SPEF: SEMANTIC PRIOR-SELECTED EXPERT FUSION FOR WORKING-MEMORY PREDICTION FROM MULTIMODAL BRAIN CONNECTOMES

Anonymous authors
Paper under double-blind review

ABSTRACT

Predicting cognition from functional and structural connectomes requires learning from thousands of correlated edges from comparatively modest neuroimaging cohorts. In order to stabilize the prediction framework, we introduce *Semantic Prior-selected Expert Fusion (SPEF)*, a transparent framework that converts cognitive task descriptions and atlas region names into a frozen semantic atlas prior, selects compact functional-connectivity (FC) and structural-connectivity (SC) expert subspaces, and hierarchically fuses the experts in a strong data-driven connectome backbone. The prior is generated independently of participant imaging and behavioral data, validated against the canonical AAL116 ordering, normalized and frozen, and then transferred to connectivity space to define compact FC and SC subspaces. On 510 Human Connectome Project participants, five-seed, five-fold repeated nested cross-validation gives Pearson $r = 0.2439$, RMSE 11.293, and MAE 9.130 for SPEF-ridge on a working memory task. Among the evaluated multimodal comparators and SPEF variants, the best Pearson correlation is 0.2439, attained by both matched SPEF-ridge and cross-task SPEF-NC-ridge, while the latter achieves the lowest RMSE (11.288) and MAE (9.121). The working-memory semantic prior also exceeds shuffled and random prior controls by $+0.0127$ and $+0.0148$ mean seed-wise correlation, respectively, while the matched and cross-task semantic priors perform similarly. Prediction-linked coefficient maps show reproducible FC and SC structure. Relative to an architecture-matched no-prior full-connectome ridge expert, SPEF increases FC edge-rank stability (0.723 vs. 0.544) and yields a substantially less adverse top-ROI perturbation contrast (-0.0017 vs. -0.0864), while other stability measures remain mixed. SPEF provides an auditable framework for introducing language-derived cognitive knowledge as an anatomically grounded inductive bias for multimodal connectome based prediction and candidate biomarker analyses.

1 INTRODUCTION

Human cognition emerges from coordinated neural activity constrained by anatomical pathways. Functional connectivity (FC) summarizes statistical dependencies among regional time series whereas structural connectivity (SC) describes the white-matter substrate that supports communication among regions (Sporns et al., 2005; Bullmore & Sporns, 2009; Honey et al., 2009). Their relationship varies across the cortex and carries behaviorally relevant information (Preti & Van De Ville, 2019). This complementarity motivates multimodal prediction models that integrate functional and structural connectomes so that functional organization and anatomical support contribute jointly to behavioral prediction.

A central challenge in this prediction task is dimensionality. With an AAL116 parcellation, each symmetric connectome contains 6,670 unique edges, already exceeding the sample size of many neuroimaging cohorts. Connectome-based predictive modeling, graph neural networks, and multimodal fusion extract useful signal from these representations (Finn et al., 2015; Shen et al., 2017; Dsouza et al., 2021), yet the learning remains ill-posed due to the sparsity of imaging data. Biological or task-specific prior knowledge offers a natural route to stabilizing learning. The key challenge is to determine how to introduce such knowledge while preserving strict train/test separation, retaining

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.

a competitive data-driven predictor, and exposing the effect of prior identity through controlled comparisons.

Recent neuroimaging work has increasingly emphasized biologically guided multimodal learning and interpretable structure–function analysis. Recent works have studied task-specific functional architectures (Gupta et al., 2022), joint neuroimaging and omics integration (Chan et al., 2022), prior-informed graph learning (Xiao et al., 2023), multimodal brain graphs (Zhang et al., 2023), modality-specific FC/SC interaction models (Xia et al., 2025b), and modularity-guided structure–function interaction learning (Xia et al., 2026). In parallel, large-scale meta-analysis and language-to-brain models demonstrate that textual descriptions of cognitive processes contain structured neuroanatomical information (Yarkoni et al., 2011; Salo et al., 2023; Ngo et al., 2022). These developments motivate a complementary use of language models: a large language model provides a frozen, external hypothesis over atlas regions that control the connectome features available to a compact expert while participant-level learning remains fully data driven. In this work, we explore how publicly available knowledge bases are incorporated into cognitive task prediction combining structural and functional connectomes.

We propose **Semantic Prior-selected Expert Fusion (SPEF)**, a framework that converts cognitive task descriptions and atlas region names into a frozen prior for task prediction. Rather than asking the language model to operate on participant data or make behavioral predictions, SPEF restricts its role to prior construction. The resulting regional scores are validated against the canonical AAL116 ordering, normalized, frozen, and propagated to connectivity space. They then define candidate FC and SC subspaces in which compact regression experts are estimated. The modality-specific experts are combined by a convex fusion stage and subsequently integrated with a data-driven FC–SC backbone. Matched-task, cross-task, shuffled, and random priors can therefore be exchanged while retaining the same downstream prediction architecture.

The study makes four contributions:

1. We construct a reproducible task-to-atlas semantic-prior, which converts natural-languaged brain knowledge into a frozen, versioned AAL116 relevance vector before participant-level model fitting.
2. We translate regional semantic relevance into connectivity-space subspaces and estimate FC and SC experts by using ridge and network-controlled ridge (NC-ridge), providing complementary isotropic and network-structured estimators within the same prior-guided feature spaces.
3. We introduce two-stage multimodal fusion that first combines FC and SC expert predictions and then combines the semantic expert with a data-driven FC–SC backbone.
4. We evaluate semantic-prior identity through matched-task, cross-task, shuffled, and random controls, compare SPEF with representative multimodal connectome predictors, and link prediction to coefficient stability and perturbation-based candidate-biomarker analysis.

