ONE SHARED GRAPH, TWO MAPPINGS: BIDIRECTIONAL GRAPH AUTOENCODER FOR BRAIN STRUCTURAL AND FUNCTIONAL CONNECTIVITY

Anonymous authors
Paper under double-blind review

ABSTRACT

Structural connectivity (SC) derived from diffusion MRI and functional connectivity (FC) derived from functional MRI provide complementary summaries of brain organization across large-scale networks. Because coherent brain function is supported by anatomically grounded, polysynaptic communication pathways, we model SC and FC as arising from a shared, biologically meaningful scaffold and propose a bidirectional graph autoencoder to estimate a shared latent graph common to both modalities. The method fits jointly trained forward and reverse mappings (SC$\rightarrow$FC and FC$\rightarrow$SC) that share a single latent adjacency matrix, which serves as the propagation operator for message passing in both modalities and links the two directions through a common graph over brain regions. On the landmark Human Connectome Project (HCP) and Adolescent Brain Cognitive Development (ABCD) data, the proposed approach improves both bidirectional reconstruction accuracy and graph reliability over existing graph-based and connectome translation methods, yielding roughly 7%-15% lower reconstruction error and about 4%-16% higher top-edge overlap. Moreover, the learned graphs exhibit strong cross-cohort consistency between HCP and ABCD, supporting the utility of the framework for interpretable brain structure-function modeling.

1 INTRODUCTION

Recent advances in neuroimaging provide rich, noninvasive measures of brain activity and brain-wide organization. Among commonly studied imaging traits, functional magnetic resonance imaging (fMRI) yields functional connectivity (FC) networks that summarize statistical dependencies in neural activity (Seguin et al., 2023), while diffusion MRI (dMRI) yields structural connectivity (SC) networks that approximate the brain's underlying white-matter wiring architecture (Grier et al., 2020). Notably, SC and FC are closely linked rather than independent. Anatomical pathways shape the routes through which neural signals propagate, and recurring patterns of coordinated activity can reinforce effective communication along these pathways (Rowland et al., 2023).

Emerging evidence further suggests that this structure–function relationship is organized around a shared scaffold that supports long-range coordination and integration across functional systems. FC is constrained by the anatomical connectome, yet it also reflects coordinated interactions that are not limited to directly connected region pairs (Honey et al., 2009; Sarwar et al., 2021; Xia et al., 2025). Additionally, scaffold-like structure has been associated with compact integrative substrates such as structural cores and rich-club systems, whose connections disproportionately facilitate global communication and coordination (Hagmann et al., 2008; van den Heuvel & Sporns, 2011). Together, these findings motivate modeling SC and FC through a shared latent graph that acts as a foundational scaffold shared by both modalities, interpretable as a set of key channels that supports distributed, brain wide coordination.

Characterizing this shared graph directly from paired SC and FC remains challenging because it is not observed in either modality. A natural approach is to use an autoencoder, which learns a low-dimensional representation that preserves the input structure while decoding it back to the observed data. In this setting, a graph autoencoder is especially well suited because it operates on network-structured inputs and can learn latent connectivity patterns that reflect consistent topology

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

across modalities rather than relying on edgewise correspondence alone (Xia et al., 2025). However, most existing graph autoencoder approaches for multimodal connectomics are trained primarily in a single direction, typically predicting FC from SC (Li et al., 2019; Amodeo et al., 2022; Cai et al., 2025), while translation frameworks that operate in both directions, such as Krakencoder (Jamison et al., 2025) and the bidirectional graph GAN of Zhou et al. (2025), share information across modalities through a latent embedding or paired generators rather than through a single graph over brain regions. In either case, no single scaffold is required to be simultaneously compatible with both modalities, which can lead to direction-specific representations, weaker cross-modal identifiability, and reduced robustness.

To address this gap, we propose a *Bidirectional Graph Autoencoder with a shared latent adjacency matrix* for SC–FC modeling (**BGAA**). BGAA learns jointly trained forward and reverse mappings, SC→FC and FC→SC, that are tied through a single latent adjacency acting as the common propagation scaffold for both modalities, so that the structure shared by the two directions is itself an interpretable graph over brain regions. This shared-graph formulation treats the structure–function relationship as a two-way constraint on a latent communication network, allowing the model to identify a compact set of edges that is jointly consistent with anatomical wiring and functional coordination. By recovering an interpretable scaffold rather than only optimizing cross-modal prediction, BGAA yields a biologically grounded representation of large-scale brain organization that can serve as a compact substrate for downstream biomedical analysis.

Contributions.

- We introduce **BGAA**, a bidirectional graph autoencoder for joint SC–FC modeling with a shared adjacency that serves as a common scaffold for both SC→FC and FC→SC mappings.
- A walk-algebra-based formulation characterizes the shared SC–FC graph as a compact scaffold that supports both modalities through multi-step communication structure.
- Direct graph learning with anatomy informed construction and explicit regularization enables joint optimization of the shared adjacency together with bidirectional reconstruction.
- Experiments on the HCP and ABCD datasets demonstrate that BGAA learns a more accurate, stable, and biologically plausible shared graph than competitive baselines, while revealing consistent system-level organization across cohorts.

2 Related Work

A substantial literature shows that structural connectivity (SC) and functional connectivity (FC) provide complementary views of large-scale brain organization and are linked through anatomically grounded but not purely one-step communication processes (Hagmann et al., 2008; Honey et al., 2009; Goñi et al., 2014; Sarwar et al., 2021; Rowland et al., 2023; Zalesky et al., 2024). At the same time, growing evidence suggests that the SC–FC relationship is organized around a shared mesoscale substrate rather than being distributed uniformly across cortex (Suárez et al., 2020; Baum et al., 2020; Vázquez-Rodríguez et al., 2019; Chan et al., 2022; van den Heuvel et al., 2012; Powell et al., 2024). These findings motivate models that recover a compact latent scaffold capable of supporting both anatomical connections and coordinated functional interactions, especially in association-rich and integrative systems.

Methodologically, our work is related to graph-based representation learning and cross-modal connectome mapping. Graph autoencoders and multimodal encoder-decoder frameworks have been used to learn latent representations from paired structural and functional connectomes and to reconstruct one modality from the other (Li et al., 2019; Liu et al., 2021; Xia et al., 2025; Jamison et al., 2025; Cai et al., 2025; Zhou et al., 2025). Diffusion-based graph propagation models capture communication over connectomes through random-walk-style aggregation (Atwood & Towsley, 2016), while higher-order architectures such as MixHop explicitly combine information across multiple hop distances (Abu-El-Haija et al., 2019). More recent connectome-specific graph learning frameworks, including Krakencoder and BrainGB, have further emphasized multimodal latent fusion and standardized graph modeling for brain networks (Jamison et al., 2025; Cui et al., 2023). BGAA builds on this line of work by centering the model on a single shared latent adjacency that is learned as a shared scaffold and used simultaneously for both SC→FC and FC→SC mappings.

3 PROBLEM FORMULATION

3.1 SETUP AND NOTATIONS

Let V be the number of brain regions. For subject $i \in \{1, \dots, N\}$, we observe a SC matrix $\mathbf{S}_i \in \mathbb{R}^{V \times V}$ derived from diffusion MRI tractography and a FC matrix $\mathbf{F}_i \in \mathbb{R}^{V \times V}$ derived from resting-state fMRI. After preprocessing, we treat both $\mathbf{S}_i$ and $\mathbf{F}_i$ as symmetric matrices with zero diagonal, where $\mathbf{S}_i$ summarizes anatomical connectivity between regions and $\mathbf{F}_i$ summarizes statistical dependence in regional activity.

Our objective is to characterize a *shared latent graph* with adjacency matrix $\mathbf{M} \in \mathbb{R}^{V \times V}$ that captures cross modal communication structure common to SC and FC. We aim to define the graph in a biologically grounded manner that reflects how anatomical wiring supports functional interactions and how functional coordination reveals effective communication supported by that wiring. Accordingly, the graph should represent polysynaptic pathways and routing structure that are jointly compatible with both SC and FC, rather than reducing to direct tracts alone or simply selecting the strongest FC correlations. Guided by this definition, we develop a statistical learning framework to estimate $\mathbf{M}$ from paired connectomes observed across subjects.

3.2 SHARED GRAPH DEFINITION

Empirical and modeling work suggests that SC and FC are mutually informative. Anatomical wiring constrains functional interactions, while functional interactions reflect effective communication supported by polysynaptic routes rather than only direct tracts. Walk based representations capture this principle because powers of an adjacency encode length-k pathways, and low degree combinations of these powers encompass communication-like transformations that have been repeatedly linked to SC–FC correspondence (Defferrard et al., 2016; Seguin et al., 2020).

