



Supervised Deep Multimodal Matrix Factorization for Interpretable Brain Network Analysis

AMJAD SEYEDI¹, LIFANG HE², SONGLIN ZHAO², AKWUM ONWUNTA³, AND
NICOLAS GILLIS¹

¹Department of Mathematics and Operational Research, University of Mons, Mons, Belgium

²Department of Computer Science and Engineering, Lehigh University, Bethlehem, PA, USA

³Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA, USA

ISE Technical Report 26T-010



Supervised Deep Multimodal Matrix Factorization for Interpretable Brain Network Analysis *

Amjad Seyedi

Dept. of Mathematics & Operational Research
University of Mons, Mons, Belgium
seyedamjad.seyedi@umons.ac.be

Lifang He

Dept. of Computer Science & Engineering
Lehigh University, Bethlehem, PA, USA
lih319@lehigh.edu

Songlin Zhao

Dept. of Computer Science & Engineering
Lehigh University, Bethlehem, PA, USA
soz223@lehigh.edu

Akwum Onwunta

Dept. of Industrial & Systems Engineering
Lehigh University, Bethlehem, PA, USA
ako221@lehigh.edu

Nicolas Gillis

Dept. of Mathematics & Operational Research
University of Mons, Mons, Belgium
nicolas.gillis@umons.ac.be

Abstract

We present Supervised Deep Multimodal Matrix Factorization (SD3MF), an interpretable framework for integrative brain network analysis that generalizes Symmetric Nonnegative Matrix Tri-Factorization (SNMTF) from unsupervised single-graph clustering to supervised prediction over populations of multimodal graphs. SD3MF learns deep hierarchical factorizations for each modality together with a shared latent representation that aligns subjects across views. An encoder–decoder formulation jointly optimizes graph reconstruction and supervised prediction, while adaptive weights enable data-driven multimodal fusion. By representing each subject through community-level interaction matrices, the model yields interpretable and discriminative features. Experiments on multimodal connectome datasets show that SD3MF consistently outperforms strong deep learning baselines such as CNNs and GNNs, while enabling biologically interpretable insights. Code for reproducibility is available at: <https://github.com/amjadseyedi/SD3MF>.

1 Introduction

Modeling the brain as an interconnected network has recently attracted strong interest because it addresses fundamental questions about cognition, information flow, and distributed brain activity [1]. Network-based methods aid early detection of neurodegenerative disorders such as Alzheimer’s and Parkinson’s diseases, and HIV-dementia [2, 3]. In parallel, multimodal brain network analysis—combining sMRI, DTI, and fMRI—has emerged as a powerful framework for studying complex brain processes [4, 5]. Integrating these modalities provides a more complete view of network architecture and pathology, supporting improved diagnosis and biomarker discovery.

* Amjad Seyedi (AS) and Nicolas Gillis (NG) gratefully acknowledge support from the European Union (ERC Consolidator Grant, eLinoR, no. 101085607). Lifang He (LH) and Akwum Onwunta (AO) gratefully acknowledge support from the U.S. Department of Energy (DOE grant DE-SC0025801, "Harnessing Nonnegative Matrix Factorization for Advanced Computational Materials Modeling").

Multimodal brain network analysis has employed a wide range of machine learning methods, from canonical correlation analysis [6, 7] and graph kernels [8] to Tensor Decompositions (TDs) [9–11] and CNNs [12–15]. Despite progress, no model consistently captures multimodal brain-network characteristics. Shallow methods struggle with nonlinear topology, Convolutional Neural Networks (CNNs) may distort graph structure, and many approaches ignore cross-modal correlations by treating modalities uniformly. Graph Convolutional Networks (GCNs) have emerged as a strong alternative for graph-structured multimodal data. Methods such as MVGCN [16] and MVS-GCN [17] improve prediction in Parkinson’s disease and autism, respectively, but often rely on prior knowledge to build shared spaces. Masking-based variants [18, 19] simplify models yet may miss complex inter-modality dependencies, while hybrid methods like GCN-SVM [20] depend on predefined graphs or KNN constructions, limiting generalization.

Despite the high predictive accuracy of CNNs and GNNs in brain network analysis, these models often operate as "black boxes". While they capture complex nonlinear patterns, they provide limited insight into identifiable, biologically meaningful brain modules and can exhibit high variance in small-sample settings typical of neuroimaging data. Conversely, traditional Matrix Factorization (MF) methods are inherently linear and offer strong interpretability through parts-based representations, but lack the capacity to model quadratic interactions and hierarchical structure in brain networks. This creates a fundamental gap between expressive but opaque models and interpretable but limited ones, motivating the need for a framework that combines hierarchical modeling power with structured and interpretable representations.

To tackle these issues, we introduce Supervised Deep Multimodal Matrix Factorization (SD3MF), an interpretable deep learning framework built as an alternative to CNN- and GNN-based models for graph classification, specifically for multimodal brain network analysis. It is a general framework that can be applied to other multimodal structured data, e.g., social networks, biomedical multi-omics data, sensor networks, and knowledge graphs. SD3MF extends nonnegative matrix tri-factorization (SNMTF) into a hierarchical architecture capable of capturing multilevel structure within each connectivity modality while simultaneously learning a shared subject representation that aligns modalities in a biologically meaningful manner. An adaptive mixture mechanism automatically determines the relative importance of each brain network modality, enabling principled, non-heuristic multimodal fusion. By integrating representation learning with supervised prediction, SD3MF ensures that the extracted low-rank network modules are both interpretable and discriminative for clinical or cognitive labels. Across various multimodal connectome datasets, SD3MF consistently outperforms state-of-the-art deep graph baselines while offering transparent, cluster-based explanations of the brain systems most relevant to the predictive task.

2 Related Work

Storing relationships among subjects, imaging modalities, and brain connections naturally leads to tensor representations. Many methods reduce these high-order structures through decomposition. MPCA [21] extends PCA to multidimensional arrays to extract features from combined imaging sources [22], while MIC [23] builds similarity tensors and applies CP decomposition to derive subject profiles. Other approaches employ t-product factorizations with sparsity constraints [24], tensor splitting for multiview clustering [25], or structured decompositions [26] to capture patterns across multiple graphs. Recent work extends tensor fusion by stacking adjacency matrices from multiple functional network estimators into higher-order tensors and applying CP or Tucker decompositions to learn shared embeddings [24]. In a complementary direction, Li et al. [27] introduce a slow-thinking module that uses TDs to construct an explicit knowledge graph, refining multimodal fusion and improving disease identification. Despite their strengths, many tensor methods rely on multilinear projections that flatten hierarchical brain network structure and fail to capture complex nonlinear cross-modal interactions. As a result, subtle but critical relationships across scales may be lost, weakening detection and interpretability.

GCNs have become central to multimodal brain network analysis [28–30]. MVGCN [16] aligns modalities in a shared spatial framework but depends on auxiliary information often unavailable in heterogeneous data. Attention-based fusion of multiple GCNs [31] requires flattened feature vectors that inflate dimensionality and obscure structure. Other strategies enforce cross-modal guidance through shared graph patterns [32], though misalignment between modalities can degrade learning. Late-fusion designs [33] delay cross-modal interaction, limiting early representation learning. Ex-

plainable GNNs for biomarker discovery [34] often rely on small cohorts, restricting generalization. Methods reshaping functional connectivity using structural priors [35] and multiview alignment [17] depend on fixed templates or consistent data availability, while rigid anatomical constraints limit robustness across datasets [20]. Lightweight masking or region-level models [19, 18] improve interpretability but may oversimplify cross-modal interactions or require stable image registration. More recent studies combine tensor models with GCNs. Some extract shared representations via TDs before graph learning [36], while others construct graph operators using Kronecker sums [37] or t-products for dynamic and multilayer graphs [38, 39]. Low-rank tensor denoising improves connectivity estimates [40], structural priors guide early GCN layers with attention refinement [41], and adversarial VAE frameworks fuse latent subspaces across modalities [42]. Complex-valued tensor GNNs track evolving regional interactions [43], and TGNet [44] integrates TDs with multilayer GCNs to model multimodal brain networks for improved disease classification.

3 Background

This section reviews SNMTF, its deep hierarchical extension, and supervised MFs, which together motivate our supervised deep multimodal MF framework.

SNMTF. SNMTF is a proximity-aware MF model that facilitates interpretability through structured nonnegative factors. Let $\mathbf{A} \in \mathbb{R}_+^{n \times n}$ be a nonnegative symmetric matrix, interpreted as the affinity matrix of a graph over n nodes, e.g., n regions of the brain. An important task in graph theory is to identify densely connected subsets of nodes, a.k.a. communities [45, 46]. Most community detection algorithms work in a single-layer fashion; that is, each node is associated with some communities, and possibly, the communities interact with each other. A popular model is SNMTF:

$$\mathbf{A} \approx \mathbf{W}\mathbf{S}\mathbf{W}^\top, \quad (1)$$

where $\mathbf{W} \in \mathbb{R}_+^{n \times r}$ contains the community memberships, that is, $W(i, k) > 0$ if node i belongs to the community k , while symmetric $\mathbf{S} \in \mathbb{R}_+^{r \times r}$ indicates the interactions between the communities, that is, $S(k, \ell)$ indicates the strength between community k and ℓ . This model is closely related to the degree-corrected block model [47] and the mixed membership stochastic block model [48].