We evaluate SPEF on working-memory prediction on 510 HCP participants with five-seed, five-fold repeated nested cross-validation. The resulting design treats semantic knowledge as a testable modeling component: prior identity can change while the estimator, folds, feature construction, and tuning protocol remain fixed.

2 Related Work

Connectome prediction and multimodal fusion. Individual behavior can be predicted from distributed FC patterns, motivating connectome fingerprinting and connectome-based predictive modeling (Finn et al., 2015; Shen et al., 2017). Multimodal graph models extend this idea by combining FC and SC. M-GCN integrates functional and structural connectomics for phenotypic prediction (Dsouza et al., 2021). IMG-GCN introduces modularity-guided structure–function interactions for cognitive prediction (Xia et al., 2024), while MS-Inter-GCN separates modality-specific information from learned FC–SC interactions (Xia et al., 2025b). Rajapakse et. al. have developed several related multimodal formulations: graph-based integration of neuroimaging and omics (Chan et al., 2022), late-fusion multimodal graph learning for neurodegeneration (Zhang et al., 2023), modality-specific and interactive FC–SC graph convolution (Xia et al., 2025b), and modularity-guided attention for

cross-ROI structure–function interactions (Xia et al., 2026). Multimodal stacking has also been studied as a way to improve reliability of brain-wide cognitive associations (Tetereva et al., 2025). SPEF focuses on a complementary regime: compact regression experts whose feature spaces are selected by an external semantic prior and integrated with a strong multimodal backbone through learned convex fusion.

Task semantics and prior knowledge. Meta-analytic resources such as Neurosynth and NiMARE aggregate associations between cognitive terms and neuroimaging coordinates (Yarkoni et al., 2011; Salo et al., 2023). Text2Brain further showed that free-form text can be translated into brain activation maps (Ngo et al., 2022). Within Rajapakse's work, task-specific versus task-general functional architectures have been decoded from HCP task fMRI (Gupta et al., 2022), graph neural networks with multiple biological priors have been investigated for multi-omics learning (Xiao et al., 2023), and meta-analysis reference maps have guided a multi-task graph transformer for neurological disease and cognitive-deficit prediction (Xia et al., 2025c). SPEF uses semantic information in a deliberately restricted role: the language model produces a versioned ROI score vector before model fitting, and that vector defines candidate FC/SC subspaces. Architecture-matched prior swaps then separate the contribution of structured semantics from the generic effects of feature restriction.

Biological guidance across data modalities. A broader line of Rajapakse's work uses domain structure to make multimodal representations more biologically organized. Multiple-prior graph learning encodes complementary omics relations (Xiao et al., 2023); modality-specific interaction models separate information unique to FC and SC from their coupling (Xia et al., 2025b); modularity-guided attention uses network organization to guide cross-ROI structure–function interactions (Xia et al., 2026); and latent bottleneck attention integrates transcriptomic, proteomic, and imaging information in neurodegeneration (Niwarthana & Rajapakse, 2026). SPEF follows the same general principle of explicit inductive structure while locating the prior outside the participant data: semantic knowledge defines candidate edge subspaces, and nested training estimates their predictive contribution. This placement gives the prior a direct experimental interpretation because matched, cross-task, shuffled, and random score vectors can be exchanged under the same statistical architecture.

Interpretable biomarkers. Interpretability in neuroimaging benefits from reproducibility alongside anatomical salience. Site-invariant biomarker analysis (Chan et al., 2023), multi-task identification of shared and phenotype-specific functional changes (Xia et al., 2025a), and dynamic-FC biomarker discovery (Girish et al., 2026) illustrate the importance of reproducibility across resampling, sites, or related outcomes. A recent systematic review led by Rajapakse emphasizes robustness as a central criterion for machine-learning-derived neuroimaging biomarkers (Chan et al., 2025). Following this principle, SPEF reconstructs prediction-linked primal coefficients, measures cross-fold stability, and complements stability with perturbation analysis on the corresponding outer-test folds.

3 METHODS

3.1 DESIGN PRINCIPLE

SPEF separates three components that are often entangled in prior-guided prediction: *semantic-prior construction*, *participant-level estimation*, and *prediction fusion*. Semantic knowledge is first converted into a frozen atlas-level score vector independently of participant imaging and behavioral data. The prior then restricts the FC and SC feature spaces in which ridge or network-controlled ridge experts are estimated. Finally, the modality-specific expert predictions are fused and combined with a data-driven FC–SC backbone. The procedures used to select subspaces, regularization parameters, and fusion weights are part of the training and validation protocol and are described in Sec. 4.2.

3.2 TASK-SEMANTIC ATLAS PRIOR CONSTRUCTION

For subject s, let $x_s^F, x_s^S \in \mathbb{R}^E$ denote vectorized FC and SC features derived under an atlas $\mathcal{A} = \{a_i\}_{i=1}^N$, and let y_s denote the behavioral target. Given a cognitive-task description τ and the atlas labels, SPEF defines a task-to-atlas semantic mapping

$$g(\tau, \mathcal{A}) \longrightarrow r = (r_1, \ldots, r_N), \tag{1}$$

where r_i represents the semantic relevance assigned to region a_i for the target cognitive process. The mapping is constructed before participant-level imaging features, behavioral outcomes, model coefficients, or evaluation partitions are used. The semantic prior therefore represents an external anatomical hypothesis rather than a participant-adaptive predictor.

Regional semantic scoring. A frozen language model receives the cognitive-task description together with the ordered atlas labels and returns one continuous relevance score for each region. The output is accepted only when every atlas label is represented exactly once and can be matched to the canonical atlas ordering. The specific language model and generation settings used in the experiments are given in Sec. 4.2.