Let $\mathbb{S}^V$ denote symmetric $V \times V$ matrices. For an integer $K \geq 1$ and an adjacency $M \in \mathbb{S}^V$, define the degree-K *walk algebra* generated by M as

$$\mathcal{A}_K(M) := \Big\{ \sum_{k=0}^{K} c_k M^k : \; c_k \in \mathbb{R} \Big\}, \tag{1}$$

where M^k aggregates information carried by length-k walks.

Given paired connectomes $\{(S_i, F_i)\}_{i=1}^N$ with $(S_i, F_i) \in \mathbb{S}^V \times \mathbb{S}^V$, we define the shared SC–FC graph as the minimal adjacency whose walk algebra can support both modalities across subjects.

Definition 3.1 (Shared SC–FC graph)**.** Let $\Omega(\cdot)$ be a complexity functional. A matrix $M^\star \in \mathbb{S}^V$ is a *shared SC–FC graph* if

$$M^\star \in \arg\min_{M \in \mathbb{S}^V} \; \Omega(M) \quad \text{s.t.} \quad S_i \in \mathcal{A}_K(M), \;\; F_i \in \mathcal{A}_K(M), \qquad i = 1, \dots, N. \tag{2}$$

Equivalently, for each subject i there exist coefficient vectors $c_i^{(S)}, c_i^{(F)} \in \mathbb{R}^{K+1}$ such that

$$S_i = \sum_{k=0}^{K} c_{i,k}^{(S)} M^k, \qquad F_i = \sum_{k=0}^{K} c_{i,k}^{(F)} M^k. \tag{3}$$

The graph $M^\star$ represents a shared communication scaffold that supports both anatomical and functional organization through multi step pathways. Each entry $M^\star_{uv}$ can be interpreted as the strength of a latent communication channel between regions u and v that is consistently useful for explaining both SC and FC, while the powers $(M^\star)^k$ summarize contributions from length-k routes. Differences between SC and FC are captured by subject specific coefficient profiles $c_i^{(S)}$ and $c_i^{(F)}$, whereas the graph itself summarizes the cross modal structure that is most essential under the chosen complexity criterion $\Omega(\cdot)$.

This definition specifies the structural role of the graph, but it does not by itself yield a practical estimator. In real data, SC and FC are noisy, heterogeneous, and only approximately compatible with a low degree walk algebra representation, so directly solving for a shared nonnegative adjacency M

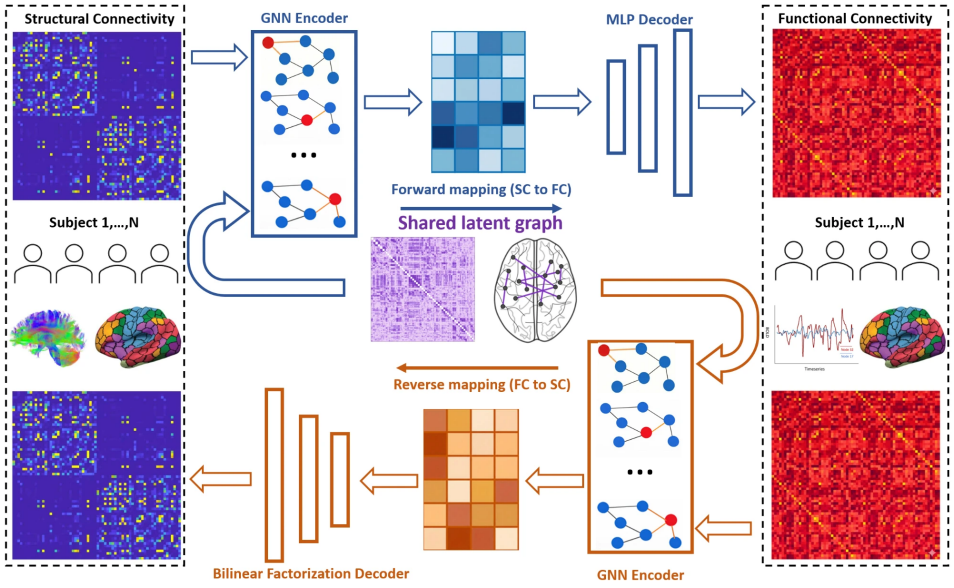


Figure 1: Overview of the proposed **BGAA** framework for learning a shared SC–FC graph. For each subject, structural connectivity (SC, left) and functional connectivity (FC, right) are linked by two mappings that share a common latent graph.

together with subject specific coefficient vectors is highly nonconvex and potentially nonidentifiable. Moreover, the definition alone does not specify how to regularize M so that the learned scaffold is sparse, stable, and biologically plausible. These challenges motivate the BGAA framework in Section 4, which treats the shared graph as a latent adjacency matrix and estimates it indirectly through a bidirectional reconstruction model with explicit sparsity regularization.

4 METHOD

Section 3.2 characterizes the shared graph as a scaffold that supports both SC and FC through multi-step pathways and modality specific transformations. We now introduce *Bidirectional Graph Autoencoder with a shared latent Adjacency matrix* for structural–functional connectivity (**BGAA**), which serves as a practical estimator of this idea. As summarized in Figure 1, BGAA learns a single population-level adjacency $\mathbf{M}$ as the shared scaffold and uses it to construct two linked mappings between SC and FC. The shared adjacency represents the common communication structure, while the two directional mappings capture modality specific transformations built on top of that scaffold.

4.1 SHARED GRAPH CONSTRUCTION

The central object in BGAA is the shared latent adjacency matrix $\mathbf{M}$, which serves as the graph used by both the SC$\rightarrow$FC and FC$\rightarrow$SC mappings. Rather than introducing a distribution over graphs, we directly parameterize and optimize a single sparse, nonnegative, symmetric graph at the population level. This design separates the role of the graph from subject specific connectomes. The graph is intended to represent a common cross modal scaffold, while the observed SC and FC matrices enter only through the bidirectional reconstruction paths introduced in Section 4.2.

We include anatomical meta features to guide the learned graph toward biologically plausible organization. In effect, these features define an anatomy informed similarity kernel over brain regions, so the graph is not learned as an unconstrained free graph.

Each brain region u is associated with a feature vector $\mathbf{x}_u \in \mathbb{R}^p$ encoding anatomical and organizational properties, including centroid coordinates in MNI space, hemisphere indicator, and canonical functional system assignment. These features are embedded into a latent space through a learnable

map

$$\mathbf{e}_u = h_\psi(\mathbf{x}_u) \in \mathbb{R}^d, \tag{4}$$

where h_ψ is a shallow multilayer perceptron. From the embeddings, we compute two projected representations

$$\mathbf{q}_u = \mathbf{Q}_e \mathbf{e}_u, \qquad \mathbf{p}_u = \mathbf{P}_e \mathbf{e}_u, \tag{5}$$

and define a symmetric edge score

$$\ell_{uv} = \frac{1}{2\sqrt{d}} \left(\mathbf{q}_u^\top \mathbf{p}_v + \mathbf{q}_v^\top \mathbf{p}_u \right) + r_{uv}, \tag{6}$$

where r_{uv} is a free learnable residual parameter on the upper triangle with $r_{vu} = r_{uv}$. The attention term supplies an anatomy informed kernel between regions, while the residual term allows deviations that are not captured by the anatomical feature map alone.

We then construct the graph directly from these scores:

$$M_{uv} = \mathrm{softplus}(\ell_{uv})(1 - \delta_{uv}), \qquad M_{vu} = M_{uv}, \tag{7}$$

which enforces symmetry, nonnegativity, and zero diagonal. After training, the learned population-level graph is given directly by the optimized matrix $\mathbf{M}$. Each entry M_{uv} can therefore be interpreted as the strength of a latent communication channel between regions u and v in the shared scaffold.

4.2 Shared-graph bidirectional autoencoder

Given the learned graph $\mathbf{M}$, BGAA combines two reconstruction paths within a single bidirectional autoencoder. The forward and reverse paths play the role of two modality specific transformations supported by the same scaffold. In this way, the shared adjacency captures common cross modal structure, while the directional mappings account for differences between how that structure is expressed in SC and FC.

Forward path: SC $\rightarrow$ FC. Subject i's structural connectivity is propagated as node features over the shared graph:

$$\mathbf{B}_i = \mathrm{GCN}(\mathbf{S}_i, \mathbf{M} \mid \Xi_s), \qquad \mathbf{B}_i \in \mathbb{R}^{V \times d}. \tag{8}$$

The resulting node embeddings are decoded into a functional connectivity estimate:

$$\widetilde{\mathbf{F}}_i = g_f(\mathbf{B}_i), \tag{9}$$

where $g_f(\cdot)$ is an MLP with symmetrized, zero-diagonal output.