Deep SNMTF. SNMTF was extended to a deeper framework [49], as it was done for standard MFs [50], enabling hierarchical representation learning. In this deep setting, the adjacency matrix is factorized in multiple stages, starting with $\mathbf{A} \approx \mathbf{W}_1 \mathbf{S}_1 \mathbf{W}_1^\top$, where \mathbf{W}_1 represents the membership of nodes to low-level communities and \mathbf{S}_1 captures the interactions among these communities. The interaction matrix \mathbf{S}_1 is then decomposed as $\mathbf{S}_1 \approx \mathbf{W}_2 \mathbf{S}_2 \mathbf{W}_2^\top$, continuing recursively until $\mathbf{S}_{L-1} \approx \mathbf{W}_L \mathbf{S}_L \mathbf{W}_L^\top$. Each layer progressively aggregates smaller communities into higher-level structures, resulting in a hierarchical graph summarization that reflects both local and global connectivity patterns. The optimization problem is formulated using the Frobenius norm [50, 49], solving

$$\min_{\Psi = \prod_{l=1}^L \mathbf{W}_l, \mathbf{S}_l} \|\mathbf{A} - \Psi \mathbf{S}_L \Psi^\top\|_F^2 \text{ s.t. } \mathbf{W}_l, \mathbf{S}_l \geq 0, \quad (2)$$

where $\Psi = \prod_{l=1}^L \mathbf{W}_l$ denotes the aggregated membership matrix. This formulation preserves the nonnegativity and interpretability of MFs while enabling multi-level community discovery within a unified deep optimization framework.

Supervised Matrix Factorization. MF is a classical unsupervised feature extraction framework that learns latent structures in complex datasets. Nonnegative MF (NMF), a popular variant, enables interpretable parts-based representations through nonnegativity constraints [51]. Recently, Supervised MF (SMF) [52, 53] extends NMF by incorporating label information to learn low-dimensional features and perform classification jointly. By aligning representation learning with predictive tasks, SMF improves the relevance of extracted features and has been applied to tabular, single-view data in areas like text classification, bioinformatics, and image analysis. Let $a_i \in \mathbb{R}$ be the activation corresponding to label y_i , defined as:

$$a_i = \begin{cases} \beta^\top \mathbf{W}^\top \mathbf{x}_i + \gamma^\top \mathbf{x}'_i & \text{for SMF-W,} \\ \beta^\top \mathbf{h}_i + \gamma^\top \mathbf{x}'_i & \text{for SMF-H,} \end{cases}$$

where (β, γ) are logistic regression coefficients associated with either the ‘filtered’ feature $\mathbf{W}^\top \mathbf{x}_i$ (SMF-W), or the learned code \mathbf{h}_i (SMF-H), and side information \mathbf{x}'_i . Let $Z := (\mathbf{W}, \mathbf{H}, \beta, \gamma)$ denote

the block parameters. To estimate Z from observed data $(\mathbf{x}_i, \mathbf{x}'_i, y_i)$ for $i \in \{1, \dots, N\}$, Lee et al. [53] consider the following problem:

$$\min_{\mathbf{W}, \mathbf{H}, \beta, \mathbf{\Gamma}} f(Z) := \xi \|\mathbf{X} - \mathbf{W}\mathbf{H}\|_F^2 + \sum_{i=1}^N \ell(y_i, a_i) \quad (3)$$

with the logistic loss $\ell(y_i, a_i) = \log(1 + \exp(a_i)) - y_i a_i$. SMF provides a principled framework for supervised representation learning that retains the transparency and interpretability of MF while achieving competitive predictive performance. A key weakness is the heuristic synchronization of Lee et al. [53] that separately updates representation and classification components.

4 Proposed Model

We propose a supervised graph representation learning framework designed for populations of weighted, undirected/directed networks with a shared node set, where each graph corresponds to an individual subject and the learning goal is graph-level prediction (graph classification). The model is applicable to a broad class of network-valued data in which meaningful low-dimensional structure can be expressed through latent groups of nodes and their interaction patterns. We focus on brain connectomes as a primary application, where nodes represent regions of interest (ROIs) and edges encode functional or structural connectivity. Our key assumption is that such networks can be compactly and discriminatively characterized by a hierarchy of latent communities and by the patterns of intra- and inter-community interactions. To this end, the proposed model learns a shared, possibly deep, node-to-community embedding that captures population-level structure, together with subject-specific community interactions that summarize each network at the community level. These low-dimensional representations are both interpretable and well-suited for supervised learning. Let $\{\mathbf{A}_i\}_{i=1}^N$ denote a collection of adjacency matrices (connectomes), where each $\mathbf{A}_i \in \mathbb{R}^{n \times n}$ corresponds to subject i over n ROIs. Associated with each subject is a label $y_i \in \mathcal{Y}$, where \mathcal{Y} denotes the label space (e.g., diagnostic categories or continuous clinical scores). A graph representation function f_θ maps \mathbf{A}_i to a compact embedding $\mathbf{S}_i = f_\theta(\mathbf{A}_i)$, and a classifier g_ϕ produces the prediction $\hat{y}_i = g_\phi(\mathbf{S}_i)$. In our framework, f_θ is implemented via (deep) SNMTF, where $\mathbf{S}_i \in \mathbb{R}^{r \times r}$ explicitly encodes community-to-community interactions.

4.1 Network Representation Learning

To extract network-level representations as subject-specific community interactions, we employ Collective Deep SNMTF. The key quantity \mathbf{S}_i summarizes both within-community cohesion (diagonal entries) and between-community coupling (off-diagonal entries). When \mathbf{S}_i is allowed to take signed values, positive entries indicate strong connectivity (positive correlation) between communities, whereas negative entries indicate anti-connectivity (negative correlation) [49]. To capture shared structure across subjects, we learn a common node-to-community mapping $\Psi \in \mathbb{R}^{n \times r}$ by solving

$$\min_{\Psi = \prod_{l=1}^L \mathbf{W}_l, \mathbf{S}_i} \sum_{i=1}^N \|\mathbf{A}_i - \Psi \mathbf{S}_i \Psi^\top\|_F^2 \quad \text{s.t.} \quad \mathbf{W}_l \geq 0, \quad (4)$$

where $\mathbf{S}_i \in \mathbb{R}^{r \times r}$ is the subject-specific representation for graph \mathbf{A}_i . We enforce nonnegativity only on the latent factors \mathbf{W}_l , which preserves the interpretability. In contrast, \mathbf{S}_i is left unconstrained and may take signed values, to represent both positive and negative interactions.

4.2 Basic Model: Supervised Deep MF

We next integrate supervised learning into SNMTF. The goal is to learn structure-aware model parameter set $\Psi = \prod_{l=1}^L \mathbf{W}_l$, interaction-aware network representation \mathbf{S}_i , and classifier parameters β jointly, so that \mathbf{S}_i remains faithful to the observed graph while becoming predictive of the label y_i . Notably, although the final classifier is linear, the underlying representation induces quadratic interactions between node memberships and network connectivity, enabling nonlinear modeling of community-level structure prior to classification. The proposed model integrates two main objectives: (1) decoder reconstruction, where the factorization $\mathbf{A}_i \approx \Psi \mathbf{S}_i \Psi^\top$ enables the model to capture community structures within each graph; (2) supervised encoding, where each network \mathbf{A}_i is encoded as $\Psi^\top \mathbf{A}_i \Psi$, and a supervised loss is applied to its vectorized form $\text{vec}(\Psi^\top \mathbf{A}_i \Psi)$, to align the learned representation with class labels y_i . We propose to solve

$$\min_{\Psi, \mathbf{S}_i, \beta} \sum_{i=1}^N \mu \|\mathbf{A}_i - \Psi \mathbf{S}_i \Psi^\top\|_F^2 + \ell(y_i, \beta^\top \text{vec}(\Psi^\top \mathbf{A}_i \Psi)) \quad \text{s.t.} \quad \Psi = \prod_{l=1}^L \mathbf{W}_l, \mathbf{W}_l \geq 0, \Psi \mathbf{1} = \mathbf{1},$$

where $\ell(\cdot)$ denotes the cross-entropy loss and μ is a regularization hyperparameter that balances the decoding (reconstruction) term and the predictive accuracy. The constraint $\Psi \mathbf{1} = \mathbf{1}$, where $\mathbf{1}$ is the vector of all ones of appropriate dimension, allows to interpret Ψ as a fuzzy membership matrix. The model hierarchically encodes network structure while preserving interpretability. Joint reconstruction and supervision ensure that the learned representations are structurally faithful and predictive. The matrix \mathbf{S}_i captures subject-specific interactions in the learned latent space, providing a compact summary of individual network variations.