Normalization and freezing. Let r_i be the validated raw score. SPEF forms the normalized prior

$$p_i = \frac{r_i - \min_j r_j}{\max_j r_j - \min_j r_j}, \qquad p_i \in [0, 1]. \tag{2}$$

The resulting vector $p = (p_1, \ldots, p_N)$ is stored together with its atlas indices, labels, task specification, model-generation metadata, and SHA256 checksum. The instantiated working-memory prior and its complete regional distribution are visualized in the supplementary material.

From regional to connectivity semantics. Because SPEF operates on connectome edges, regional relevance is transferred to an undirected connection (i, j) through

$$q_{ij} = p_i p_j, \qquad i < j. \tag{3}$$

An edge therefore receives a high semantic score when both incident regions receive high task relevance. The resulting regional and edge scores define the prior-guided candidate subspaces. Coherent cross-task, shuffled, and random priors can be substituted without changing the downstream estimator, enabling controlled evaluation of semantic organization. Their experimental instantiation is described in Sec. 4.2.

3.3 Prior-guided subspaces and regression experts

For each modality, SPEF considers two families of candidate subspaces. A *direct top-K* subspace retains the K highest-scoring edges under q. An ***ROI-incident*** subspace first selects the M highest-scoring regions under p and retains every edge incident to at least one selected region. FC and SC subspaces are selected independently using training data only. The candidate values of K and M used in the experiments are specified in Sec. 4.2.

Given a selected subspace $\mathcal{M}_m$ for modality $m \in \{F, S\}$, SPEF considers two regression estimators. The first is ordinary ridge regression,

$$\hat{\beta}_m = \arg\min_{\beta} \|y - X_{m,\mathcal{M}_m}\beta\|_2^2 + \lambda_m \|\beta\|_2^2, \tag{4}$$

where $\lambda_m \geq 0$ controls isotropic regularization.

The second estimator is *network-controlled ridge (NC-ridge)*, which augments the isotropic ridge penalty with a structured smoothness term. A prior-derived line-graph Laplacian L_p relates connectivity features sharing high-prior anatomical regions,

$$\hat{\beta}_m = \arg\min_{\beta} \|y - X_{m,\mathcal{M}_m}\beta\|_2^2 + \lambda_m \|\beta\|_2^2 + \rho\lambda_m \beta^\top L_p \beta, \tag{5}$$

where $\rho \geq 0$ controls the strength of network structure relative to ordinary ridge; $\rho = 0$ recovers ridge. The construction and candidate regularization values used in the experiments are reported in Sec. 4.2.

3.4 Expert and hierarchical fusion

The modality-specific predictions are first combined into a semantic expert; v controls the relative contribution of FC and SC.

$$\hat{y}_E = v\hat{y}_F + (1 - v)\hat{y}_S, \qquad 0 \leq v \leq 1, \tag{6}$$

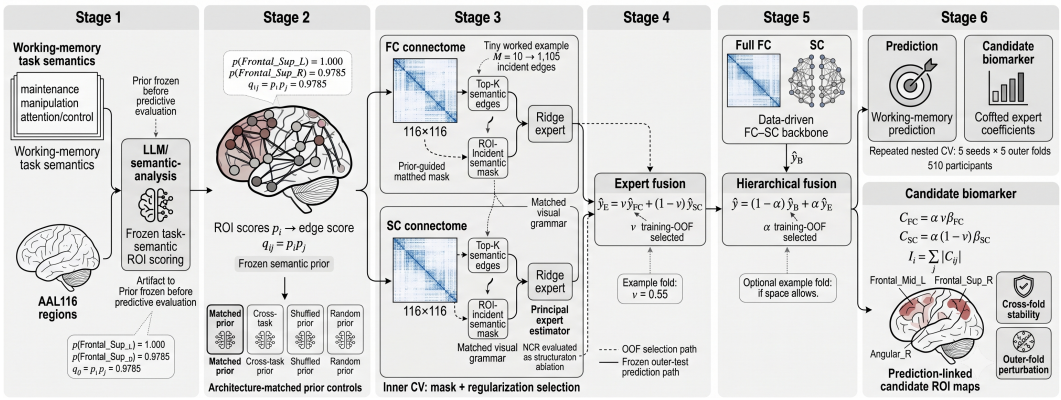


Figure 1: SPEF workflow. A frozen task-semantic atlas prior defines compact FC and SC subspaces in which ridge or NC-ridge experts are estimated. Modality-specific predictions are combined into a semantic expert and subsequently fused with a data-driven FC–SC backbone. Prediction-linked primal coefficients provide candidate anatomical maps. Training and validation procedures used to estimate subspaces, regression parameters, and fusion weights are described in Sec. 4.2.

The semantic expert is subsequently combined with the data-driven multimodal backbone,

$$\hat{y} = (1 - \alpha)\hat{y}_B + \alpha\hat{y}_E, \qquad 0 \leq \alpha \leq 1, \tag{7}$$

where $\hat{y}_B$ denotes the backbone prediction and α controls the contribution of the semantic expert. Estimation of v and α is a training procedure and is therefore described together with the repeated nested evaluation protocol in Sec. 4.2.

3.5 Architecture-matched semantic controls and candidate biomarkers

We evaluate four priors under identical features, folds, candidate masks, hyperparameter grids, preprocessing, and fusion rules: (i) the matched working-memory semantic prior; (ii) a cross-task fluid-intelligence semantic prior; (iii) a deterministic shuffle of the working-memory scores across ROI labels; and (iv) a frozen random ROI prior. Ridge and NC-ridge each receive the same four prior identities. This design quantifies the contribution of structured semantic organization while controlling feature-count restriction through priors with matched dimensionality and identical downstream learning.