Reverse path: FC $\rightarrow$ SC. Functional connectivity is propagated over the same graph:

$$\mathbf{Q}_i = \mathrm{GCN}(\mathbf{F}_i, \mathbf{M} \mid \Xi_f), \qquad \mathbf{Q}_i \in \mathbb{R}^{V \times d}. \tag{10}$$

To decode SC, we use a symmetric bilinear mixture

$$\widetilde{\mathbf{S}}_i = \sum_{j=1}^{J} \mathbf{Q}_i \, \mathbf{T}_j \, \mathbf{Q}_i^\top, \tag{11}$$

where $\{\mathbf{T}_j\}_{j=1}^J$ are learnable symmetric $d \times d$ matrices. This decoder is aligned with the topology of structural connectivity in two ways. First, SC is naturally modeled as a symmetric affinity between regions, and Eq. (11) guarantees symmetry by construction since each $\mathbf{T}_j$ is symmetric. Second, the bilinear form ties each predicted edge to pairwise compatibility between region embeddings, so edges emerge from shared latent factors that encode system-level organization. As the components enter linearly, the mixture is as expressive as a single bilinear form with $\sum_j \mathbf{T}_j$, and J mainly shapes optimization and the effective weight decay (Appendix C.3). This low-rank parameterization stabilizes estimation when SC is sparse and heavy-tailed.

Shared message passing. We use standard normalized message passing with the same graph in both directions. Define

$$\bar{\mathbf{M}} := \mathbf{M} + \mathbf{I}, \qquad \mathbf{D} := \mathrm{diag}\big(\bar{\mathbf{M}}\mathbf{1}\big), \qquad \widehat{\mathbf{M}} := \mathbf{D}^{-1/2}\,\bar{\mathbf{M}}\,\mathbf{D}^{-1/2}. \tag{12}$$

An L-layer GCN updates node features by

$$\mathbf{H}^{(\ell+1)} = \sigma\Big(\widehat{\mathbf{M}}\,\mathbf{H}^{(\ell)}\mathbf{W}^{(\ell)}\Big), \qquad \ell = 0, \dots, L-1, \tag{13}$$

with $\mathbf{H}^{(0)} = \mathbf{S}_i$ for the forward path and $\mathbf{H}^{(0)} = \mathbf{F}_i$ for the reverse path. We take $\mathbf{B}_i = \mathbf{H}^{(L)}$ for the SC→FC direction and $\mathbf{Q}_i = \mathbf{H}^{(L)}$ for the FC→SC direction.

This encoder structure is naturally aligned with the walk-based graph formulation. Each propagation step aggregates information from neighboring regions along the shared graph, and stacking layers builds increasingly multi-step features from repeated application of the same graph operator. Because $\widehat{\mathbf{M}}$ is the only structural operator shared by the two paths, joint training favors an adjacency whose multi-step walks inform both SC and FC, the property formalized in Definition 3.1 and measured directly by the walk-algebra fit in Section 5.

4.3 Training objective

We optimize BGAA as a direct graph learning problem. Let

$$\Theta := \{\psi, \mathbf{Q}_e, \mathbf{P}_e, \mathbf{r}, \Xi_s, \Xi_f, g_f, \{\mathbf{T}_j\}_{j=1}^{J}\}$$

collect all trainable parameters other than the induced adjacency matrix $\mathbf{M}$. We then jointly optimize the graph and the bidirectional autoencoder by minimizing

$$\min_{\Theta} \sum_{i=1}^{N} \Big(\|\mathbf{F}_i - \widetilde{\mathbf{F}}_i\|_F^2 + \|\mathbf{S}_i - \widetilde{\mathbf{S}}_i\|_F^2 \Big) + \lambda_1 \|\mathbf{M}\|_1 + \gamma \|\mathbf{M}\mathbf{1}\|_2^2 + \mu \|\mathbf{r}\|_1, \tag{14}$$

where $\mathbf{M}$ is constructed from Eq. (7). The first two terms are the forward and reverse reconstruction losses, which require the same graph to be informative for both SC→FC and FC→SC mapping.

The three regularization terms serve distinct roles. The ℓ_1 penalty $\lambda_1\|\mathbf{M}\|_1$ promotes graph sparsity. As softplus keeps all entries positive, this penalty shrinks weak edges toward zero and concentrates weight on a small set of dominant edges. The degree penalty $\gamma\|\mathbf{M}\mathbf{1}\|_2^2$ discourages concentration of graph edges at hub nodes and instead promotes a more distributed scaffold. The residual penalty $\mu\|\mathbf{r}\|_1$ encourages the model to explain graph edges through the anatomy informed structure learning module rather than the free residual, which makes the two component decomposition more identifiable.

Taken together, this objective provides a direct learning criterion for a compact shared scaffold that is informative for both modalities. The bidirectional reconstruction terms favor a graph that supports both mappings, while the regularization terms favor structured, and biologically plausible solutions.

5 Experiment

5.1 Data overview

We evaluated our model on paired diffusion MRI and resting-state fMRI data from two large cohorts: the Human Connectome Project (HCP) S1200 release (Van Essen et al., 2013) and the Adolescent Brain Cognitive Development (ABCD) study (Casey et al., 2018). For each subject i, we constructed a structural connectivity matrix $\mathbf{S}_i \in \mathbb{R}^{68\times68}$ and a functional connectivity matrix $\mathbf{F}_i \in \mathbb{R}^{68\times68}$ using the Desikan–Killiany (DK) parcellation ($V = 68$). Structural connectivity $\mathbf{S}_i$ was derived from dMRI tractography by summarizing white-matter connectivity strength between each pair of regions (e.g., streamline-based weights), followed by symmetrization, removal of self-connections, and subject-wise normalization to reduce inter-subject scale variability. Functional connectivity $\mathbf{F}_i$ was computed from preprocessed resting-state fMRI by extracting regional time series and forming Pearson correlation matrices, which were Fisher z-transformed and symmetrized with zero diagonal. After restricting to subjects with both modalities available and passing quality control, our analyses included $N = 1065$ individuals from HCP and $N = 2360$ individuals from ABCD, which were split into training/validation/test sets within each cohort.

5.2 BASELINE COMPARISON

We compare BGAA against four representative baselines that capture complementary ways of modeling SC and FC. **DGNN** is a diffusion based graph neural network that aggregates random walk style propagation signals and serves as a strong communication oriented baseline for connectome prediction (Atwood & Towsley, 2016). **MixHop** is a higher order GNN that mixes information from multiple hop distances within each layer, making it a relevant baseline for polysynaptic structure–function relationships (Abu-El-Haija et al., 2019). **Krak** is a multimodal connectome translation and fusion framework that learns shared latent representations across structural and functional connectivity, providing a strong autoencoding baseline for cross modal reconstruction (Jamison et al., 2025). **BrainGB** is a benchmark framework for brain network analysis with graph neural networks, and we include it as a strong connectome specific baseline that standardizes GNN modeling choices across structural and functional brain graphs (Cui et al., 2023). Together, these methods allow us to compare BGAA against diffusion style propagation, multi hop message passing, multimodal latent fusion, and connectome specific GNN modeling. All baselines are trained in both directions on the same splits (Appendix B.2).

We evaluate each method from three complementary perspectives that together provide a more complete view of model quality: bidirectional reconstruction accuracy (BRA), test–retest stability (TRS), and walk-algebra fit (WAF). BRA evaluates the full end-to-end prediction pipeline by measuring how well the model reconstructs both modalities on held-out subjects, and we report its total value together with the directional components for FC and SC. TRS evaluates the learned graph itself by refitting each method on two independent training replicates and comparing the resulting graphs using the Spearman correlation of edge weights, denoted by pTRS_{ρ}, and the Jaccard overlap of the top 1% and top 10% graph edges by weight, denoted by $\text{pTRS}_{J}(1\%)$ and $\text{pTRS}_{J}(10\%)$. WAF evaluates whether the learned graph is also structurally consistent with the walk-based formulation in Section 3.2. Given a learned graph $\widehat{\mathbf{M}}$, we reconstruct each subject's SC and FC by fitting low-degree linear combinations of the powers of $\widehat{\mathbf{M}}$ after training, and then measure the resulting reconstruction error. We report the SC component WAF_S, the FC component WAF_F, and their sum WAF. In this sense, WAF tests the quality of the graph as a shared multi-step communication scaffold, independently of the flexibility of the encoder and decoder modules.