4.3 Final Model: Supervised Deep Multimodal MF

Traditional brain network analysis has largely focused on single-modality data, capturing only isolated aspects of brain organization. To address this limitation, we develop a multimodal version of our model that performs model-level fusion, jointly learning from multiple modalities within a single end-to-end architecture. This approach enables a more comprehensive view of brain networks by leveraging complementary information across data types. The model fuses different (e.g., structural and functional) connectivity data into a shared representation space as follows:

$$\begin{aligned} \min_{\Psi^{(m)}, \mathbf{S}_i, \beta, \alpha} \sum_{i=1}^N \sum_{m=1}^M \mu \|\mathbf{A}_i^{(m)} - \Psi^{(m)} \mathbf{S}_i \Psi^{(m)\top}\|_F^2 + \ell(y_i, \beta^\top \sum_{m=1}^M \alpha^{(m)} \text{vec}(\Psi^{(m)\top} \mathbf{A}_i^{(m)} \Psi^{(m)})) \\ \text{s.t. } \mathbf{W}_l^{(m)} \geq 0, \Psi^{(m)} \mathbf{1} = \mathbf{1}, \sum_{m=1}^M \alpha_m = 1, \end{aligned} \quad (5)$$

where $\mathbf{A}_i^{(m)}$ is the graph associated to the i th subject in mode m , $\Psi^{(m)} = \prod_{l=1}^L \mathbf{W}_l^{(m)}$ is the membership matrix of mode m , $\alpha \in \mathbb{R}^M$ control the contribution of each mode in the final classifier. The fused feature vector, $\mathbf{v}_i = \sum_{m=1}^M \alpha^{(m)} \text{vec}(\Psi^{(m)\top} \mathbf{A}_i^{(m)} \Psi^{(m)})$, is constructed as a convex combination of modality-specific community-level summaries. Intuitively, $\Psi^{(m)\top} \mathbf{A}_i^{(m)} \Psi^{(m)} \in \mathbb{R}^{r \times r}$ aggregates ROI-to-ROI connectivity into community-to-community interactions for modality m , yielding an interpretable low-dimensional representation of intra- and inter-community coupling. The classifier then maps the fused representation to a probability via the logistic link, i.e., $\hat{y}_i = \sigma(\beta^\top \mathbf{v}_i)$ for binary classification (or softmax in the multi-class setting), and the parameters $\{\Psi^{(m)}\}_{m=1}^M$, $\{\mathbf{S}_i\}_{i=1}^N$, β , and α are learned jointly by minimizing the multimodal objective in (5).

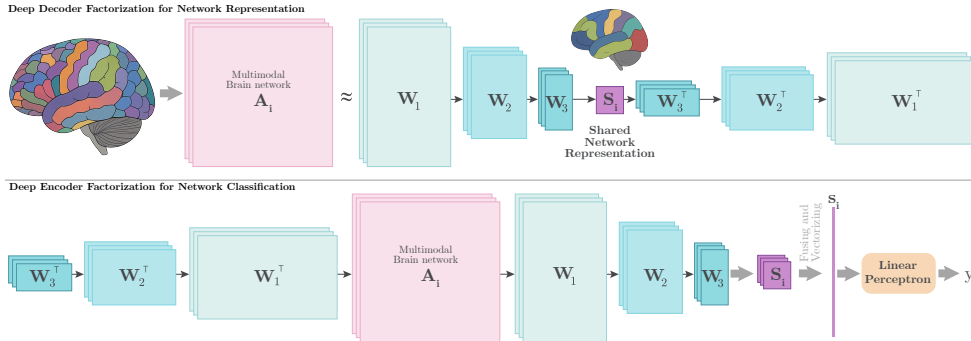


Figure 1: SD3MF architecture: The decoder (top) reconstructs each of the three modality-specific connectomes from a shared subject representation, while the encoder (bottom) produces a discriminative representation fed to a linear classifier.

Figure 1 illustrates the SD3MF pipeline. For each subject i and modality m , we learn a deep, nonnegative node-to-community mapping parameterized by a product of L factor matrices, $\Psi^{(m)} = \prod_{l=1}^L \mathbf{W}_l^{(m)}$. The decoder (top) reconstructs the observed connectome via a shared low-rank community interaction matrix \mathbf{S}_i , i.e., $\mathbf{A}_i^{(m)} \approx \Psi^{(m)} \mathbf{S}_i \Psi^{(m)\top}$, which preserves community structure and yields an interpretable representation. In parallel, the encoder (bottom) produces the community-level summary $\Psi^{(m)\top} \mathbf{A}_i^{(m)} \Psi^{(m)} \in \mathbb{R}^{r \times r}$, which is vectorized and optionally fused across modalities using weights $\{\alpha^{(m)}\}_{m=1}^M$. The fused representation is then fed to a linear classifier to predict y_i . Training jointly optimizes reconstruction and supervised objectives, encouraging \mathbf{S}_i to remain structurally faithful while being discriminative for classification.

4.4 Optimization


The optimization problem is inherently non-convex due to the deep MF and shared latent representations across multiple views. Rather than relying on explicit regularization, we rely on implicit regularization emerging from gradient-based optimization with small learning rates and small-scale initialization, which biases the learned factors toward low-rank and well-conditioned solutions [54]. Each deep basis was initialized using orthogonal matrices scaled to small magnitudes (10^{-3} in our experiments). Empirically, we observed that small-scale initialization substantially improved optimization stability and final predictive performance, likely by promoting smoother optimization dynamics during the early stages of training. The shared latent matrices S_i were initialized from the view-averaged projections of the input matrices. The classifier head was initialized using Xavier uniform initialization, also at a small scale, to ensure balanced optimization dynamics across modules. Training was performed using stochastic gradient descent (SGD) with a fixed learning rate of 10^{-5} and a maximum of 30,000 iterations. We specifically favor SGD over adaptive methods like Adam to preserve the gradient noise structure, which is essential for the implicit rank-minimization bias required in matrix factorization [54]. The relatively large iteration count allows the small learning rate and implicit regularization dynamics to fully evolve while avoiding instability that could arise from aggressive updates in this highly non-convex setting. Training exhibits stable convergence in practice; a detailed analysis of the optimization behavior is provided in Appendix A. Additional details, including model efficiency in comparison to GNN-based models, are provided in Appendix B.

5 Experiments

We assessed our approach on three real clinical neuroimaging datasets spanning multiple imaging modalities. These datasets are of small to moderate scale, reflecting typical cohort sizes in multi-modal brain studies. To limit bias, major confounds (e.g., age and gender) were handled during recruitment and preprocessing so group distributions remained comparable. Acquisition was also kept consistent within each dataset to support fair comparisons. Specifically, the HIV dataset consists of connectivity networks over 90 ROIs from fMRI and DTI modalities with 35 healthy and 35 diseased subjects; the BP dataset includes 82 ROIs from fMRI and DTI modalities with 45 healthy and 52 diseased subjects; and the PPMI dataset contains 84 ROIs with structural connectivity derived from PICo, Hough, and Protrackx tractography, with 149 healthy and 569 diseased subjects; complete acquisition/preprocessing details are in Appendix C. For each dataset, we employ a stratified random split with 80% of the samples used for training and 20% for evaluation, repeated over 10 runs with different random seeds, and report mean \pm standard deviation.

Table 1: Performance comparison on HIV, BP, and PPMI datasets. Figure 2: Sensitivity vs. specificity of top methods.

Method	HIV		BP		PPMI	
	ACC	AUC	ACC	AUC	ACC	AUC
M2E	50.61 \pm 15.84	51.53 \pm 13.68	57.78 \pm 12.61	53.63 \pm 11.82	76.98 \pm 8.65	71.53 \pm 8.34
MIC	55.63 \pm 15.28	56.61 \pm 13.43	51.21 \pm 13.78	50.12 \pm 16.78	78.03 \pm 8.36	72.42 \pm 7.93
MPCA	67.24 \pm 11.56	66.92 \pm 12.46	56.92 \pm 13.33	56.86 \pm 13.69	81.25 \pm 6.34	72.36 \pm 7.18
MK-SVM	67.71 \pm 13.18	69.89 \pm 10.36	60.12 \pm 10.83	56.78 \pm 12.86	81.68 \pm 5.06	75.96 \pm 6.89
3D-CNN	74.31 \pm 18.81	73.53 \pm 16.41	63.33 \pm 11.21	61.62 \pm 10.26	82.24 \pm 5.32	75.65 \pm 6.11
GAT	68.58 \pm 13.51	67.31 \pm 14.32	61.31 \pm 15.04	59.93 \pm 13.54	82.28 \pm 5.95	75.19 \pm 5.57
GCN	70.16 \pm 12.54	69.94 \pm 12.91	64.44 \pm 15.71	64.24 \pm 16.45	83.10 \pm 6.01	75.33 \pm 6.79
DiffPool	71.42 \pm 14.78	71.08 \pm 15.12	62.22 \pm 12.83	62.54 \pm 13.41	82.19 \pm 6.23	72.75 \pm 2.18
MVGCN	74.29 \pm 11.27	73.75 \pm 12.63	62.22 \pm 16.83	62.64 \pm 16.89	82.54 \pm 5.15	76.48 \pm 6.69
MVS-GCN	78.57 \pm 16.66	77.50 \pm 12.94	65.00 \pm 16.03	63.70 \pm 15.31	80.28 \pm 1.41	76.54 \pm 8.03
MaskGNN	75.71 \pm 11.41	79.20 \pm 12.43	64.95 \pm 10.34	64.71 \pm 10.31	80.43 \pm 5.68	72.24 \pm 5.14
GCN-SVM	62.86 \pm 6.55	63.33 \pm 15.00	61.67 \pm 5.15	57.33 \pm 8.21	78.87 \pm 4.86	70.25 \pm 9.59
SGCN	77.14 \pm 14.09	80.47 \pm 12.67	63.89 \pm 17.24	65.18 \pm 16.73	81.09 \pm 5.17	74.78 \pm 7.69
TGNet	81.39 \pm 13.41	82.08 \pm 14.81	67.78 \pm 12.28	66.31 \pm 10.24	83.94\pm6.09	77.93 \pm 6.86
SD3MF	87.50\pm3.60	95.31\pm2.62	74.00\pm4.90	85.05\pm2.34	82.37 \pm 2.07	79.89\pm2.53



Results. Table 1 reports the classification performance of all compared methods, including the same baselines considered in [44], which represent state-of-the-art approaches for supervised and unsupervised graph-based classification. Baseline results are taken from prior work evaluated on the same datasets, while SD3MF is implemented and evaluated using the same preprocessing and experimental protocol to ensure fair comparison. For our model, the architectural and optimization hyperparameters are selected following the protocol described in the Appendix D. Overall, SD3MF consistently achieves the best or near-best performance across all datasets and metrics, showing both superior predictive accuracy and strong ranking capability.