For a fitted split with expert weight $\alpha > 0$, the expert's primal coefficients are mapped back to the complete E-dimensional connectivity space. Prediction-linked modality maps are

$$c_F = \alpha v \beta_F, \qquad c_S = \alpha(1 - v)\beta_S. \tag{8}$$

ROI importance sums absolute incident coefficients,

$$I_i = \sum_{j \neq i} |c_F(i, j)| + \sum_{j \neq i} |c_S(i, j)|. \tag{9}$$

Here I_i is a coefficient-derived importance magnitude rather than a probability; normalization is used only for visualization. FC and SC stability are evaluated separately after excluding folds in which the corresponding modality is inactive; folds with $\alpha = 0$ are excluded from coefficient-map comparisons. We evaluate cross-fold edge-rank correlation, top-k ROI overlap, and coefficient-sign consistency. Perturbation faithfulness replaces edges incident to selected ROIs by their outer-training means, leaves model parameters frozen, and measures the resulting change in outer-test RMSE relative to random ROI masks. To directly evaluate the effect of semantic subspace restriction on candidate-biomarker discovery, we additionally construct a no-prior full-connectome ridge expert. It uses the complete FC and SC feature spaces but otherwise shares the same outer folds, ridge grid, modality fusion, hierarchical fusion, coefficient reconstruction, stability definitions, and perturbation protocol as SPEF-ridge. The comparison therefore isolates the effect of replacing full-connectome estimation with prior-guided feature spaces.

4 EXPERIMENTAL DESIGN

4.1 COHORT, CONNECTOMES, AND TARGET

We study 510 Human Connectome Project Young Adult participants (Van Essen et al., 2012) with complete FC, SC, and NIH Toolbox List Sorting working-memory scores (`ListSort_Unadj`). Resting-state fMRI follows HCP minimal preprocessing (Glasser et al., 2013; Smith et al., 2013). Regional time series are extracted with AAL116, FC is computed as Pearson correlation, Fisher-z transformed, symmetrized, and vectorized over the 6,670 upper-triangular edges. SC is reconstructed from diffusion MRI using FSL/MRtrix3 tools (Sotiropoulos et al., 2013; Jenkinson et al., 2012; Tournier et al., 2019), anatomically constrained tractography and SIFT2 weighting (Smith et al., 2015), followed by AAL116 connectome construction. SC weights undergo $\log(1+x)$ transformation, subject-wise maximum scaling, symmetrization, and upper-triangle vectorization. The supplement provides the executable preprocessing sequence and release-script mapping.

4.2 REPEATED NESTED EVALUATION AND FUSION-WEIGHT ESTIMATION

Experimental settings. The experiments use the AAL116 atlas ($N = 116$, $E = 6{,}670$) and NIH List Sorting as the working-memory target. The working-memory semantic prior is generated with Qwen3.8-27B using temperature 0.2 and seed 42; fluid intelligence provides the coherent cross-task prior. Candidate direct subspaces use $K \in \{100, 300, 600, 1200\}$ and ROI-incident subspaces use $M \in \{5, 10, 15\}$. Ridge uses $\lambda \in \{10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 100, 1000\}$. For NC-ridge, L_p links edges associated with the ten highest-prior regions and $\rho \in \{0, 0.1, 0.3, 1\}$. Fusion weights use $v, \alpha \in \{0, 0.05, \ldots, 1\}$.

The final protocol uses five outer seeds and five outer folds per seed. Within every outer-training partition, three-fold inner cross-validation selects the FC and SC subspaces and their ridge or NC-ridge regularization parameters. All fold-local standardization is fitted exclusively on the corresponding training data.

After subspace and estimator selection, three-fold cross-fitting within the outer-training partition generates out-of-fold predictions from the FC and SC experts. These training-only predictions select $v \in \{0, 0.05, \ldots, 1\}$ in Eq. 6. The data-driven FC–SC backbone independently produces cross-fitted training predictions, and the paired backbone/expert predictions select $\alpha \in \{0, 0.05, \ldots, 1\}$ in Eq. 7. The selected experts are then refitted on the complete outer-training partition and the frozen v and α are applied once to the outer-test fold.

The same 25 outer partitions are reused for every prior identity, both regression estimators, and all comparison methods. The primary estimand is the mean of the five seed-wise outer-CV Pearson correlations; seed SD, RMSE, and MAE are reported alongside it. No outer-test outcome contributes to feature selection, hyperparameter selection, early stopping, or fusion-weight estimation.

Illustrative outer fold. To illustrate the complete training and evaluation sequence, consider seed 7171, fold 0. The working-memory prior assigns Frontal_Sup_L $p_i = 1.000$ and Frontal_Sup_R $p_j = 0.9785$, giving $q_{ij} = 0.9785$ for their connection. Within the outer-training partition, inner cross-validation selects an FC ROI-incident subspace with $M = 10$ (1,105 edges) and an SC direct top-K subspace with $K = 1200$. The fitted ridge experts use modality weight $v = 0.55$, while hierarchical fusion selects $\alpha = 0.35$. Applying these frozen choices to the outer-test fold gives $r = 0.2924$. This example traces one prior score through subspace selection, expert estimation, fusion-weight selection, and final outer-fold prediction; additional worked examples are provided in the supplement.