With these evaluation criteria in place, we implement BGAA and all competing baselines on the HCP and ABCD datasets and summarize their performance in Table 1. Across datasets, BGAA delivers the strongest overall profile on BRA, TRS, and WAF, which highlights three complementary advantages of the model. The lower BRA values show that the shared-graph bidirectional architecture captures cross-modal dependence more effectively than competing approaches. In particular, using a single latent adjacency for both SC→FC and FC→SC leads to more accurate reconstruction than methods based only on diffusion style propagation, multi-hop aggregation, or general latent fusion, because the same scaffold is optimized to support both directions simultaneously. The improvement over Krak, which also learns both directions but ties them through a shared latent vector, together with the ablation in Appendix C.1, indicates that sharing the adjacency itself contributes to this gain. The stronger TRS results further show that BGAA does not gain accuracy by fitting unstable graphs, as its learned adjacency is more reproducible across data perturbations, especially among the highest-weight edges, and is therefore stably determined by the data. The WAF results show that low-degree powers of the BGAA graph reconstruct both SC and FC more accurately than those of competing graphs, indicating closer alignment with the walk-based formulation and a multi-step structure that is more informative for both modalities.

5.3 ANALYSIS OF THE LEARNED SCAFFOLD

To further examine the structure of the learned graph beyond the metrics in Table 1, Figure 2 provides a complementary view of its organization across cohorts at both the node and functional-system levels. The node-level weight matrices show that the strongest weights are not scattered broadly, but concentrate on a small number of recurring system pairs that appear in both HCP and ABCD. This cross-cohort consistency makes the learned scaffold more plausible as a population-level organizing structure rather than a dataset-specific artifact. The lower correlation matrix further reinforces this point by showing high agreement between the two cohorts at the functional-system level, especially

Table 1: Estimates with standard deviations for BRA, TRS, and WAF on HCP and ABCD (details in Appendix B.2–B.3).

		HCP					ABCD				
Metric	**Criterion**	**BGAA**	**DGNN**	**MixHop**	**Krak**	**BrainGB**	**BGAA**	**DGNN**	**MixHop**	**Krak**	**BrainGB**
BRA	$\mathrm{BRA}_F \downarrow$	0.77 (0.12)	1.33 (0.25)	0.89 (0.19)	0.86 (0.18)	0.89 (0.19)	1.08 (0.13)	2.57 (0.24)	1.48 (0.19)	1.15 (0.13)	1.38 (0.17)
	$\mathrm{BRA}_S \downarrow$	1.22 (0.26)	2.41 (0.39)	1.54 (0.37)	1.27 (0.33)	1.30 (0.29)	1.65 (0.24)	3.01 (0.36)	2.19 (0.40)	2.06 (0.29)	2.23 (0.30)
	$\mathrm{BRA} \downarrow$	1.99 (0.37)	3.74 (0.61)	2.43 (0.55)	2.13 (0.41)	2.19 (0.40)	2.73 (0.36)	5.58 (0.55)	3.67 (0.54)	3.21 (0.38)	3.61 (0.45)
TRS	$\mathrm{pTRS}_p \uparrow$	0.53 (0.08)	0.40 (0.15)	0.41 (0.09)	0.38 (0.07)	0.41 (0.11)	0.41 (0.06)	0.39 (0.14)	0.29 (0.11)	0.43 (0.09)	0.41 (0.09)
	$\mathrm{pTRS}_J(1\%) \uparrow$	0.50 (0.03)	0.38 (0.14)	0.44 (0.09)	0.43 (0.08)	0.48 (0.07)	0.45 (0.04)	0.35 (0.10)	0.29 (0.06)	0.39 (0.09)	0.40 (0.05)
	$\mathrm{pTRS}_J(10\%) \uparrow$	0.58 (0.08)	0.50 (0.17)	0.45 (0.15)	0.40 (0.09)	0.43 (0.14)	0.46 (0.10)	0.37 (0.15)	0.30 (0.18)	0.42 (0.10)	0.37 (0.11)
WAF	$\mathrm{WAF}_F \downarrow$	1.10 (0.19)	2.21 (0.33)	1.09 (0.29)	1.69 (0.29)	1.91 (0.27)	1.23 (0.17)	3.29 (0.29)	1.49 (0.31)	1.57 (0.28)	1.74 (0.30)
	$\mathrm{WAF}_S \downarrow$	1.48 (0.30)	3.48 (0.45)	1.57 (0.37)	2.02 (0.33)	2.88 (0.28)	1.78 (0.25)	2.92 (0.39)	2.43 (0.30)	2.54 (0.26)	3.01 (0.27)
	$\mathrm{WAF} \downarrow$	2.58 (0.47)	5.69 (0.73)	2.66 (0.64)	3.71 (0.54)	4.79 (0.49)	3.01 (0.40)	6.21 (0.64)	3.92 (0.60)	4.11 (0.51)	4.75 (0.55)

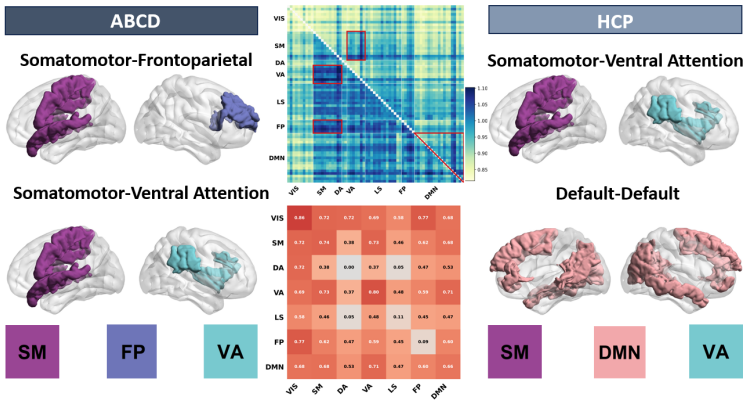


Figure 2: Cross-cohort organization of the learned shared SC–FC graph. The upper matrix shows node-level edge weights, with nodes ordered by functional system, the upper triangle from HCP, and the lower triangle from ABCD. Red boxes mark the three system pairs with the largest mean edge weight, which are visualized in the left and right panels. The lower matrix shows the correlation between the HCP and ABCD graphs.

for the most prominent system pairs. As system membership is only one soft node feature of the kernel and does not fix which pairs are strong, this concentration is learned from the data.

The highlighted system pairs also align well with current biological literature on large-scale brain organization. Strong within-system and selective cross-system edge weights involving Somatomotor, Ventral Attention, and Default-related systems are consistent with the view that the structure–function relationship is supported by a compact set of anatomically grounded, polysynaptic communication routes rather than only direct tracts or the largest FC correlations (Honey et al., 2009; Goñi et al., 2014; Suárez et al., 2020). In particular, Ventral Attention and Default-related interactions have been repeatedly linked to integrative coordination, salience-driven switching, and large-scale cognitive reconfiguration (Goulden et al., 2014; Tosoni et al., 2023; Menon, 2023), while strong Somatomotor involvement is consistent with the relatively stable within-system structure–function correspondence often observed in sensorimotor cortex (Vázquez-Rodríguez et al., 2019), and the concentration of weight on association systems agrees with the view that large-scale communication relies on a limited set of high-capacity routes (van den Heuvel et al., 2012; Power et al., 2011; Yeo et al., 2011). These results therefore suggest that BGAA is recovering a biologically plausible scaffold that supports both within-network coherence and targeted cross-network integration.

We further examined whether the learned graph carries information about individual differences in cognition. Within each cohort, cognitive scores were adjusted for age, sex, acquisition site, and head motion, subjects were split at the median of the adjusted score, and BGAA was fitted separately to each group, with system-level changes tested against a label-permutation null (Figure 3, Appendix E). The system pairs that gain weight in the high-cognition graphs mainly involve visual, attention, and control systems, with the largest change within the frontoparietal system (Figure 3a). This pattern reflects the central role of frontoparietal and attention networks in flexible cognitive control and the

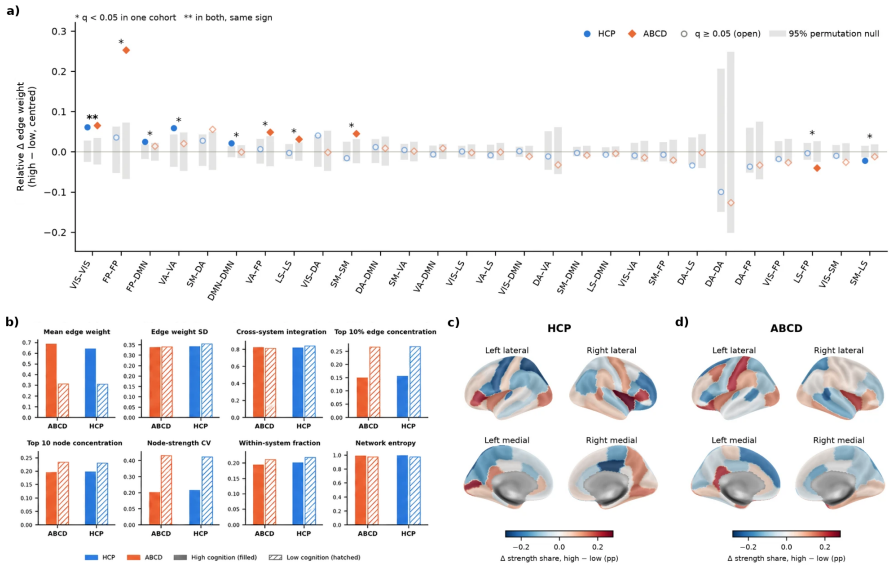


Figure 3: Differences between graphs fitted to high- and low-cognition groups. (a) Centred change in mean edge weight (high − low) for each system pair; grey bars show the 95% permutation null, and filled markers indicate Benjamini–Hochberg $q < 0.05$. (b) Global properties of the high- (filled) and low-cognition (hatched) graphs. (c, d) Change in each region's share of node strength. VIS: visual; SM: somatomotor; DA: dorsal attention; VA: ventral attention; LS: limbic; FP: frontoparietal; DMN: default mode.