On HIV, a challenging and high-variance dataset, shallow multimodal methods perform poorly, while graph-based models yield substantial gains, highlighting the importance of relational modeling. TGNet is the strongest baseline (ACC 81.39%, AUC 82.08%), but SD3MF significantly outperforms all competitors (ACC 87.50%, AUC 95.31%) with markedly lower variance, indicating more effective multimodal integration and improved stability. On BP, GCN-based methods outperform kernel and CNN baselines, with TGNet again providing the best baseline performance (ACC 67.78%, AUC 66.31%). SD3MF achieves a clear margin (ACC 74.00%, AUC 85.05%), with a particularly large AUC gain, suggesting more discriminative representations in difficult or imbalanced settings and improved robustness. PPMI is comparatively less noisy, with higher baseline accuracies across methods. While TGNet attains the highest baseline ACC (83.94%), SD3MF obtains the best AUC (79.89%) with competitive accuracy (82.37%) and the smallest variance among top methods, indicating better calibration and stability on structured data. Across all datasets, graph-based approaches consistently outperform traditional multimodal methods, confirming the necessity of relational learning. SD3MF further improves both accuracy and robustness, particularly in low-sample regimes common in medical imaging. To further evaluate clinical performance, Figure 2 illustrates the sensitivity–specificity trade-off across top-performing methods and datasets. SD3MF consistently lies in the top-right region, demonstrating a more favorable balance between sensitivity and specificity, which is particularly important in clinical settings where both false positives and false negatives are critical.

Ablation Study. To better understand the behavior of SD3MF, we perform several ablation studies.

Effect of decoder factorization: We analyze the influence of the decoder reconstruction term by evaluating different values of the regularization parameter μ in independent runs. As shown in Figure 3, at $\mu=0$, where the decoder reconstruction term is removed, performance is substantially lower, confirming that decoder factorization plays a crucial role in regularizing the representation and improving generalization. Both ACC and AUC increase steadily as μ grows from very small values to $\mu \in [1, 10]$, indicating that incorporating decoder-based reconstruction promotes the learning of structurally faithful and more discriminative community representations. When μ becomes too large, performance degrades, suggesting that excessive emphasis on reconstruction weakens the supervised signal and leads to suboptimal classification. This behavior highlights the importance of balancing decoder factorization and supervised encoding: a moderate reconstruction weight enables SD3MF to preserve network structure while enhancing predictive performance.




Figure 3: Influence of the regularization parameter μ on model performance. Both ACC and AUC improve as $\mu \in [1, 10]$ and decrease for larger values.

Effect of Modality Types: To evaluate the contribution of multimodal fusion, we conduct an ablation study by training and testing the model using each modality independently and comparing the results with the full multimodal configuration (Table 2). Treating each modality in isolation effectively removes the fusion component of SD3MF, allowing us to isolate the impact of integrating complementary neural connectivity information. As reported in Table 2, SD3MF consistently achieves improved performance when combining modalities across all datasets. On the HIV dataset, the DTI-only variant outperforms the fMRI-only model, suggesting that structural connectivity provides more stable disease-related patterns in this cohort, while the multimodal model further improves performance by incorporating complementary functional information. A similar trend is observed on the BP dataset, where neither modality alone reaches the performance of the fused model, highlighting the importance of cross-modal interactions. For the PPMI dataset, although individual tractography-derived modalities already yield competitive results, their combination leads to the best overall accuracy and AUC, indicating that different tractography methods encode partially distinct yet complementary structural information. Importantly, these patterns are reflected in the learned modality weights $\alpha^{(m)}$,


which adaptively assign higher importance to more informative modalities while still preserving contributions from weaker ones, providing a principled and data-driven fusion mechanism within SD3MF. Overall, this ablation study confirms that the performance gains of SD3MF are not solely due to strong single-modal modeling, but critically rely on its adaptive multimodal integration strategy.

Effect of Model Depth: Figure 4 evaluates the effect of model depth by comparing the shallow S3MF variant with the deep SD3MF on the HIV, BP, and PPMI datasets in terms of ACC and AUC. Across all three datasets, SD3MF consistently achieves higher performance, suggesting that deep (multi-layer) factorization learns more informative community-level representations than a single-layer decomposition. In our setting, deepness is an overparameterization [55] that prevents converging to poor local minima [56]. The improvements are particularly strong on HIV (78.60→87.50 ACC; 83.33→95.31 AUC) and BP (70.00→74.00 ACC; 72.73→85.05 AUC), where deeper modeling appears to better handle noise and capture cross-community interaction patterns relevant for diagnosis. On PPMI, SD3MF also yields steady gains (78.50→82.37 ACC; 75.30→79.89 AUC), indicating that depth remains beneficial even when single-model performance is already competitive. Overall, these results support the use of deep, hierarchical structure in SD3MF to capture multi-scale organization in brain networks and improve downstream classification.

Table 2: Evaluating the SD3MF model across single-modal and multimodal neural connectivity data.

Dataset	Modality	Accuracy (%)	AUC (%)
HIV	fMRI	77.16±3.22	89.58±4.89
	DTI	82.13±6.84	91.15±2.62
	Both	87.50±3.60	95.31±2.62
BP	fMRI	71.00±5.50	78.18±2.50
	DTI	72.00±2.70	76.57±2.66
	Both	74.00±4.90	85.05±2.34
PPMI	PICo	78.93±2.48	75.66±2.44
	Hough	76.54±1.87	69.32±2.26
	Protrackx	74.07±2.07	71.06±2.12
	All	82.37±2.07	79.89±2.53

Figure 4: Shallow (S3MF) vs. deep (SD3MF) models in terms of ACC and AUC.



Interpretation of Learned Representations. We provide an intrinsic interpretability analysis by linking learned parameters to neurobiologically meaningful community structure and discriminative regional patterns. Specifically, we interpret (i) community organization and network composition from the membership matrices $\{\Psi^{(m)}\}_{m=1}^M$ and subject-specific interaction matrices \mathbf{S}_i of 2 groups, and (ii) salient ROIs from ROI-level scores derived from $\{\Psi^{(m)}\}_{m=1}^M$. Both analyses are conducted in a modality-specific manner and under the weighted multimodal fusion setting to enable comprehensive interpretation across imaging modalities. We illustrate these interpretations on the HIV dataset.

Interpreting Communities and Networks: We summarize the learned community composition on the HIV cohort for each modality (DTI, fMRI) and the multimodal setting using the corresponding basis matrix $\Psi^{(m)}$. The number of communities is set to 10, aligning with established evidence that the brain functional architecture is organized into approximately ten primary representative networks [57]. To make the community structure explicit, we assign each ROI p to its dominant community via a hard membership rule $g^{(m)}(p) = \arg \max_k \Psi_{p,k}^{(m)}$, and report the dominant ROIs per community (C1, C2, ..., C10) accordingly, as shown in Figure 5. Across modality settings, SD3MF consistently groups ROIs with coherent anatomical locations and functional roles into the same community. Multiple communities are enriched with occipital and early visual cortices, consistent with canonical visual resting-state networks [58]. For example, Figure 6 highlights an fMRI-derived community (C7 by fMRI) that jointly groups the pallidum with lateral temporal cortices (including the inferior temporal gyrus and temporal pole), consistent with prior evidence that HIV preferentially affects basal ganglia and temporal-limbic systems and is associated with disruptions spanning motor control and higher-order cognition [59, 60]. These patterns align with the modular and hierarchical organization of large-scale brain networks [61] and provide a principled, interpretable partition of ROIs into communities that are directly optimized for the downstream classification objective. To further facilitate qualitative inspection, Figure 5 provides a Circos visualization of ROI relationships, where ROIs are grouped by their dominant community. We reconstruct the connectivity matrices for both groups using the SD3MF approximation, compute edge-wise group differences, and select the top 20 most discriminative connections. All selected connections are validated by a two-sample t -test. The resulting chords illustrate how salient ROI-to-ROI interactions are distributed within and across groups. We observe modality-specific emphases that remain neurobiologically plausible:

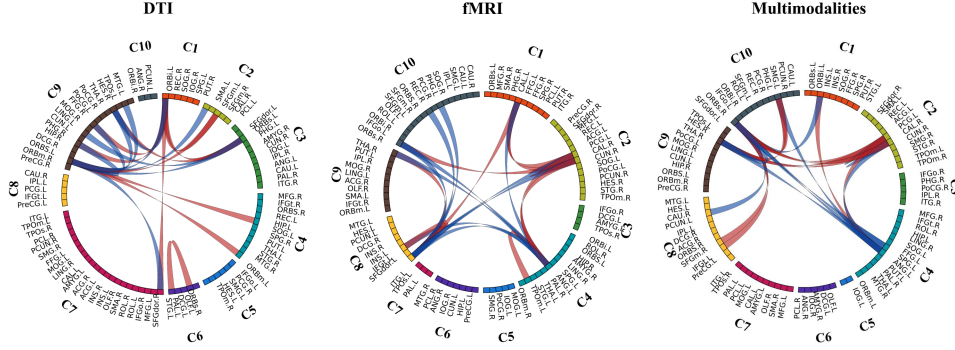


Figure 5: Interpretation of learned communities for the HIV cohort. ROIs are grouped by latent communities (C1 to C10), and ribbons indicate the relationship between ROIs, where red and blue denote relatively high and low loading values, respectively.




