121422 (2025)
4. Dan, T., Kim, M., Kim, W.H., Wu, G.: Uncovering structural-functional coupling alterations for neurodegenerative diseases. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 87–96. Springer (2023)
5. Dong, X., Li, Q., Wang, X., He, Y., Zeng, D., Chu, L., Zhao, K., Li, S.: How brain structure–function decoupling supports individual cognition and its molecular mechanism. Human Brain Mapping **45**(2), e26575 (2024)
6. Dsouza, N.S., Nebel, M.B., Crocetti, D., Robinson, J., Mostofsky, S., Venkataraman, A.: M-gcn: A multimodal graph convolutional network to integrate functional and structural connectomics data to predict multidimensional phenotypic characterizations. In: Medical imaging with deep learning. pp. 119–130. PMLR (2021)
7. Esteban, O., Markiewicz, C.J., Blair, R.W., Moodie, C.A., Isik, A.I., Erramuzpe, A., Kent, J.D., Goncalves, M., DuPre, E., Snyder, M., et al.: fmriprep: a robust preprocessing pipeline for functional mri. Nature methods **16**(1), 111–116 (2019)
8. Fotiadis, P., Parkes, L., Davis, K.A., Satterthwaite, T.D., Shinohara, R.T., Bassett, D.S.: Structure–function coupling in macroscale human brain networks. Nature Reviews Neuroscience **25**(10), 688–704 (2024)
9. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. In: First conference on language modeling (2024)
10. Jack Jr, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., L. Whitwell, J., Ward, C., et al.: The alzheimer's disease neuroimaging initiative (adni): Mri methods. Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine **27**(4), 685–691 (2008)
11. Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M.: Fsl. Neuroimage **62**(2), 782–790 (2012)

12. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment. NeuroImage **146**, 1038–1049 (2017)
13. Li, W., Zhou, H., Yu, J., Song, Z., Yang, W.: Coupled mamba: Enhanced multimodal fusion with coupled state space model. Advances in Neural Information Processing Systems **37**, 59808–59832 (2024)
14. Marek, K., Jennings, D., Lasch, S., Siderowf, A., Tanner, C., Simuni, T., Coffey, C., Kieburtz, K., Flagg, E., Chowdhury, S., et al.: The parkinson progression marker initiative (ppmi). Progress in neurobiology **95**(4), 629–635 (2011)
15. Pipoli, V., Saporita, A., Marchesini, K., Grana, C., Ficarra, E., Bolelli, F.: Im-fuse: A mamba-based fusion block for brain tumor segmentation with incomplete modalities. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 225–235. Springer (2025)
16. Rampášek, L., Galkin, M., Dwivedi, V.P., Luu, A.T., Wolf, G., Beaini, D.: Recipe for a general, powerful, scalable graph transformer. Advances in Neural Information Processing Systems **35**, 14501–14515 (2022)
17. Routier, A., Burgos, N., Díaz, M., Bacci, M., Bottani, S., El-Rifai, O., Fontanella, S., Gori, P., Guillon, J., Guyot, A., et al.: Clinica: An open-source software platform for reproducible clinical neuroscience studies. Frontiers in neuroinformatics **15**, 689675 (2021)
18. Shang, S., Wang, L., Xu, Y., Zhang, H., Chen, L., Dou, W., Yin, X., Ye, J., Chen, Y.C.: Optimization of structural connectomes and scaled patterns of structural-functional decoupling in parkinson's disease. Neuroimage **284**, 120450 (2023)
19. Suárez, L.E., Markello, R.D., Betzel, R.F., Misic, B.: Linking structure and function in macroscale brain networks. Trends in cognitive sciences **24**(4), 302–315 (2020)
20. Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M.: Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. Neuroimage **15**(1), 273–289 (2002)
21. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al.: The wu-minn human connectome project: an overview. Neuroimage **80**, 62–79 (2013)
22. Wang, C., Tsepa, O., Ma, J., Wang, B.: Graph-mamba: Towards long-range graph sequence modeling with selective state spaces. arXiv preprint arXiv:2402.00789 (2024)
23. Wang, J., Lang, J., Yang, L.Z., Li, H.: Learning dynamic brain network representation based on graph mamba architecture. In: 2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). pp. 5143–5150. IEEE (2024)
24. Xia, J., Chan, Y.H., Girish, D., Rajapakse, J.C.: Interpretable modality-specific and interactive graph convolutional network on brain functional and structural connectomes. Medical Image Analysis **102**, 103509 (2025)
25. Yang, Y., Ye, C., Guo, X., Wu, T., Xiang, Y., Ma, T.: Mapping multi-modal brain connectome for brain disorder diagnosis via cross-modal mutual learning. IEEE Transactions on Medical Imaging **43**(1), 108–121 (2023)
26. Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., et al.: The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of neurophysiology (2011)

27. Zarkali, A., McColgan, P., Leyland, L.A., Lees, A.J., Rees, G., Weil, R.S.: Organisational and neuromodulatory underpinnings of structural-functional connectivity decoupling in patients with parkinson's disease. Communications biology **4**(1), 86 (2021)
28. Zhang, X., He, L., Chen, K., Luo, Y., Zhou, J., Wang, F.: Multi-view graph convolutional network and its applications on neuroimage analysis for parkinson's disease. In: AMIA annual symposium proceedings. vol. 2018, p. 1147 (2018)