coordination among salience, control, and default mode systems that supports higher-order cognition (Goulden et al., 2014; Tosoni et al., 2023; Menon, 2023). At the network level, higher cognition is associated with stronger edge weights that are less concentrated on a few dominant edges and regions, while cross-system integration is largely preserved (Figure 3b–d). These results indicate that better cognitive performance relies on a communication scaffold that spreads load across many channels instead of depending on a small set of hubs, consistent with the view that efficient large-scale communication balances high-capacity routes with broadly shared pathways (van den Heuvel et al., 2012). Because these differences appear in both a young adult and a developmental cohort, and structure–function correspondence matures during youth and relates to executive function (Baum et al., 2020; Chan et al., 2022), the shared graph provides a compact and interpretable readout of the structural substrate of cognition across development. More broadly, system-level edge weights offer a natural basis for stratifying individuals and for tracking altered structure–function relationships in neurodevelopmental and psychiatric conditions.

6 Discussion

BGAA models the structure–function relationship through a shared latent graph that supports both SC→FC and FC→SC mapping. Across cohorts, this shared-graph formulation yields more accurate, more reproducible, and more biologically coherent representations than strong graph-based baselines. Several limitations and future directions nevertheless remain. Although the results are stable across the parcellations in Appendix C.2, the main experiments rely on one family of graph-learning and message-passing operators, and the recovered scaffold should also be assessed under related graph encoder architectures. Applying the framework to clinical populations, task settings, and longitudinal data is a natural next step beyond the two resting-state cohorts studied here. Finally, BGAA learns a population-level scaffold, and hierarchical extensions that capture subject-specific deviations around it will further strengthen its biomedical utility for interpretable connectome analysis.

ETHICS STATEMENT

This work analyzes de-identified neuroimaging data from the Human Connectome Project (HCP) and the Adolescent Brain Cognitive Development (ABCD) study, both of which were collected under institutional review board approval with informed consent (and parental permission for ABCD participants). Data were accessed under the respective data use agreements, and no attempt was made to re-identify participants. The learned graphs are population-level summaries and are not intended for individual clinical decision making.

REPRODUCIBILITY STATEMENT

The model and training objective are fully specified in Section 4. Data sources, subject selection, and SC/FC construction are described in Appendix A; data splits, preprocessing, optimizer settings, and all hyperparameters are given in Appendix B.1; and computational resources are reported in Appendix B.4. Ablations, parcellation sensitivity, and hyperparameter sensitivity are reported in Appendix C. Source code will be released upon publication.

AI USE STATEMENT

Generative AI tools were used to assist with converting the manuscript's LaTeX source to the ICLR 2027 template, checking typesetting, and polishing the writing. All AI-assisted content was reviewed by the authors, who take full responsibility for the final content of this work.

REFERENCES

Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *international conference on machine learning*, pp. 21–29. PMLR, 2019.

Carlo Amodeo, Igor Fortel, Olusola Ajilore, Liang Zhan, Alex Leow, and Theja Tulabandhula. Unified embeddings of structural and functional connectome via a function-constrained structural graph variational auto-encoder. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 406–415. Springer, 2022.

James Atwood and Don Towsley. Diffusion-convolutional neural networks. *Advances in neural information processing systems*, 29, 2016.

Graham L Baum, Zaixu Cui, David R Roalf, Rastko Ciric, Richard F Betzel, Bart Larsen, Matthew Cieslak, Philip A Cook, Cedric H Xia, Tyler M Moore, 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.

Qing Cai, Hannah Thomas, Vanessa Hyde, Pedro Luque Laguna, Carolyn B McNabb, Krish D Singh, Derek K Jones, and Eirini Messaritaki. Decoding brain structure-function dynamics in health and in psychosis via an autoencoder. *Scientific Reports*, 15(1):40052, 2025.

B. J. Casey, Tariq Cannonier, May I Conley, Alexandra O Cohen, Deanna M Barch, Mary M Heitzeg, Mary E Soules, Theresa Teslovich, Danielle V Dellarco, Hugh Garavan, et al. The adolescent brain cognitive development (ABCD) study: imaging acquisition across 21 sites. *Developmental cognitive neuroscience*, 32:43–54, 2018.

Shi Yu Chan, Zi Yan Ong, Zhen Ming Ngoh, Yap Seng Chong, Juan H. Zhou, Marielle V. Fortier, Lourdes M. Daniel, Anqi Qiu, Michael J. Meaney, and Ai Peng Tan. Structure-function coupling within the reward network in preschool children predicts executive functioning in later childhood. *Developmental Cognitive Neuroscience*, 55:101107, 2022. doi: 10.1016/j.dcn.2022.101107.

Hejie Cui, Wei Dai, Yanqiao Zhu, Xuan Kan, Antonio Aodong Chen Gu, Joshua Lukemire, Liang Zhan, Lifang He, Ying Guo, and Carl Yang. Braingb: A benchmark for brain network analysis with graph neural networks. *IEEE Transactions on Medical Imaging*, 42(2):493–506, 2023. doi: 10.1109/TMI.2022.3218745.

Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems*, 29, 2016.

Joaquín Goñi, Martijn P van den Heuvel, Andrea Avena-Koenigsberger, Nieves Velez de Mendizabal, Richard F Betzel, Alessandra Griffa, Patric Hagmann, Bernat Corominas-Murtra, Jean-Philippe Thiran, and Olaf Sporns. Resting-brain functional connectivity predicted by analytic measures of network communication. *Proceedings of the National Academy of Sciences*, 111(2):833–838, 2014.

Nia Goulden, Aygul Khusnulina, Nicholas J. Davis, Robert M. Bracewell, Arun L. Bokde, Jonathan P. McNulty, and Paul G. Mullins. The salience network is responsible for switching between the default mode network and the central executive network: replication from dcm. *NeuroImage*, 99: 180–190, 2014. doi: 10.1016/j.neuroimage.2014.05.052.

Mark D Grier, Jan Zimmermann, and Sarah R Heilbronner. Estimating brain connectivity with diffusion-weighted magnetic resonance imaging: promise and peril. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 5(9):846–854, 2020.

Patric Hagmann, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J. Honey, Van J. Wedeen, and Olaf Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7):e159, 2008.

Christopher J Honey, Olaf Sporns, Leila Cammoun, Xavier Gigandet, Jean-Philippe Thiran, Reto Meuli, and Patric Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009.

Keith W. Jamison, Zijin Gu, Qinxin Wang, Ceren Tozlu, Mert R. Sabuncu, and Amy Kuceyeski. Krakencoder: a unified brain connectome translation and fusion tool. *Nature Methods*, 22(7): 1583–1592, 2025. doi: 10.1038/s41592-025-02706-2.

Yang Li, Rasoul Shafipour, Gonzalo Mateos, and Zhengwu Zhang. Mapping brain structural connectivities to functional networks via graph encoder-decoder with interpretable latent embeddings. In *2019 ieee global conference on signal and information processing (globalsip)*, pp. 1–5. IEEE, 2019.

Meimei Liu, Zhengwu Zhang, and David B Dunson. Graph auto-encoding brain networks with applications to analyzing large-scale brain imaging datasets. *Neuroimage*, 245:118750, 2021.

Vinod Menon. 20 years of the default mode network: A review and synthesis. *Neuron*, 111(16): 2469–2487, 2023. doi: 10.1016/j.neuron.2023.04.023.

Nathaniel J Powell, Bettina Hein, Deyue Kong, Jonas Elpelt, Haleigh N Mulholland, Matthias Kaschube, and Gordon B Smith. Common modular architecture across diverse cortical areas in early development. *Proceedings of the National Academy of Sciences*, 121(11):e2313743121, 2024.