Figure 6: An fMRI-derived community (C7) by SD3MF on the HIV dataset: the left pallidum, left temporal pole, and left inferior temporal gyrus.

Figure 7: Surface visualization of modality-specific and multimodal salient ROIs for distinguishing HIV from healthy controls. Results are rendered on lateral, medial, and ventral views.

DTI-derived communities more prominently co-localize subcortical nuclei with frontal and temporal regions, consistent with the known sensitivity of diffusion MRI to fronto-subcortical and white-matter related alterations in HIV [62]; in contrast, fMRI-derived communities exhibit clearer grouping of canonical cortical systems, consistent with resting-state fMRI capturing intrinsic functional network organization [63, 58]. Notably, the multimodal communities preserve these canonical motifs while integrating complementary ROIs across modalities, yielding communities that are simultaneously anatomically coherent and functionally interpretable [64, 65].

Interpreting Salient ROIs: To interpret the salient ROIs for distinguishing HC and HIV, we derive an ROI-level importance score from the learned membership matrices $\{\Psi^{(m)}\}_{m=1}^M$, as shown in Figure 7. Specifically, for each modality m , we compute a scalar saliency score for each ROI by aggregating its loadings across communities on HIV group (i.e., collapsing $\Psi_{p,:}^{(m)}$ to a single value per ROI), and we also compute the weighted score under the multimodal setting by aggregating across modalities. After ranking these ROI scores in descending order, we visualize the top 20 most salient brain regions identified by each modality and by the multimodal analysis. BrainNet Viewer [66] is used to render the top ROIs in lateral, medial, and ventral views of the brain surface for each modality, enabling direct comparison of modality-specific and multimodal ROI saliency patterns.

We found that the salient ROIs for distinguishing HIV and HC span fronto-executive cortices (e.g., SFGdor, MFG, IFGperc/IFGtriang, and orbitofrontal regions), salience/limbic-temporal nodes (insula, PHG, STG, and amygdala), subcortical nuclei, and visual cortices. This pattern is largely consistent with prior HIV neuroimaging evidence showing preferential involvement of frontal and cingulate/insula regions as well as limbic structures in HIV, which has been linked to cognitive decline [67, 59]. Modality-specific findings were also in line with established modality sensitivities: DTI primarily highlighted fronto-subcortical circuits and basal ganglia/thalamic structures, consistent with reports of HIV-associated white-matter and fronto-subcortical integrity loss [62], whereas fMRI emphasized cortical network nodes consistent with dissociated functional abnormalities [68]. Under our setting, these ROIs were jointly confirmed and strengthened across modalities, and the additional emphasis on visual-network regions is supported by HIV-related findings in visual cortices [68].

6 Conclusion

We proposed SD3MF, a supervised deep multimodal matrix factorization that generalizes SNMTF into a deep, end-to-end model for brain network classification. SD3MF learns compact and interpretable community-level representations, while consistently improving accuracy and AUC over strong tensor-based, CNN, and GNN baselines, and having lower variance, across multiple multimodal connectome datasets. Ablation studies further confirm that these gains arise from effective multimodal fusion, increased model depth, and properly balancing reconstruction with supervision. While the model provides interpretable, community-level representations and strong empirical performance, the use of a linear supervised head reflects a trade-off between interpretability and modeling flexibility. Future work will explore more expressive yet interpretable alternatives, such as Kolmogorov-Arnold Networks (KAN) [69]. In addition, extending SD3MF to temporal or longitudinal multimodal networks would enable modeling the evolution of community structure over time. More broadly, SD3MF is applicable to a wider range of multimodal structured data beyond brain networks.

References

- [1] Haoteng Tang, Guixiang Ma, Yanfu Zhang, Kai Ye, Lei Guo, Guodong Liu, Qi Huang, Yalin Wang, Olusola Ajilore, Alex D. Leow, Paul M. Thompson, Heng Huang, and Liang Zhan. A comprehensive survey of complex brain network representation. *Meta-Radiology*, 1(3):100046, 2023.
- [2] Haijing Sun, Anna Wang, and Shanshan He. Temporal and spatial analysis of Alzheimer’s disease based on an improved convolutional neural network and a resting-state fMRI brain functional network. *International Journal of Environmental Research and Public Health*, 19(8): 4508, 2022.
- [3] Rong Zhou, Houliang Zhou, Li Shen, Brian Y. Chen, Yu Zhang, and Lifang He. Integrating multimodal contrastive learning and cross-modal attention for Alzheimer’s disease prediction in brain imaging genetics. In *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 1806–1811, 2023.
- [4] Emine Elif Tulay, Barış Metin, Nevzat Tarhan, and Mehmet Kemal Arıkan. Multimodal neuroimaging: basic concepts and classification of neuropsychiatric diseases. *Clinical EEG and Neuroscience*, 50(1):20–33, 2019.
- [5] Sidong Liu, Weidong Cai, Siqi Liu, Fan Zhang, Michael Fulham, Dagan Feng, Sonia Pujol, and Ron Kikinis. Multimodal neuroimaging computing: a review of the applications in neuropsychiatric disorders. *Brain Informatics*, 2(3):167–180, 2015.
- [6] Adrian Benton, Huda Khayrallah, Biman Gujral, Dee Ann Reisinger, Sheng Zhang, and Raman Arora. Deep generalized canonical correlation analysis. In *Workshop on Representation Learning for NLP*, pages 1–6, 2019.
- [7] Rong Zhou, Houliang Zhou, Brian Y Chen, Li Shen, Yu Zhang, and Lifang He. Attentive deep canonical correlation analysis for diagnosing Alzheimer’s disease using multimodal imaging genetics. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 681–691. Springer, 2023.
- [8] Asif Salim, S.S. Shiju, and S. Sumitra. Design of multi-view graph embedding using multiple kernel learning. *Engineering Applications of Artificial Intelligence*, 90:103534, 2020.
- [9] Yipu Zhang, Li Xiao, Gemeng Zhang, Biao Cai, Julia M. Stephen, Tony W. Wilson, Vince D. Calhoun, and Yu-Ping Wang. Multi-paradigm fMRI fusion via sparse tensor decomposition in brain functional connectivity study. *IEEE Journal of Biomedical and Health Informatics*, 25(5): 1712–1723, 2021.
- [10] Irina Belyaeva, Ben Gabrielson, Yu-Ping Wang, Tony W. Wilson, Vince D. Calhoun, Julia M. Stephen, and Tülay Adalı. Learning spatiotemporal brain dynamics in adolescents via multimodal MEG and fMRI data fusion using joint tensor/matrix decomposition. *IEEE Transactions on Biomedical Engineering*, 71(7):2189–2200, 2024.