4.3 MULTIMODAL COMPARISON METHODS

SPEF is compared with representative multimodal connectome predictors under the identical 510-participant outer partitions. The comparison includes a concatenated FC+SC ridge baseline, M-GCN (Dsouza et al., 2021), IMG-GCN (Xia et al., 2024), and MS-Inter-GCN (Xia et al., 2025b), together with the data-driven FC–SC backbone. Graph models use their validated architectures and training-only validation for checkpoint selection; each outer-test fold is evaluated once. This design compares

Table 1: Working-memory prediction under identical 510-participant repeated nested-CV outer folds. For each SPEF estimator family, the strongest observed prior configuration is reported; the complete semantic-prior analysis is given in Table 2. Pearson values are mean ± seed SD. Bold denotes the best displayed value in each column; ties at the reported Pearson precision are all bolded.

Method	Pearson $r \uparrow$	RMSE ↓	MAE ↓
FC+SC ridge (Hoerl & Kennard, 1970)	0.2081 ± 0.0214	11.911	9.504
M-GCN (Dsouza et al., 2021)	0.0350 ± 0.0396	11.563	9.370
IMG-GCN (Xia et al., 2024)	0.0973 ± 0.0188	11.578	9.367
MS-Inter-GCN (Xia et al., 2025b)	0.0712 ± 0.0294	11.553	9.355
Data-driven FC–SC backbone	0.2430 ± 0.0248	11.403	9.195
SPEF-ridge / working memory	$\mathbf{0.2439} \pm 0.0129$	11.293	9.130
SPEF-NC-ridge / cross-task	$\mathbf{0.2439} \pm 0.0191$	**11.288**	**9.121**

methods on the same participants, target, FC/SC inputs, and test partitions rather than importing performance values obtained from different cohorts or cognitive targets.

4.4 Control logic and reporting estimand

The four prior identities form a progressively stronger diagnostic of what the semantic scores contribute. The matched working-memory prior encodes the target-task hypothesis. The cross-task prior preserves a coherent cognitive-neuroanatomical organization while changing the task description to fluid intelligence. The shuffled prior preserves the empirical set of working-memory scores while reassigning them across ROI labels, and the random prior supplies a frozen atlas-sized reference vector. Because every downstream choice is shared, differences among these conditions trace prior organization rather than estimator capacity. We report the mean seed-wise Pearson correlation as the primary repeated-CV estimand, accompanied by seed SD, RMSE, MAE, and the number of seeds with a positive paired difference. Full fold-wise values, selected configurations, and the prediction-ensemble sensitivity estimand are provided in the supplement.

5 Results

5.1 Comparison with multimodal connectome predictors

Table 1 compares SPEF with representative multimodal connectome predictors under identical outer partitions. The matched SPEF-ridge model and the cross-task SPEF-NC-ridge model both attain Pearson $r = 0.2439$ at the reported precision. SPEF-NC-ridge with the cross-task prior obtains the lowest RMSE (11.288) and MAE (9.121), while the matched SPEF-ridge formulation attains RMSE 11.293 and MAE 9.130. Both SPEF estimator families outperform the evaluated multimodal graph comparators on the present working-memory task.

5.2 Semantic-prior controls

Table 2 reports the principal prediction results. The data-driven backbone reaches $r = 0.2430 \pm 0.0248$. SPEF-ridge with the matched working-memory prior reaches $r = 0.2439 \pm 0.0129$, RMSE 11.293, and MAE 9.130. The mean correlation difference relative to the backbone is $+0.0010$, with positive paired differences in three of five outer-CV seeds. The cross-task semantic prior reaches a nearly identical $r = 0.2439$, indicating that both structured cognitive priors define useful expert subspaces at this sample size.

The architecture-matched shuffled and random controls separate semantic organization from generic restriction. Within ridge, the matched prior exceeds the shuffled control by $+0.0127$ mean seed-wise r (4/5 seeds) and the random control by $+0.0148$ (4/5 seeds); the matched-minus-cross difference is $+0.00002$. NC-ridge exhibits the same ordering against shuffled and random controls, while its cross-task semantic prior attains the strongest RMSE and MAE in the table. Together, these comparisons concentrate the clearest effect on *structured semantic priors versus shuffled or random organization*, while task-specific and cross-task semantic structure remain closely matched. Full control results are reported in the supplement.

Table 2: Architecture-matched semantic-prior controls and regression-estimator variants. Pearson values are mean ± seed SD over five outer-CV seeds.

Method / prior	Pearson $r \uparrow$	RMSE $\downarrow$	MAE $\downarrow$
Data-driven FC–SC backbone	0.2430 ± 0.0248	11.403	9.195
SPEF-ridge / working memory	0.2439 ± 0.0129	11.293	9.130
SPEF-ridge / cross-task	0.2439 ± 0.0214	11.292	9.132
SPEF-NC-ridge / working memory	0.2386 ± 0.0180	11.338	9.155
SPEF-NC-ridge / cross-task	0.2439 ± 0.0191	11.288	9.121