Jonathan D Power, Alexander L Cohen, Steven M Nelson, Gagan S Wig, Kelly Anne Barnes, Jessica A Church, Alecia C Vogel, Timothy O Laumann, Fran M Miezin, Bradley L Schlaggar, et al. Functional network organization of the human brain. *Neuron*, 72(4):665–678, 2011.

James M Rowland, Thijs L van der Plas, Matthias Loidolt, Robert M Lees, Joshua Keeling, Jonas Dehning, Thomas Akam, Viola Priesemann, and Adam M Packer. Propagation of activity through the cortical hierarchy and perception are determined by neural variability. *Nature Neuroscience*, 26(9):1584–1594, 2023.

Tabinda Sarwar, Ye Tian, B. T. Thomas Yeo, Kotagiri Ramamohanarao, and Andrew Zalesky. Structure-function coupling in the human connectome: A machine learning approach. *NeuroImage*, 226:117609, 2021. doi: 10.1016/j.neuroimage.2020.117609.

Caio Seguin, Ye Tian, and Andrew Zalesky. Network communication models improve the behavioral and functional predictive utility of the human structural connectome. *Network Neuroscience*, 4(4): 980–1006, 2020. doi: 10.1162/netn_a_00161.

Caio Seguin, Olaf Sporns, and Andrew Zalesky. Brain network communication: concepts, models and applications. *Nature reviews neuroscience*, 24(9):557–574, 2023.

Laura E. Suárez, Ross D. Markello, Richard F. Betzel, and Bratislav Misic. Linking structure and function in macroscale brain networks. *Trends in Cognitive Sciences*, 24(4):302–315, 2020. doi: 10.1016/j.tics.2020.01.008.

Annalisa Tosoni, Paolo Capotosto, Antonello Baldassarre, Sara Spadone, and Carlo Sestieri. Neuroimaging evidence supporting a dual-network architecture for the control of visuospatial attention in the human brain: a mini review. *Frontiers in Human Neuroscience*, 17:1250096, 2023. doi: 10.3389/fnhum.2023.1250096.

Martijn P van den Heuvel and Olaf Sporns. Rich-club organization of the human connectome. *Journal of Neuroscience*, 31(44):15775–15786, 2011. doi: 10.1523/JNEUROSCI.3539-11.2011.

Martijn P van den Heuvel, René S Kahn, Joaquín Goñi, and Olaf Sporns. High-cost, high-capacity backbone for global brain communication. *Proceedings of the National Academy of Sciences*, 109 (28):11372–11377, 2012. doi: 10.1073/pnas.1203593109.

David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, et al. The wu-minn human connectome project: an overview. *Neuroimage*, 80:62–79, 2013.

Bertha Vázquez-Rodríguez, Laura E. Suárez, Ross D. Markello, Golia Shafiei, Casey Paquola, Patric Hagmann, Martijn P. van den Heuvel, Boris C. Bernhardt, R. Nathan Spreng, and Bratislav Misic. Gradients of structure–function tethering across neocortex. *Proceedings of the National Academy of Sciences*, 116(42):21219–21227, 2019. doi: 10.1073/pnas.1903403116.

Jing Xia, Yi Hao Chan, Deepank Girish, and Jagath C. Rajapakse. Interpretable modality-specific and interactive graph convolutional network on brain functional and structural connectomes. *Medical Image Analysis*, 102:103509, 2025.

B. T. Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, Bruce Fischl, Hesheng Liu, and Randy L Buckner. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3):1125–1165, 2011.

Andrew Zalesky, Tabinda Sarwar, Ye Tian, Yuanzhe Liu, B. T. Thomas Yeo, and Kotagiri Ramamohanarao. Predicting an individual's functional connectivity from their structural connectome: Evaluation of evidence, recommendations, and future prospects. *Network Neuroscience*, 8(4): 1291–1309, 2024. doi: 10.1162/netn_a_00400.

Tong Zhou, Chen Ding, Changhong Jing, Feng Liu, Kevin Hung, Hieu Pham, Mufti Mahmud, Zhihan Lyu, Sibo Qiao, Shuqiang Wang, et al. BG-GAN: Generative AI enable representing brain structure-function connections for Alzheimer's disease. *IEEE Transactions on Consumer Electronics*, 71(2):3087–3098, 2025. doi: 10.1109/TCE.2025.3543943.

A Description of ABCD and HCP dataset

A.1 HCP dataset

We apply our model to data from the Human Connectome Project (HCP), a landmark multi-modal study that has collected high-quality neuroimaging together with extensive behavioral and genetic measures. In this work, we use the WU–Minn HCP minimally processed S1200 release, which includes diffusion MRI (dMRI), T1-weighted MRI, and resting-state fMRI for over 1,000 healthy young adults aged 22–37 years. We restrict to subjects with both modalities available and passing quality control, yielding 1,065 individuals for the analyses reported in the main text. Subjects were excluded if they lacked complete dMRI or resting-state fMRI data or were flagged by the HCP quality-control review. The HCP sample contains twins and non-twin siblings, and family structure was accounted for in the analyses.

A.2 ABCD dataset

We also evaluate our model on data from the Adolescent Brain Cognitive Development (ABCD) Study, a large, longitudinal cohort designed to characterize brain and behavioral development in youth. ABCD has acquired multi-modal MRI together with extensive phenotypic and environmental measures across thousands of participants. In the present analyses, we use imaging data from the baseline visit, acquired across 21 sites, for subjects with both diffusion MRI (dMRI) and resting-state fMRI available and passing quality control. Subjects were excluded if they failed modality-specific quality control for T1-weighted, dMRI, or resting-state fMRI data or had missing regions in either connectome. Family structure among siblings was likewise accounted for. After applying modality-specific and connectome-level QC filters and ensuring paired availability, our final sample includes 2,360 individuals for the analyses reported in the main text.

A.3 SC and FC construction

Structural connectivity (SC). For each subject, we construct a structural connectivity matrix from minimally processed dMRI and T1-weighted MRI. We first perform anatomical parcellation using the Desikan–Killiany (DK) atlas with $V = 68$ cortical regions of interest (ROIs). Whole-brain tractography is then generated in the subject's native space, and streamlines are assigned to ROI pairs to form edge-wise tract sets. To improve robustness at the gray–white matter interface and reduce spurious links, we adopt standard tractography post-processing steps, including (i) dilation of each cortical ROI to better capture termination points near white matter, (ii) separating streamlines that traverse multiple ROIs into ROI-to-ROI segments when appropriate, and (iii) removing obvious outlier streamlines based on implausible geometry. For each ROI pair (u, v), we summarize the connection strength using the mean fractional anisotropy (FA) along the retained streamlines connecting the two regions. This yields a symmetric $V \times V$ SC matrix with a zero diagonal. To reduce inter-subject scale variability, we additionally apply subject-wise normalization, and we ensure the final SC matrices are comparable across individuals.

Functional connectivity (FC). We construct functional connectivity from resting-state fMRI using the same DK parcellation for consistency with SC. For each subject, regional time series are extracted by averaging the preprocessed fMRI signal within each ROI. We then compute pairwise Pearson correlations between ROI time series, producing a dense $V \times V$ correlation matrix. To stabilize variance and facilitate learning, correlations are Fisher z-transformed. The resulting FC matrices are symmetrized and set to have a zero diagonal. In this representation, each entry summarizes the degree of synchronized spontaneous fluctuation between two regions during rest, providing a complementary (and typically denser) view of brain-wide organization relative to tractography-based SC.

B Implementation and Computational Details

B.1 Training and Evaluation Details

For HCP, the data were split with random seed 13 into 865 training scans, 100 validation scans, and 100 held-out test scans, corresponding to an 80/10/10 train-validation-test split. The same ratio was

used for ABCD, giving 1,888/236/236 training, validation, and test subjects. Structural connectivity matrices were clipped at zero, symmetrized, diagonal-cleared, transformed by $\log(1+x)$, and z-scored using only training off-diagonal entries. Functional connectivity matrices were symmetrized, diagonal-cleared, and z-scored using only training off-diagonal entries. The graph was parameterized exactly as the deterministic anatomy-informed shared adjacency described in Section 4. Training used stochastic gradient descent with Nesterov momentum. Hyperparameters were selected by validation reconstruction loss from six SGD search trials, each trained for 100 epochs, which varied the learning rate, momentum, and the three penalty weights. The final configuration used learning rate 0.01, momentum 0.95, Nesterov enabled, weight decay 0.0001, graph ℓ_1 penalty $\lambda_1 = 0.0001$, degree penalty $\gamma = 10^{-6}$, residual penalty $\mu = 10^{-5}$, and batch size 64. The final model was then trained for 1000 epochs using this configuration.