- [11] Lifang He, Kun Chen, Wanwan Xu, Jiayu Zhou, and Fei Wang. Boosted sparse and low-rank tensor regression. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [12] Shen Wang, Lifang He, Bokai Cao, Chun-Ta Lu, Philip S Yu, and Ann B Ragin. Structural deep brain network mining. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 475–484, 2017.
- [13] Uğur Demir, Mohammed Amine Gharsallaoui, and Islem Rekik. Clustering-based deep brain multigraph integrator network for learning connectional brain templates. In *International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging*, pages 109–120, 2020.
- [14] Xiaoyi Chen, Pengfei Ke, Yuanyuan Huang, Jing Zhou, Hehua Li, Runlin Peng, Jiayuan Huang, Liqin Liang, Guolin Ma, Xiaobo Li, et al. Discriminative analysis of schizophrenia patients using graph convolutional networks: A combined multimodal MRI and connectomics analysis. *Frontiers in Neuroscience*, 17:1140801, 2023.
- [15] Kai Zhang, Rong Zhou, Eashan Adhikarla, Zhiling Yan, Yixin Liu, Jun Yu, Zhengliang Liu, Xun Chen, Brian D Davison, Hui Ren, et al. A generalist vision–language foundation model for diverse biomedical tasks. *Nature Medicine*, 30(11):3129–3141, 2024.
- [16] Xi Zhang, Lifang He, Kun Chen, Yuan Luo, Jiayu Zhou, and Fei Wang. Multi-view graph convolutional network and its applications on neuroimage analysis for Parkinson’s disease. In *AMIA annual symposium proceedings*, volume 2018, page 1147, 2018.
- [17] Guangqi Wen, Peng Cao, Huiwen Bao, Wenju Yang, Tong Zheng, and Osmar Zaiane. MVS-GCN: A prior brain structure learning-guided multi-view graph convolution network for autism spectrum disorder diagnosis. *Computers in Biology and Medicine*, 142:105239, 2022.
- [18] Gang Qu, Ziyu Zhou, Vince D. Calhoun, Aiyang Zhang, and Yu-Ping Wang. Integrated brain connectivity analysis with fMRI, DTI, and sMRI powered by interpretable graph neural networks. *Medical Image Analysis*, 103:103570, 2025.
- [19] Houliang Zhou, Lifang He, Brian Y. Chen, Li Shen, and Yu Zhang. Multi-modal diagnosis of Alzheimer’s disease using interpretable graph convolutional networks. *IEEE Transactions on Medical Imaging*, 44(1):142–153, 2025.
- [20] Yuefeng Ma, Tengfei Zhang, Zeqi Wu, Xiaochen Mu, Xun Liang, and Lanzhen Guo. Multi-view brain networks construction for Alzheimer’s disease diagnosis. In *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 889–892, 2023.
- [21] Haiping Lu, Konstantinos N. Plataniotis, and Anastasios N. Venetsanopoulos. MPCA: Multilinear principal component analysis of tensor objects. *IEEE Transactions on Neural Networks*, 19(1): 18–39, 2008.
- [22] Yanqiao Zhu, Hejie Cui, Lifang He, Lichao Sun, and Carl Yang. Joint embedding of structural and functional brain networks with graph neural networks for mental illness diagnosis. In *International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 272–276, 2022.
- [23] Weixiang Shao, Lifang He, and Philip S Yu. Clustering on multi-source incomplete data via tensor modeling and factorization. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 485–497. Springer, 2015.
- [24] Ming Yin, Junbin Gao, Shengli Xie, and Yi Guo. Multiview subspace clustering via tensorial t-product representation. *IEEE Transactions on Neural Networks and Learning Systems*, 30(3): 851–864, 2019.
- [25] Guixiang Ma, Lifang He, Chun-Ta Lu, Weixiang Shao, Philip S. Yu, Alex D. Leow, and Ann B. Ragin. Multi-view clustering with graph embedding for connectome analysis. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, page 127–136, 2017.

- [26] Ye Liu, Lifang He, Bokai Cao, Philip S. Yu, Ann B. Ragin, and Alex D. Leow. Multi-view multi-graph embedding for brain network clustering analysis. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [27] Guanghui Li, Qinghua Huang, Chunying Liu, Guanying Wang, Lingli Guo, Ruonan Liu, and Longzhong Liu. Fully automated diagnosis of thyroid nodule ultrasound using brain-inspired inference. *Neurocomputing*, 582:127497, 2024.
- [28] Jeremy Kawahara, Colin J. Brown, Steven P. Miller, Brian G. Booth, Vann Chau, Ruth E. Grunau, Jill G. Zwicker, and Ghassan Hamarneh. Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage*, 146:1038–1049, 2017.
- [29] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2017. URL <https://arxiv.org/abs/1609.02907>.
- [30] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [31] Anees Kazi, Shayan Shekarforoush, S Arvind Krishna, Hendrik Burwinkel, Gerome Vivar, Benedict Wiestler, Karsten Kortüm, Seyed-Ahmad Ahmadi, Shadi Albarqouni, and Nassir Navab. Graph convolution based attention model for personalized disease prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 122–130. Springer, 2019.
- [32] Yanwu Yang, Chenfei Ye, Xutao Guo, Tao Wu, Yang Xiang, and Ting Ma. Mapping multi-modal brain connectome for brain disorder diagnosis via cross-modal mutual learning. *IEEE Transactions on Medical Imaging*, 43(1):108–121, 2024.
- [33] 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.
- [34] Hejie Cui, Wei Dai, Yanqiao Zhu, Xiaoxiao Li, Lifang He, and Carl Yang. Interpretable graph neural networks for connectome-based brain disorder analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 375–385, 2022.
- [35] Liangliang Liu, Yu-Ping Wang, Yi Wang, Pei Zhang, and Shufeng Xiong. An enhanced multi-modal brain graph network for classifying neuropsychiatric disorders. *Medical Image Analysis*, 81:102550, 2022.
- [36] Xiaoxiao Li, Nicha C Dvornek, Yuan Zhou, Juntang Zhuang, Pamela Ventola, and James S Duncan. Graph neural network for interpreting task-fMRI biomarkers. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 485–493, 2019.
- [37] Tong Zhang, Wenming Zheng, Zhen Cui, and Yang Li. Tensor graph convolutional neural network, 2018. URL <https://arxiv.org/abs/1803.10071>.
- [38] Osman Asif Malik, Shashanka Ubaru, Lior Horesh, Misha E Kilmer, and Haim Avron. Tensor graph neural networks for learning on time varying graphs. In *Proceedings of NIPS Workshop*, 2019.
- [39] Zhichao Huang, Xutao Li, Yunming Ye, and Michael K Ng. Mr-gcn: Multi-relational graph convolutional networks based on generalized tensor product. In *IJCAI*, volume 20, pages 1258–1264, 2020.
- [40] Ananya Samanta, Monalisa Sarma, and Debasis Samanta. ALERT: Atlas-based low estimation rank tensor approach to detect Autism spectrum disorder. In *International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 1–4, 2023.
- [41] Yilin Sang and Wan Li. Classification study of Alzheimer’s disease based on self-attention mechanism and dti imaging using gcn. *IEEE Access*, 12:24387–24395, 2024.

- [42] Qiankun Zuo, Ning Zhong, Yi Pan, Huisi Wu, Baiying Lei, and Shuqiang Wang. Brain structure-function fusing representation learning using adversarial decomposed-VAE for analyzing MCI. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:4017–4028, 2023.
- [43] Yanwu Yang, Guoqing Cai, Chenfei Ye, Yang Xiang, and Ting Ma. Tensor-based complex-valued graph neural network for dynamic coupling multimodal brain networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [44] Zhaoming Kong, Rong Zhou, Xinwei Luo, Songlin Zhao, Ann B Ragin, Alex D Leow, and Lifang He. TGNet: tensor-based graph convolutional networks for multimodal brain network analysis. *BioData Mining*, 17(1):55, 2024.
- [45] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.
- [46] Amjad Seyedi and Nicolas Gillis. Encoder-decoder symmetric nonnegative matrix tri-factorization for graph clustering. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 501–505, 2026.
- [47] Brian Karrer and Mark EJ Newman. Stochastic blockmodels and community structure in networks. *Physical Review E*, 83(1):016107, 2011.
- [48] Edoardo M Airoldi, David M Blei, Stephen E Fienberg, and Eric P Xing. Mixed membership stochastic blockmodels. *Journal of Machine Learning Research*, 9(Sep):1981–2014, 2008.
- [49] Akram Hajiveisheh, Seyed Amjad Seyedi, and Fardin Akhlaghian Tab. Deep asymmetric nonnegative matrix factorization for graph clustering. *Pattern Recognition*, 148:110179, 2024.
- [50] George Trigeorgis, Konstantinos Bousmalis, Stefanos Zafeiriou, and Björn W Schuller. A deep matrix factorization method for learning attribute representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(3):417–429, 2016.
- [51] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.
- [52] Joowon Lee, Hanbaek Lyu, and Weixin Yao. Exponentially convergent algorithms for supervised matrix factorization. In *Advances in Neural Information Processing Systems*, volume 36, pages 76947–76959, 2023.
- [53] Joowon Lee, Hanbaek Lyu, and Weixin Yao. Supervised matrix factorization: local landscape analysis and applications. In *International Conference on Machine Learning*, 2024.
- [54] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [55] Pierre De Handschutter and Nicolas Gillis. A consistent and flexible framework for deep matrix factorizations. *Pattern Recognition*, 134:109102, 2023.
- [56] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252, 2019.
- [57] Stephen M Smith, Peter T Fox, Karla L Miller, David C Glahn, P Mickle Fox, Clare E Mackay, Nicola Filippini, Kate E Watkins, Roberto Toro, Angela R Laird, et al. Correspondence of the brain’s functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106(31):13040–13045, 2009.
- [58] BT 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, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 2011.
- [59] Sarah M Israel, Shiva Hassanzadeh-Behbahani, Peter E Turkeltaub, David J Moore, Ronald J Ellis, and Xiong Jiang. Different roles of frontal versus striatal atrophy in HIV-associated neurocognitive disorders. *Human Brain Mapping*, 40(10):3010–3026, 2019.