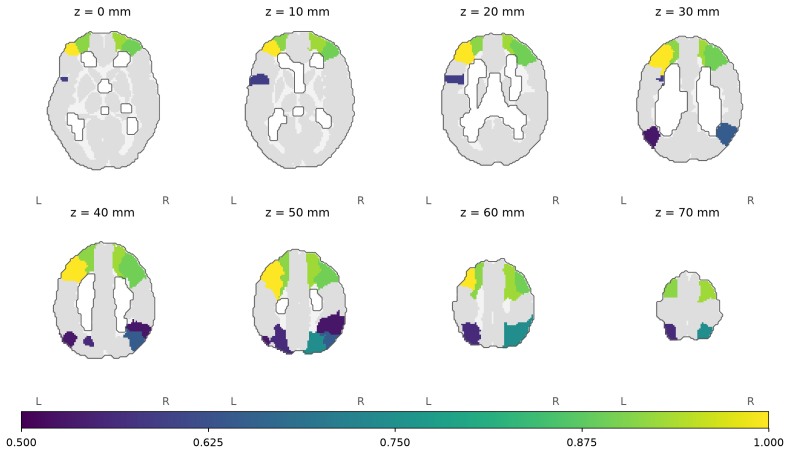


Figure 2: Spatial distribution of prediction-linked candidate biomarkers identified by matched SPEF-ridge. Multimodal ROI importance is derived from the fitted FC and SC coefficient maps, $I_i = I_F(i) + I_S(i)$, and averaged across valid outer folds. For visualization, importance is normalized by the maximum ROI importance, and only regions with normalized importance ≥ 0.5 are shown in color; remaining regions are gray. The normalized importance is a coefficient-derived magnitude, not a probability.

5.3 What the expert contributes

The prediction ablations clarify the hierarchy. The matched-prior FC ridge expert alone reaches mean seed-wise $r = 0.2448$, while the SC expert alone reaches $r = 0.0412$. Direct FC–SC expert fusion reaches $r = 0.2003$, and the final hierarchical fusion reaches $r = 0.2439$ by combining expert and backbone predictions through training OOF estimates. Across the 25 outer folds, the matched ridge expert receives mean FC modality weight $v = 0.714$ and mean hierarchical weight $\alpha = 0.246$; three folds select $\alpha = 0$. Thus, the learned hierarchy emphasizes FC within the semantic expert and keeps the data-driven backbone as the dominant prediction path across folds.

Mask selection also reveals the scale preferred by inner CV. For the matched ridge model, FC most often selects ROI-incident masks with $M = 10$ or 15 (9 folds each), with direct top-1200 selected in 7 folds. SC favors the compact ROI-incident $M = 5$ mask in 12 folds and direct top-1200 in 5 folds. In the NC-ridge ratio sweep, the ridge endpoint $\rho = 0$ provides the strongest aggregate setting, supporting SPEF-ridge as the principal formulation and NC-ridge as the network-controlled alternative within the same prior-guided subspaces.

5.4 Stable prediction-linked candidate biomarker maps

Figure 3 summarizes corrected modality-specific stability together with outer-fold perturbation behavior. For the matched-prior ridge model, mean FC edge-rank Spearman is 0.723 and SC is 0.554. The corresponding top-10 ROI Jaccard values are 0.629 for FC and 0.410 for SC, with multimodal top-10 Jaccard 0.480 and multimodal sign consistency 0.833. These values are computed from primal

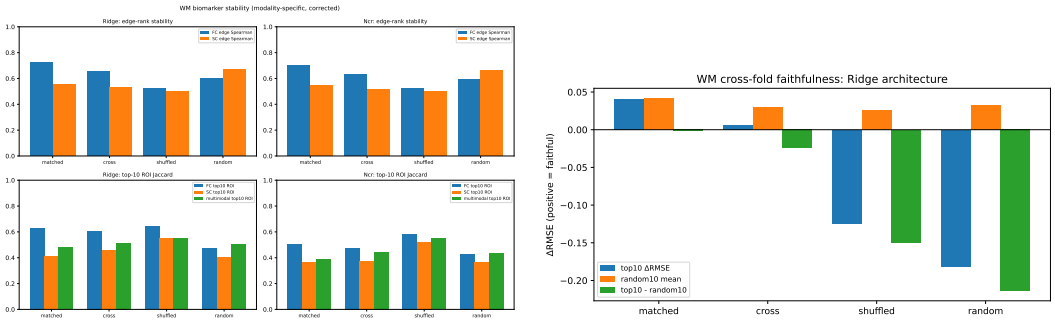


Figure 3: Prediction-linked biomarker analysis. Left: corrected modality-specific coefficient stability and top-10 ROI overlap across valid outer folds. Right: outer-fold perturbation for the ridge architecture; positive ΔRMSE indicates prediction degradation after masking. Stability and perturbation provide complementary evidence about the learned candidate maps.

maps c_F and c_S after excluding inactive modality folds, so each modality metric compares valid coefficient maps only.

A direct architecture-matched comparison with the no-prior full-connectome ridge expert clarifies what semantic restriction changes. SPEF-ridge increases FC edge-rank stability from 0.544 to 0.723. The no-prior expert, however, has higher SC edge-rank stability (0.714 vs. 0.554) and higher FC, SC, and multimodal top-10 ROI Jaccard overlap ($0.680/0.607/0.653$ vs. $0.629/0.410/0.480$), showing that semantic selection does not provide a uniform stability advantage.

The perturbation behavior differs more clearly. SPEF's top-10 ROIs produce mean $\Delta\text{RMSE} = +0.0404$ compared with $+0.0421$ for random 10-ROI masks, giving a top-minus-random contrast of -0.0017. For the no-prior expert, masking the top-10 ROIs gives $\Delta\text{RMSE} = -0.0401$ despite a random-mask mean of $+0.0463$, yielding a markedly more adverse contrast of -0.0864. Thus, prior-guided restriction improves FC reproducibility and avoids the strongly counter-faithful top-ranked pattern obtained from the full-connectome expert, while its biomarker advantage is not uniform across all stability measures. Full comparisons are reported in the supplement.

Figure 2 localizes the strongest coefficient-derived candidate biomarkers on the AAL116 atlas, displaying only regions whose normalized importance is at least 0.5. The highest multimodal ROI importances are Frontal_Mid_L, Frontal_Sup_R, Frontal_Sup_L, Frontal_Mid_R, Parietal_Sup_R, Angular_R, Frontal_Inf_Oper_L, Parietal_Sup_L, Angular_L, and Parietal_Inf_R. The resulting pattern is concentrated in frontal and parietal association cortex, consistent with a frontoparietal working-memory hypothesis. Because the semantic prior defines the expert subspace, we interpret this anatomical agreement as a mechanistic trace of prior propagation; prospective biological validation forms a distinct subsequent stage.