B.2 Baseline Implementation

All baselines were trained on the same data splits and preprocessed matrices as BGAA, and their hyperparameters were tuned on the validation set with the same six-trial search budget as BGAA. Because DGNN, MixHop, and BrainGB are designed for a single input graph, each was trained as two direction-specific predictors, SC→FC and FC→SC, using the input connectome as node features and the same decoder families as BGAA. Krak was used with its native bidirectional translation between SC and FC. For TRS and WAF, a population-level $V \times V$ graph was extracted for each baseline from the propagation or fusion structure that the method learns, averaged over the two directions, and was symmetrized, made nonnegative, and assigned a zero diagonal in the same way as $\mathbf{M}$, so that all methods are compared on matrices of the same form.

B.3 Evaluation Metrics and Uncertainty

Bidirectional reconstruction accuracy. For the N_{test} held-out subjects,

$$\mathrm{BRA} = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \Big(\|\widetilde{\mathbf{F}}_i - \mathbf{F}_i\|_F^2 + \|\widetilde{\mathbf{S}}_i - \mathbf{S}_i\|_F^2 \Big),$$

computed on the standardized matrices of Appendix B.1, with BRA_F and BRA_S denoting the two terms.

Test–retest stability. Here test–retest refers to the re-estimation stability of the population graph rather than to repeated scans of the same subjects. Two training replicates were formed by randomly splitting the training set into two disjoint halves, and each method was fitted to both to obtain $\widehat{\mathbf{M}}^{(A)}$ and $\widehat{\mathbf{M}}^{(B)}$. With $\mathrm{vec}_\triangle(\cdot)$ denoting the vector of upper-triangular entries and $E^{(\cdot)}(p)$ the set of the top-$p\%$ edges by weight,

$$\mathrm{pTRS}_\rho = \mathrm{Spearman}\big(\mathrm{vec}_\triangle(\widehat{\mathbf{M}}^{(A)}), \mathrm{vec}_\triangle(\widehat{\mathbf{M}}^{(B)})\big), \qquad \mathrm{pTRS}_J(p) = \frac{|E^{(A)}(p) \cap E^{(B)}(p)|}{|E^{(A)}(p) \cup E^{(B)}(p)|}.$$

Walk-algebra fit. Given a learned graph $\widehat{\mathbf{M}}$,

$$\mathrm{WAF}_S = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \min_{c \in \mathbb{R}^{K+1}} \Big\| \mathbf{S}_i - \sum_{k=0}^{K} c_k \widehat{\mathbf{M}}^k \Big\|_F^2,$$

and WAF_F is defined analogously for $\mathbf{F}_i$. The coefficients are obtained by least squares for each subject, with $K = 3$, and each graph is rescaled to unit spectral norm before its powers are computed so that methods with different weight scales are comparable.

Uncertainty. Values in Table 1 are means with standard deviations over 20 repetitions that differ in random initialization and training replicate. The 95% confidence intervals in Table 2 were obtained by bootstrapping held-out subjects with 1,000 resamples.

Table 2: Ablation results on HCP and ABCD. Lower BRA is better. Diff denotes the increase in BRA relative to full BGAA, with 95% confidence intervals.

	HCP		ABCD	
Variant	**BRA**	**Diff (95% CI)**	**BRA**	**Diff (95% CI)**
BGAA (full)	1.99	–	2.73	–
Fixed $\mathbf{M} = \mathbf{I}$	2.33	0.34 (0.12, 0.56)	3.19	0.46 (0.17,0.75)
w/o shared graph	2.25	0.26 (0.08, 0.44)	3.28	0.55 (0.19,0.91)
w/o attention module	2.11	0.12 (0.03, 0.21)	2.91	0.18 (0.06,0.30)

Table 3: Sensitivity to parcellation choice on HCP and ABCD. Lower BRA is better. Higher pTRS_ρ is better.

		DK-68	Destrieux-148	Schaefer-200	Schaefer-500
HCP	BRA ↓	1.99	2.01	2.10	2.09
	pTRS_ρ ↑	0.53	0.49	0.56	0.53
ABCD	BRA ↓	2.73	2.69	2.81	2.85
	pTRS_ρ ↑	0.41	0.47	0.60	0.55

B.4 Computational Resources

Training was executed on a local CPU-only compute worker using PyTorch 2.5.1+cpu. CUDA was not available, so no GPU acceleration was used. The machine exposed 2 logical CPU cores, and both data loading and model optimization ran in the main PyTorch process on CPU. At the post-run resource check, the system had approximately 1.0 TiB total RAM with about 920 GiB available, so memory capacity was not a limiting factor for the completed ABCD run, although peak memory was not separately profiled. The full BGAA training pipeline, including the 6-trial hyperparameter search and the final 1000-epoch optimization, completed in roughly 38–45 minutes of wall-clock time.

C Additional Experiments

C.1 Ablation Study

Table 2 shows that each major component of BGAA contributes to performance, with the full model achieving the lowest BRA on both HCP and ABCD. Replacing the learned graph with the identity matrix leads to a substantial degradation, indicating that meaningful latent graph structure is essential for cross-modal reconstruction and that performance cannot be explained by local message passing alone. Removing the shared graph constraint also causes a clear loss in accuracy, which supports the central design choice of using one common adjacency to support both SC→FC and FC→SC rather than allowing the two directions to drift toward separate latent structures. By comparison, removing the attention module produces a smaller but still consistent decline, suggesting that anatomy-informed graph construction further improves the quality of the learned scaffold even when the overall bidirectional architecture is retained. Taken together, the ablation results show that the gains of BGAA arise from the combination of a learned latent graph, a shared cross-modal scaffold, and anatomy-guided structure construction.

C.2 Sensitivity to Parcellation Choice

Table 3 summarizes the sensitivity of BGAA to parcellation choice across two anatomical atlases and two functional atlases spanning a range of spatial resolutions. Across both HCP and ABCD, the

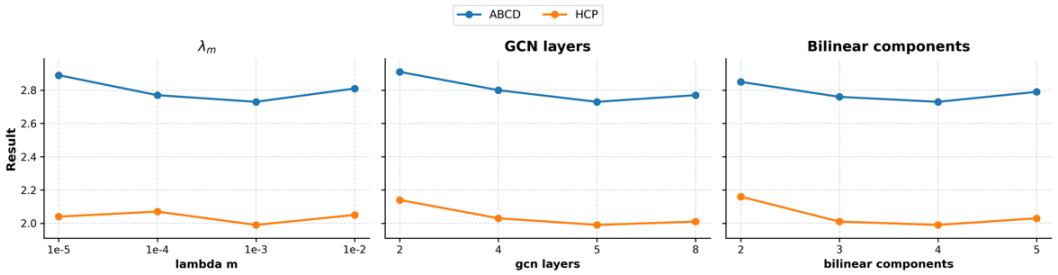


Figure 4: Sensitivity of BGAA bidirectional reconstruction accuracy (BRA) to key hyperparameters on ABCD and HCP. From left to right, the panels vary the graph sparsity penalty λ_1, the number of GCN layers, and the number of bilinear decoder components.

model remains stable in overall performance, with BRA varying only modestly across parcellations and pTRS_ρ remaining consistently moderate to high. This suggests that the shared-graph formulation is not tied to a single cortical partition and can recover comparable large-scale structure–function organization under different region definitions.

A more detailed pattern emerges when comparing predictive accuracy and graph stability. On HCP, the coarser DK-68 parcellation gives the lowest BRA, while Schaefer-200 yields the highest pTRS_ρ, indicating a mild tradeoff between reconstruction simplicity and graph reproducibility at finer spatial scales. On ABCD, Destrieux-148 gives the best BRA, whereas Schaefer-200 again produces the strongest pTRS_ρ, with Schaefer-500 remaining competitive. Taken together, these results suggest that finer functional parcellations can improve the reproducibility of the learned graph, whereas coarser or intermediate parcellations may be slightly more favorable for reconstruction. Importantly, the relative stability across all four parcellations indicates that the main conclusions of BGAA do not depend on a particular atlas choice.

C.3 Sensitivity to Graph and Architecture Hyperparameters

Figure 4 summarizes the sensitivity of BGAA in terms of bidirectional reconstruction accuracy (BRA) to three central hyperparameters that control graph sparsity and decoder capacity. Across both ABCD and HCP, BRA remains relatively stable over the tested ranges, with the best results consistently occurring in the interior rather than at the most extreme settings. In particular, moderate graph regularization, a mid-range number of GCN layers, and an intermediate number of bilinear components all lead to the strongest reconstruction accuracy. This pattern suggests that BGAA is not highly brittle to small changes in these design choices and that its performance is driven by a broad region of reasonable hyperparameter settings rather than a single finely tuned configuration.