- [60] Paul M Thompson and Neda Jahanshad. Novel neuroimaging methods to understand how HIV affects the brain. *Current HIV/AIDS Reports*, 12(2):289–298, 2015.
- [61] Olaf Sporns and Richard F Betzel. Modular brain networks. *Annual Review of Psychology*, 67(1):613–640, 2016.
- [62] Talia R Seider, Assawin Gongvatana, Adam J Woods, Huaihou Chen, Eric C Porges, Tiffany Cummings, Stephen Correia, Karen Tashima, and Ronald A Cohen. Age exacerbates HIV-associated white matter abnormalities. *Journal of Neurovirology*, 22(2):201–212, 2016.
- [63] Michael D Fox and Marcus E Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8(9):700–711, 2007.
- [64] Xinhui Jin, Xinyu Liang, and Gaolang Gong. Functional integration between the two brain hemispheres: evidence from the homotopic functional connectivity under resting state. *Frontiers in Neuroscience*, 14:932, 2020.
- [65] Olga Marshall, Jinsoo Uh, Daniel Lurie, Hanzhang Lu, Michael P Milham, and Yulin Ge. The influence of mild carbon dioxide on brain functional homotopy using resting-state fMRI. *Human Brain Mapping*, 36(10):3912–3921, 2015.
- [66] Mingrui Xia, Jinhui Wang, and Yong He. Brainnet viewer: a network visualization tool for human brain connectomics. *PLoS One*, 8(7):e68910, 2013.
- [67] Mingxiang Xu and Xing-Da Ju. A meta-analysis of gray matter volume abnormalities in HIV patients. *Psychiatry Research: Neuroimaging*, 335:111722, 2023.
- [68] Zhong Li, Xingxing Jin, Meng Zhang, Hongxia Wang, Wangyi Liu, Beiran Wang, Baolin Wu, and Xuekun Li. Divergent structural and functional brain alterations in HIV-infected patients: a multimodal meta-analysis. *Frontiers in Neurology*, 16:1618408, 2025.
- [69] Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljagic, Thomas Y. Hou, and Max Tegmark. KAN: Kolmogorov–Arnold networks. In *International Conference on Learning Representations*, 2025.
- [70] Nicolas Gillis. *Nonnegative Matrix Factorization*. SIAM, 2020.
- [71] Sanjeev Arora, Rong Ge, Yoni Halpern, David Mimno, Ankur Moitra, David Sontag, Yichen Wu, and Michael Zhu. Learning topic models—provably and efficiently. *Communications of the ACM*, 61(4):85–93, 2018.
- [72] Kejun Huang, Xiao Fu, and Nikolaos D Sidiropoulos. Anchor-free correlated topic modeling: Identifiability and algorithm. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- [73] Xiao Fu, Kejun Huang, Nicholas D Sidiropoulos, Qingjiang Shi, and Mingyi Hong. Anchor-free correlated topic modeling. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 41(5):1056–1071, 2018.
- [74] Ann B. Ragin, Hongyan Du, Renee Ochs, Ying Wu, Christina L. Sammet, Alfred Shoukry, and Leon G. Epstein. Structural brain alterations can be detected early in hiv infection. *Neurology*, 79(24):2328–2334, 2012.
- [75] Bokai Cao, Xiangnan Kong, Jingyuan Zhang, Philip S Yu, and Ann B Ragin. Identifying HIV-induced subgraph patterns in brain networks with side information. *Brain Informatics*, 2(4):211–223, 2015.
- [76] Olusola Ajilore, Nathalie Vizueta, Patricia Walshaw, Liang Zhan, Alex Leow, and Lori L. Altshuler. Connectome signatures of neurocognitive abnormalities in euthymic bipolar I disorder. *Journal of Psychiatric Research*, 68:37–44, 2015.
- [77] Susan Whitfield-Gabrieli and Alfonso Nieto-Castanon. Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. *Brain Connectivity*, 2(3):125–141, 2012.

- [78] Liang Zhan, Jiayu Zhou, Yalin Wang, Yan Jin, Neda Jahanshad, Gautam Prasad, Talia M Nir, Cassandra D Leonardo, Jieping Ye, Paul M Thompson, et al. Comparison of nine tractography algorithms for detecting abnormal structural brain networks in Alzheimer’s disease. *Frontiers in Aging Neuroscience*, 7:48, 2015.

Appendix

A Optimization model

In this section, we discuss further our optimization model. Collecting all parameters into $\theta = \{\mathbf{W}_l^{(m)}\}_{l,m}, \{\mathbf{S}_i\}_{i=1}^N, \beta, \alpha$, and we denote $\Psi^{(m)} = \prod_{l=1}^L \mathbf{W}_l^{(m)}$. Given the matrices $\mathbf{A}_i^{(m)}$ ’s and the labels y_i ’s, the SD3MF optimization model is

$$\min_{\theta \in \mathcal{C}} F(\theta) := \underbrace{\sum_{i=1}^N \sum_{m=1}^M \mu \|\mathbf{A}_i^{(m)} - \Psi^{(m)} \mathbf{S}_i \Psi^{(m)\top}\|_F^2}_{F_{\text{rec}}(\theta)} + \underbrace{\sum_{i=1}^N \ell(y_i, \beta^\top v_i)}_{F_{\text{cls}}(\theta)}, \quad (6)$$

where $v_i = \sum_{m=1}^M \alpha^{(m)} \text{vec}(\Psi^{(m)\top} \mathbf{A}_i^{(m)} \Psi^{(m)})$ is the fused vector, and the feasible set is

$$\mathcal{C} = \left\{ \theta \mid W_l^{(m)} \geq 0, \Psi^{(m)} = \prod_{l=1}^L W_l^{(m)}, \Psi^{(m)} \mathbf{1} = \mathbf{1}, \sum_{m=1}^M \alpha_m = 1, \alpha_m \geq 0 \right\}.$$

This is a non-smooth and highly non-convex problem (both the objective function and the feasible set are non-convex). This makes a theoretical convergence analysis of stochastic gradient descent hard and out of the scope of this paper. Although (6) is highly non convex, it exhibits partial convexity with respect to some blocks of variables, namely the \mathbf{S}_i ’s, β and α . This could be leveraged in further study of optimization algorithms and convergence behavior.

In any case, our experimental results show that the proposed optimization strategy converges on all considered data sets; see Section A.3 and Figure 8. In the following, we discuss some of aspects of this optimization model.

A.1 Effect of Non-negativity Constraints

The constraint $\Psi^{(m)} \geq 0$ provides improved identifiability; this follows from the vast literature on NMF [70, Chapter 4]. For example, we have the following results.

Theorem 1 (Identifiability under separability, [71]). *Consider $\mathbf{A} = \Psi \mathbf{S} \Psi^\top$ with $\Psi \in \mathbb{R}_+^{n \times r}$, $\text{rank}(\mathbf{S}) = r$, and Ψ is separable, that is, Ψ contains a non-singular r -by- r diagonal submatrix. Any factorization of \mathbf{A} of the same form will recover Ψ and \mathbf{S} , up to permutations and scalings of the columns of Ψ and counter scaling of rows and columns of \mathbf{S} .*

In brain network terms, separability means that each community has at least one anchor region that belongs exclusively to that community. This is a mild and neurobiologically plausible assumption, since many brain regions are strongly associated with a single functional network [58]. We note that an alternative identifiability route relies on the sufficiently scattered condition (which relaxes separability, and requires some degree of sparsity in Ψ) combined with volume minimization of \mathbf{S} [72, 73]; this would be an interesting direction of further research: use volume minimization for relaxed identifiability conditions.

The deep product $\Psi^{(m)} = \prod_l \mathbf{W}_l^{(m)}$ introduces additional factorization ambiguity between layers: the individual $\mathbf{W}_l^{(m)}$ are not uniquely identifiable [55]. However, the non-negativity of each $\mathbf{W}_l^{(m)}$ ensures that $\Psi^{(m)}$ itself remains nonnegative, so the identifiability of $\Psi^{(m)}$ under separability is preserved. This eliminates the sign and rotational ambiguities present in unconstrained deep factorizations at the level of the end-to-end mapping, which is the quantity that determines both the reconstruction and the classifier output.




Figure 8: Convergence behavior of the proposed SD3MF with learning rate 10^{-5} on three datasets showing reconstruction loss, classification loss, and training accuracy.

A.2 Landscape Smoothing via Over-parameterization

Deep factorization acts as over-parameterization that smooths the loss landscape. For the reconstruction component, define the per-subject loss

$$f_i(\Psi, \mathbf{S}_i) = \|\mathbf{A}_i - \Psi \mathbf{S}_i \Psi^\top\|_F^2.$$

With $\Psi = \mathbf{W}_1 \mathbf{W}_2 \mathbf{W}_3$, the Hessian with respect to $(\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3)$ at a critical point has fewer negative eigenvalues than the Hessian of f_i with respect to a single matrix Ψ [55, 56]. Informally, the product parameterization “stretches” narrow valleys into wider basins, making it easier for gradient descent to escape saddle points and poor local minima.

Note that in principle, the shallow model ($L = 1$) has the same expressive power as the deep model for classification, since the classifier operates on $\Psi^\top \mathbf{A}_i \Psi$ regardless of whether Ψ is parameterized as a single matrix or a product of factors. The consistent empirical improvement of the deep model over the shallow variant, see Figure 4, therefore provides direct evidence that depth improves optimization rather than expressiveness: the over-parameterized product structure helps gradient descent reach better stationary points by smoothing the loss landscape.

These results do not guarantee global optimality—we are dealing with a non-convex problem most likely NP-hard—but they collectively justify the stable convergence observed empirically and confirm that the architectural choices in SD3MF (deep parameterization, small initialization, non-negativity, reconstruction regularization) are theoretically grounded.

A.3 Empirical convergence

Figure 8 shows the evolution of the training metrics of SD3MF on the HIV, BP, and PPMI datasets. SD3MF exhibits stable convergence under SGD with a small learning rate (10^{-5}) and small-scale (near-zero) initialization, which provides implicit regularization by biasing deep MF toward low-rank, well-conditioned solutions without explicit penalties. During training, the classification loss decreases steadily and training accuracy rises and then plateaus, indicating reliable optimization of the supervised objective. The reconstruction loss initially drops but later shows a mild increase, consistent with the joint objective: as the model becomes more label-discriminative, it can sacrifice reconstruction fidelity to better align the learned community interactions with the classification task. Overall, the curves suggest controlled, non-divergent training dynamics where small initialization and gentle updates regularize learning while the optimizer navigates the reconstruction–prediction balance.