6 Conclusion

We presented SPEF, which converts cognitive-task semantics into a frozen anatomical prior that defines compact FC and SC subspaces for ridge and NC-ridge experts before multimodal fusion with a data-driven backbone. On 510 HCP participants, SPEF-ridge and cross-task SPEF-NC-ridge achieved the best Pearson correlation, while the latter achieves the lowest RMSE and MAE among the evaluated multimodal graph comparators and multimodal baselines, while matched semantic priors consistently outperformed shuffled and random controls. The close performance of matched and cross-task priors suggests that coherent higher-cognitive neuroanatomical structure contributes more clearly than fine task specificity at the present sample size. Candidate-map analysis further showed higher FC edge-rank stability than an architecture-matched no-prior full-connectome ridge model and a substantially less adverse perturbation contrast, although SC and ROI-overlap stability did not improve uniformly. These findings support task-semantic priors as an auditable inductive bias for multimodal connectome prediction and prediction-linked candidate biomarker analysis.

AI USE STATEMENT

Qwen3.8-27B was used to generate the frozen task-semantic anatomical prior from the working-memory task description and AAL116 region names. The resulting ROI prior was fixed before model evaluation and was subsequently transformed into edge-level scores for prior-guided feature selection and structured controls. Generative AI tools assisted with code inspection, artifact validation, figure preparation, and manuscript language editing. The authors take full responsibility for the methodology, scientific interpretation, and submitted content.

REPRODUCIBILITY STATEMENT

The supplementary material and accompanying code release document the complete SPEF evaluation protocol, including cohort construction, preprocessing, task-semantic prior provenance, ROI-to-edge prior construction, prior-selected FC and SC feature spaces, nested model selection, out-of-fold modality fusion, hierarchical fusion with the data-driven backbone, architecture-matched semantic controls, coefficient reconstruction, candidate-biomarker extraction, stability analysis, perturbation-based faithfulness analysis, and statistical reporting. All hyperparameter grids, fold and seed definitions, fitting scopes, deterministic tie-breaking rules, and evaluation metrics used for the reported experiments are specified explicitly. Reproducibility scripts are organized around the analysis stages described in the supplement and generate the manuscript tables, figures, biomarker summaries, and audit artifacts from the frozen inputs.

ETHICS STATEMENT

This study performs secondary analysis of Human Connectome Project Young Adult data under the applicable HCP Data Use Terms. No participants were recruited or contacted as part of this work, and no attempt was made to re-identify participants. Data access and analysis followed the applicable institutional requirements and HCP usage conditions.

REFERENCES

Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience*, 10(3):186–198, 2009.

Yi Hao Chan, Conghao Wang, Wei Kwek Soh, and Jagath C Rajapakse. Combining neuroimaging and omics datasets for disease classification using graph neural networks. *Frontiers in Neuroscience*, 16:866666, 2022.

Yi Hao Chan, Wei Chee Yew, Qian Hui Chew, Kang Sim, and Jagath C Rajapakse. Elucidating salient site-specific functional connectivity features and site-invariant biomarkers in schizophrenia via deep neural networks. *Scientific Reports*, 13(1):21047, 2023.

Yi Hao Chan, Deepank Girish, Sukrit Gupta, Jing Xia, Chockalingam Kasi, Yinan He, Conghao Wang, and Jagath C Rajapakse. Discovering robust biomarkers of psychiatric disorders from resting-state functional mri via graph neural networks: A systematic review. *NeuroImage*, pp. 121422, 2025.

Niharika Shimona Dsouza, Mary Beth Nebel, Deana Crocetti, Joshua Robinson, Stewart Mostofsky, and Archana Venkataraman. 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.

Emily S Finn, Xilin Shen, Dustin Scheinost, Monica D Rosenberg, Jessica Huang, Marvin M Chun, Xenophon Papademetris, and R Todd Constable. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature neuroscience*, 18(11):1664–1671, 2015.

Deepank Girish, Yi Hao Chan, Sukrit Gupta, Jing Xia, and Jagath C Rajapakse. Encoding and decoding of brain dynamic functional connectivity for adhd diagnosis. *IEEE Journal of Biomedical and Health Informatics*, 2026.

Matthew F. Glasser, Stamatios N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R. Polimeni, David C. Van Essen, and Mark Jenkinson. The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80:105–124, 2013.

Sukrit Gupta, Marcus Lim, and Jagath C Rajapakse. Decoding task specific and task general functional architectures of the brain. *Human Brain Mapping*, 43(9):2801–2816, 2022.

Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.

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.

Mark Jenkinson, Christian F Beckmann, Timothy EJ Behrens, Mark W Woolrich, and Stephen M Smith. Fsl. *Neuroimage*, 62(2):782–790, 2012.

Gia H Ngo, Minh Nguyen, Nancy F Chen, and Mert R Sabuncu. A transformer-based neural language model that synthesizes brain activation maps from free-form text queries. *Medical image analysis*, 81:102540, 2022.

Amashi Niwarthana and Jagath C Rajapakse. Labma: Latent-bottleneck attention-based multimodal architecture for integration of transcriptomics, proteomics, and mri in neurodegeneration. *Computers in Biology and Medicine*, 210:111740, 2026.

Maria Giulia Preti and Dimitri Van De Ville. Decoupling of brain function from structure reveals regional behavioral specialization in humans. *Nature communications*, 10(1):4747, 2019.