A more detailed view shows a consistent bias toward moderate model complexity. Very small or very large values of λ_1 slightly worsen BRA, indicating that both under-regularized and over-regularized graphs are less effective than a balanced sparse scaffold. Increasing the number of GCN layers improves BRA up to a point, after which the gains saturate or reverse, consistent with the idea that multi-step propagation is beneficial but excessive depth can dilute useful structure. A similar trend appears for the number of bilinear components, where intermediate values give the best BRA and additional components provide little benefit. Taken together, these results support the robustness of BGAA and suggest that its empirical gains do not depend on unusually sensitive hyperparameter choices.

D Biological Validation

To quantitatively assess the biological validity of the learned graph, we compared it with two well-established integrative structures in the human connectome: rich-club connections and structural core connections. For each cohort, graph edges were ranked by weight and the top 5% were selected to form the set E_{top}. The rich-club edge set E_{RC} consists of all connections whose two endpoints are cortical rich-club nodes, defined as the cortical rich-club regions reported by van den Heuvel &

Table 4: Overlap between the learned graph and established integrative structures.

Cohort	Rich-club overlap	Core overlap
HCP	0.87	0.71
ABCD	0.67	0.64

Sporns (2011), including precuneus, superior frontal, and superior parietal cortex, while the structural core edge set E_{core} consists of all connections whose two endpoints are structural core regions identified by Hagmann et al. (2008). Both reference sets are defined from structural hub properties rather than from the functional-system labels used as node features, so this comparison does not depend on the model inputs. We then quantified overlap as the fraction of known integrative edges captured by the graph,

$$\text{Overlap}_{\text{RC}} = \frac{|E_{\text{top}} \cap E_{\text{RC}}|}{|E_{\text{RC}}|}, \qquad \text{Overlap}_{\text{core}} = \frac{|E_{\text{top}} \cap E_{\text{core}}|}{|E_{\text{core}}|}.$$

To assess significance, we generated a null distribution by randomly sampling $|E_{\text{top}}|$ edges 1000 times with matched degree distribution and recomputing the same overlap. In both cohorts, the observed overlap significantly exceeded the null expectation ($p < 0.001$).

Table 4 shows strong overlap between the learned graph and both rich-club and structural core connections in HCP and ABCD. The highest-weight graph edges concentrate on canonical rich-club hubs such as precuneus, superior frontal, and superior parietal cortex, as well as structural core regions including posterior cingulate and isthmus cingulate. These patterns are consistent with prior accounts of long-range integrative relay architecture and the posterior medial scaffold that supports default-mode coordination, providing quantitative support that BGAA recovers a biologically meaningful communication scaffold rather than an arbitrary latent graph.

E Details of the Cognitive-Group Analysis

Cognitive performance was measured by the NIH Toolbox Cognition Total Composite score in both cohorts. To reduce confounding, the scores were adjusted for age, sex, acquisition site, and head motion. Subjects were split at the cohort-specific median of the adjusted score, giving 532/533 high- and low-cognition subjects in HCP and 1,180/1,180 in ABCD, and BGAA was fitted separately to each group with the configuration in Appendix B.1. Because both group graphs are fitted with identical node features, any structure contributed by the features is common to the two groups and cancels in the high − low difference.

For each of the 28 pairs formed by the seven functional systems, we computed the mean edge weight in each group and took the difference (high − low). To separate system-specific effects from a global change in edge weight, each difference was centred by subtracting the cohort-wide mean change across all edges. Significance was assessed against a null distribution from 5,000 random reassignments of system labels to regions, which keeps both graphs and the number of regions per system fixed, and p-values were corrected with the Benjamini–Hochberg procedure across the 28 pairs within each cohort. In Figure 3a, pairs are ordered by their average standardized change across the two cohorts, a single asterisk marks $q < 0.05$ in one cohort, and a double asterisk marks $q < 0.05$ in both cohorts with the same sign.

The global properties in Figure 3b are computed from the graph $\mathbf{M}$ with node strength $s_u = \sum_v M_{uv}$. Mean edge weight and edge weight SD are taken over the upper-triangular entries. Cross-system integration is the mean participation coefficient of the regions with respect to the seven systems. Top 10% edge concentration is the share of total weight carried by the strongest 10% of edges, and top 10 node concentration is the share of total strength held by the 10 strongest regions. Node-strength CV is the coefficient of variation of s_u across regions, within-system fraction is the share of total weight on edges whose two endpoints belong to the same system, and network entropy is the Shannon entropy of the normalized edge-weight distribution divided by its maximum value. In Figure 3c,d, the regional change is the difference in $s_u / \sum_v s_v$ between groups, expressed in percentage points and displayed on the inflated fsaverage surface with the medial wall shown in grey.

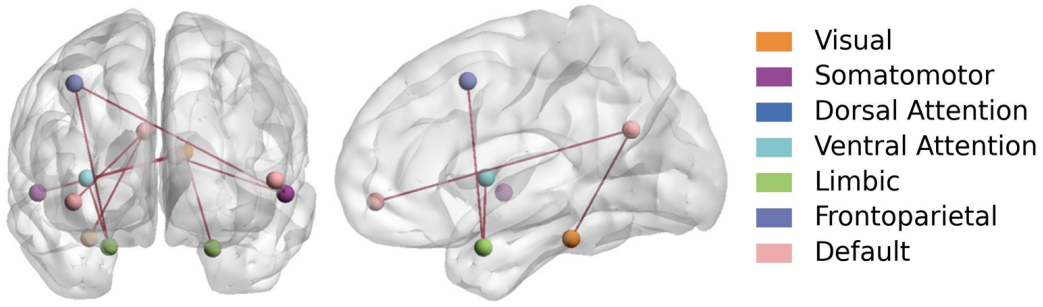


Figure 5: HCP top graph edges. Visualization of the 10 highest-weight edges in the learned graph for the HCP cohort, shown on a cortical surface rendering. Nodes are colored by large-scale functional system (legend), and lines indicate the strongest graph connections linking regions across and within systems.

F SUPPLEMENTARY REAL DATA ANALYSIS

Figures 5 and 6 visualize the ten highest-weight edges in the learned graph for HCP and ABCD, respectively. Because these edges represent the most salient entries of the shared scaffold, they provide an interpretable summary of which region-to-region channels the model consistently prioritizes when linking structural and functional organization. Across both cohorts, the top edges are predominantly long-range and cross-lobar, rather than being restricted to short-range neighboring regions. This pattern is consistent with the view that the most informative cross-modal scaffold is dominated by a compact set of high-capacity connections that support distributed coordination across the cortex, rather than reflecting only local adjacency or dense within-lobe wiring.

A second consistent feature is the system-level composition of these edges. In both figures, the top connections frequently involve association systems (Default Mode, Ventral Attention, and Frontoparietal) and their interactions with visual territories, indicating that the strongest graph channels concentrate on networks that are commonly implicated in integration, control, and flexible coordination. Interpreted biologically, this suggests that the inferred scaffold emphasizes edges that are well-positioned to link internally coherent systems with higher-order association circuits that support context-dependent coordination. In other words, the dominant graph edges appear to bridge functional specializations across systems, highlighting a small number of channels that plausibly sustain large-scale coordination.

Finally, comparing Figures 5 and 6 suggests cross-cohort robustness at the level of *edge types* even when the exact anatomical endpoints vary. Both cohorts highlight a top-edge set enriched for association-network participation and inter-system links, implying that the learned graph is not driven by cohort-specific noise or idiosyncratic preprocessing artifacts. The agreement across a young adult cohort (HCP) and a developmental cohort (ABCD) supports the interpretation that the graph is capturing conserved organizational principles of the structure–function scaffold.

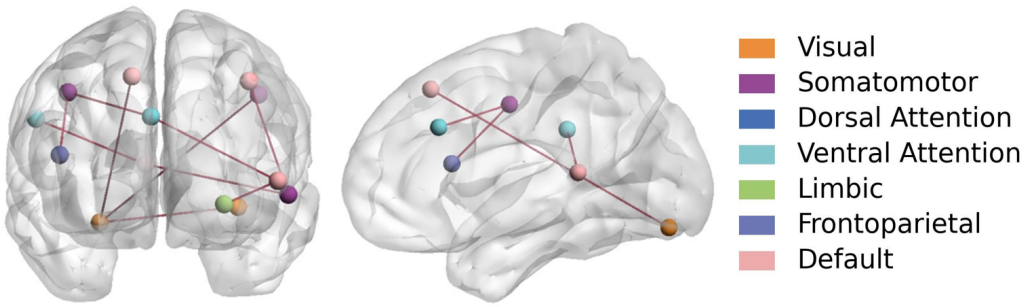


Figure 6: ABCD top graph edges. Visualization of the 10 highest-weight edges in the learned graph for the ABCD cohort, shown on a cortical surface rendering. Nodes are colored by large-scale functional system (legend), and lines indicate the strongest graph connections linking regions across and within systems.