B Model Efficiency

In the deep formulation of SD3MF, each modality is parameterized by a product of three nonnegative factor matrices $\mathbf{W}_1 \in \mathbb{R}^{n \times r_1}$, $\mathbf{W}_2 \in \mathbb{R}^{r_1 \times r_2}$, and $\mathbf{W}_3 \in \mathbb{R}^{r_2 \times r_3}$. With $(r_1, r_2, r_3) = (30, 20, 10)$, the number of parameters per modality is $nr_1 + r_1r_2 + r_2r_3 = 30n + 800$, which scales linearly with the number of nodes and remains independent of graph convolution depth or feature dimensionality. The intermediate ranks r_1 and r_2 are chosen to enable a gradual hierarchical dimensionality reduction, while the deepest rank $r_3 = 10$ reflects the assumed number of communities and aligns with evidence

that brain functional architecture comprises approximately ten major functional networks [57]. In addition, the subject-specific interaction matrices contribute only $Nr_3^2 = 100N$ parameters, and the final linear classifier introduces a negligible number of additional weights. In contrast, GNN-based multimodal architectures employ multiple graph convolution layers with weight matrices of size $O(d_{in}d_{out})$ per layer, together with parameters for neighborhood aggregation, modality pooling, and fully connected classifiers, resulting in substantially higher parameter counts and memory consumption. Moreover, repeated message-passing operations incur additional runtime and memory overhead proportional to the number of edges and layers. By relying on low-rank matrix factorizations and closed-form projections, SD3MF avoids explicit neighborhood aggregation and yields a compact architecture with favorable parameter efficiency and computational complexity.

C Detailed Dataset Description

This section provides detailed information on the datasets used in our experiments, including cohort composition, imaging acquisition protocols, preprocessing pipelines, brain parcellation schemes, and network construction procedures.

Human Immunodeficiency Virus Infection (HIV): Data were collected in the Early HIV Infection Study at Northwestern University, including both fMRI and DTI [74]. We used a standard preprocessing workflow [75]. fMRI processing (via DPARSF toolbox) included realignment to the first volume, slice-timing correction, normalization to the MNI template, and spatial smoothing with an 8-mm Gaussian kernel. Subject-specific whole-brain networks were then constructed by parcellating the brain into 90 cerebral regions (excluding 26 cerebellar regions) and computing pairwise connectivity using correlation coefficients.

Bipolar Disorder (BP): This dataset comes from the UCLA Ahmanson-Lovelace Brain Mapping Center and includes 52 euthymic bipolar I participants and 45 age- and gender-matched healthy controls, with both fMRI and DTI modalities [76]. Resting-state fMRI was acquired on a 3T Siemens Trio scanner using a T2* EPI gradient-echo sequence with IPAT; DTI was acquired on a Siemens 3T Trio scanner. Brain networks were generated using the functional connectivity (CONN) toolbox [77]. Raw EPI data were realigned and co-registered, followed by normalization and smoothing. Nuisance effects from motion, white matter, and CSF were regressed out, and connectivity matrices were computed from pairwise BOLD correlations across 82 FreeSurfer-derived cortical/subcortical gray-matter regions.

Parkinson’s Progression Markers Initiative (PPMI): This dataset is derived from the Parkinson’s Progression Markers Initiative (PPMI), a large collaborative study aimed at improving Parkinson’s disease therapeutics. We use DTI data from 754 subjects. The diffusion data are corrected for head motion and eddy-current distortions, skull-stripped, and linearly registered using FreeSurfer following [78]. We parcellated 84 ROIs from the T1 images with FreeSurfer, and using these ROIs we built three structural connectivity matrices per subject via whole-brain probabilistic tractography: Probabilistic Index of Connectivity (PICO), Hough voting, and Probtracx [78].

D Implementation Details

We implemented SD3MF in PyTorch and optimized all parameters end-to-end using stochastic gradient descent with a learning rate of 10^{-5} and a maximum of 30,000 iterations. For each modality, the deep node-to-community mapping $\Psi^{(m)}$ was parameterized by a three-layer factorization with hidden widths ([30, 20, 10]), and nonnegativity was enforced through ReLU activations followed by row normalization. All deep factor matrices were initialized using orthogonal initialization scaled to small magnitudes (10^{-3}) to stabilize early training and induce implicit regularization. The shared interaction matrices S_i were initialized from view-averaged projections using the initial bases and symmetrized at each iteration. Modality weights $\alpha^{(m)}$ were learned via a softmax over trainable logits to ensure a simplex constraint. The linear classifier head was initialized with Xavier initialization at a small gain (10^{-3}) to balance optimization dynamics across modules. All experiments were implemented in PyTorch (Python) and executed on a laptop with an NVIDIA RTX 4060 GPU with 8 GB VRAM; the model can also be trained and evaluated on CPU with comparable behavior but longer runtimes.

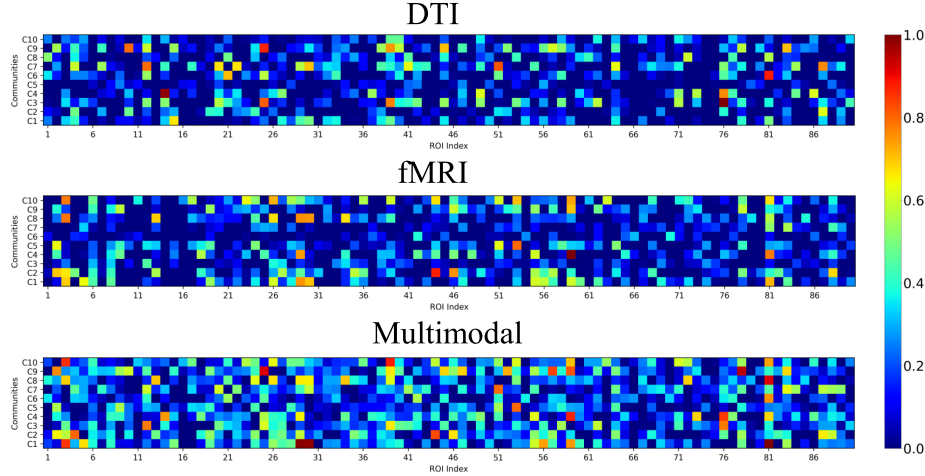


Figure 9: Visualization of the learned membership matrices $\Psi^{(m)}$ for DTI, fMRI, and Multimodal settings on the HIV dataset.




Figure 10: Approximation matrices $\hat{\mathbf{A}}_{\text{HIV}}^{(m)} = \Psi^{(m)} \bar{\mathbf{S}}_{\text{HIV}} \Psi^{(m)\top}$ in ROI space for the HIV cohort under DTI, fMRI, and multimodal settings. Entries represent approximated ROI-to-ROI connectivity strengths. ROIs are reordered by dominant community assignments (C1–C10) shown on the axes for visualization.

E Learned Membership Matrix

Figure 9 visualizes the learned membership matrices $\Psi^{(m)}$ for DTI, fMRI, and the multimodal setting on the HIV cohort. Each row corresponds to a latent community, and each column corresponds to an ROI, with color intensity indicating the strength of ROI-to-community membership. This representation provides a compact view of how SD3MF organizes ROIs into modality-specific and multimodal community structures. We observe that the learned memberships exhibit sparse and structured activation patterns, where subsets of ROIs consistently show strong loadings within specific communities, suggesting the emergence of coherent functional and anatomical modules rather than diffuse or uniform assignments. Compared to single-modality settings, the multimodal membership matrix shows denser and more balanced activation across communities, indicating complementary integration of modality-specific information and enhanced community expressiveness. These patterns support that the proposed hierarchical factorization captures meaningful modular organization of the brain connectome and provides an interpretable latent representation that is subsequently leveraged for downstream classification.

F Connectivity Analysis via Approximation Matrix

To inspect the ROI-to-ROI connectivity captured by SD3MF on the HIV dataset, we visualize the approximation matrices reconstructed in the original ROI space. For each modality in HIV group, we

compute the group-averaged interaction matrix $\bar{\mathbf{S}}_{\text{HIV}}$ and reconstruct the connectivity approximation as

$$\hat{\mathbf{A}}_{\text{HIV}}^{(m)} = \mathbf{\Psi}^{(m)} \bar{\mathbf{S}}_{\text{HIV}} \mathbf{\Psi}^{(m)\top}. \quad (7)$$

Here, $\hat{\mathbf{A}}_{\text{HIV}}^{(m)}$ represents the SD3MF-induced approximation of HIV-related ROI connectivity. In Figure 10, ROIs are reordered according to their dominant community assignments (shown on the axes) for visualization.

As shown in Figure 10, the approximation matrices exhibit structured block-like patterns rather than diffuse connectivity, indicating coherent subnetwork organization captured by the model. The DTI approximation shows smoother and more globally distributed patterns, while the fMRI approximation presents more heterogeneous and localized structures. The multimodal result integrates both characteristics and yields more coherent connectivity patterns, demonstrating the benefit of combining complementary structural and functional information.