Taylor Salo, Tal Yarkoni, Thomas E. Nichols, Jean-Baptiste Poline, Murat Bilgel, Katherine L. Bottenhorn, Dorota Jarecka, James D. Kent, Adam Kimbler, Dylan M. Nielson, Kendra M. Oudyk, Julio A. Peraza, Alexandre Pérez, Puck C. Reeders, Julio A. Yanes, and Angela R. Laird. Nimare: neuroimaging meta-analysis research environment. *Aperture Neuro*, 3:1–32, 2023.

Xilin Shen, Emily S Finn, Dustin Scheinost, Monica D Rosenberg, Marvin M Chun, Xenophon Papademetris, and R Todd Constable. Using connectome-based predictive modeling to predict individual behavior from brain connectivity. *nature protocols*, 12(3):506–518, 2017.

Robert E Smith, Jacques-Donald Tournier, Fernando Calamante, and Alan Connelly. Sift2: Enabling dense quantitative assessment of brain white matter connectivity using streamlines tractography. *Neuroimage*, 119:338–351, 2015.

Stephen M. Smith, Christian F. Beckmann, Jesper Andersson, Edward J. Auerbach, Janine Bijsterbosch, Gwenaëlle Douaud, Eugene Duff, David A. Feinberg, Ludovica Griffanti, Michael P. Harms, Michael Kelly, Timothy Laumann, Karla L. Miller, Steen Moeller, Steve Petersen, Jonathan Power, Gholamreza Salimi-Khorshidi, Abraham Z. Snyder, An T. Vu, Mark W. Woolrich, Junqian Xu, Essa Yacoub, Kamil Uğurbil, David C. Van Essen, and Matthew F. Glasser. Resting-state fmri in the human connectome project. *Neuroimage*, 80:144–168, 2013.

Stamatios N. Sotiropoulos, Saad Jbabdi, Junqian Xu, Jesper L. Andersson, Steen Moeller, Edward J. Auerbach, Matthew F. Glasser, Moises Hernandez, Guillermo Sapiro, Mark Jenkinson, David A. Feinberg, Essa Yacoub, Christophe Lenglet, David C. Van Essen, Kamil Ugurbil, and Timothy E.J. Behrens. Advances in diffusion mri acquisition and processing in the human connectome project. *Neuroimage*, 80:125–143, 2013.

Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: a structural description of the human brain. *PLoS computational biology*, 1(4):e42, 2005.

Alina Tetereva, Annchen R Knodt, Tracy R Melzer, William van der Vliet, Bryn Gibson, Ahmad R Hariri, Ethan T Whitman, Jean Li, Farzane Lal Khakpoor, Jeremiah Deng, David Ireland, Sandhya Ramrakha, and Narun Pat. Improving predictability, reliability, and generalizability of brain-wide associations for cognitive abilities via multimodal stacking. *PNAS nexus*, 4(6):pgaf175, 2025.

J-Donald Tournier, Robert Smith, David Raffelt, Rami Tabbara, Thijs Dhollander, Maximilian Pietsch, Daan Christiaens, Ben Jeurissen, Chun-Hung Yeh, and Alan Connelly. Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *Neuroimage*, 202: 116137, 2019.

D.C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T.E.J. Behrens, R. Bucholz, A. Chang, L. Chen, M. Corbetta, S.W. Curtiss, S. Della Penna, D. Feinberg, M.F. Glasser, N. Harel, A.C. Heath, L. Larson-Prior, D. Marcus, G. Michalareas, S. Moeller, R. Oostenveld, S.E. Petersen, F. Prior, B.L. Schlaggar, S.M. Smith, A.Z. Snyder, J. Xu, and E. Yacoub. The human connectome project: a data acquisition perspective. *Neuroimage*, 62(4):2222–2231, 2012.

Jing Xia, Yi Hao Chan, Deepank Girish, and Jagath C Rajapakse. Img-gcn: Interpretable modularity-guided structure-function interactions learning for brain cognition and disorder analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 470–480. Springer, 2024.

Jing Xia, Yi Hao Chan, Deepank Girish, Qian Hui Chew, Kang Sim, and Jagath C Rajapakse. Disentangling shared and unique brain functional changes associated with clinical severity and cognitive phenotypes in schizophrenia via deep learning. *Communications Biology*, 8(1):1215, 2025a.

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, 2025b.

Jing Xia, Yi Hao Chan, and Jagath C Rajapakse. Meta-analysis guided multi-task graph transformer network for diagnosis of neurological disease and cognitive deficits. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 459–468. Springer, 2025c.

Jing Xia, Yi Hao Chan, Deepank Girish, Conghao Wang, and Jagath C Rajapakse. Learning brain structure-function interactions via modularity-guided attention-based graph neural networks. *Information Fusion*, pp. 104426, 2026.

Shunxin Xiao, Huibin Lin, Conghao Wang, Shiping Wang, and Jagath C Rajapakse. Graph neural networks with multiple prior knowledge for multi-omics data analysis. *IEEE journal of biomedical and health informatics*, 27(9):4591–4600, 2023.

Tal Yarkoni, Russell A Poldrack, Thomas E Nichols, David C Van Essen, and Tor D Wager. Large-scale automated synthesis of human functional neuroimaging data. *Nature methods*, 8(8):665–670, 2011.

Yanteng Zhang, Xiaohai He, Yi Hao Chan, Qizhi Teng, and Jagath C Rajapakse. Multi-modal graph neural network for early diagnosis of alzheimer's disease from smri and pet scans. *Computers in Biology and Medicine*, 164:107328, 2023.