Hyperbolic Kernel GCN with Structure-Function Connectivity Coupling for Neurocognitive Impairment Analysis

Meimei Yang[1], Yongheng Sun[1], Qianqian Wang[1], Wei Wang[2], Hong-Jun Li[2], and Mingxia Liu[1*]

[1] Department of Radiology and Biomedical Research Imaging Center, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA
[2] Department of Radiology, Beijing Youan Hospital, Capital Medical University, Beijing 100069, China
*Corresponding author (mingxia_liu@med.unc.edu)

Abstract. Diffusion tensor imaging (DTI) and functional MRI (fMRI) provide complementary views of the brain by revealing the physical structure connectivity (SC) between brain regions and functional connectivity (FC) between those regions during neural processing. Previous evidence has shown that fusing the two modalities facilitates the identification of abnormal connectivity associated with neurocognitive disorders. However, existing fusion approaches are generally performed in Euclidean space and thus cannot effectively capture the intrinsic hierarchical organization of structural/functional brain networks. To this end, we propose a novel hyperbolic kernel graph convolutional network with SC-FC Coupling (HKC) for neurocognitive impairment analysis. The HKC consists of a *hyperbolic kernel graph convolutional network* for extracting local-to-global features from DTI and fMRI, an *SC-FC coupling module* that models global SC-FC interactions based on encoded DTI and fMRI features, and a *hyperbolic neural network predictor* for classification. Our HKC captures both local and global dependencies among structurally and functionally connected brain regions while preserving the hierarchical organization of brain networks. We evaluate HKC on paired DTI and fMRI data from 68 individuals with HIV-associated asymptomatic neurocognitive impairment and 69 healthy controls, with experimental results suggesting its superiority over state-of-the-art methods. Additionally, HKC identifies key SC-FC patterns in ANI, highlighting the visual network and fronto-cerebellar connections as critical biomarkers.

Keywords: Hyperbolic kernel · Connectivity · Cognitive impairment

1 Introduction

Diffusion tensor imaging (DTI) and functional MRI (fMRI) provide complementary views of the brain [24], as they reveal the structure connectivity (SC) between brain regions and the function connectivity (FC) among those regions

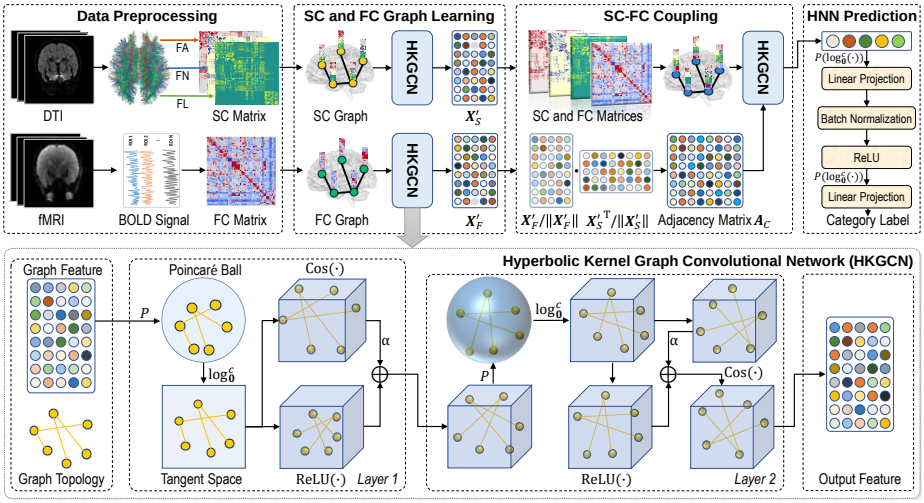

Fig. 1. Illustration of the proposed HKC framework, consisting of (1) a hyperbolic kernel graph convolutional network (HKGCN) to extract features from fMRI and DTI, (2) an SC-FC coupling module for feature fusion by capturing global SC-FC interactions, and (3) a hyperbolic neural network (HNN) for prediction.

during neural processing. Existing studies have shown that integrating fMRI and DTI can improve the identification of abnormal connectivity associated with neurocognitive disorders. However, current fusion methods are generally formulated in Euclidean space, which struggles to capture the intrinsic non-Euclidean hierarchical nature of structural or functional brain networks, limiting their ability to model complex cross-modality interactions.

Extensive evidence has shown that the brain exhibits a hierarchical organization, spanning both structure and function connectivity networks [23, 30]. Most existing methods are designed in Euclidean space, causing distortion when embedding hierarchical data [20]. Hyperbolic space, with its negative curvature and exponential volume growth, provides a theoretically grounded and low-distortion solution for encoding brain SC and FC networks [12]. Nevertheless, existing methods [38] primarily rely on hyperbolic graph neural network (HGCN) for feature encoding and fusion, with high computational complexity due to exponential and logarithmic mappings in the hyperbolic space [3]. Moreover, these methods usually fail to explicitly capture local-to-global dependencies among structurally and functionally connected brain regions.

To this end, we introduce a novel Hyperbolic Kernel graph convolutional network with SC-FC Coupling (HKC) for automated neurocognitive impairment analysis. HKC consists of three key components: (1) a hyperbolic kernel graph convolutional network (HKGCN) to extract features from fMRI and DTI, (2) an SC-FC coupling module to capture global SC-FC interactions based on HKGCN-encoded features, and (3) a hyperbolic neural network predictor for classification. Specifically, we design HKGCN by incorporating hyperbolic ker-

nels to avoid complex hyperbolic operations while modeling complex hierarchical structures in brain networks. Furthermore, this hyperbolic kernel enables us to capture the local-to-global interactions within brain networks. Additionally, the SC-FC coupling module leverages the HKGCN-extracted SC and FC features to explicitly model the global dependencies between them. This enables a more precise characterization of SC-FC interactions, improving the interpretability of neurocognitive analysis. To the best of our knowledge, this is among the first attempts to integrate hyperbolic kernels with GCNs for neurocognitive impairment analysis. Experiments on 137 subjects with paired DTI and fMRI data validate the superiority of HKC over several state-of-the-art (SOTA) methods.

2 Methodology

Data and Image Preprocessing. An HIV-associated neurocognitive impairment dataset is involved, with paired resting-state fMRI and DTI from 68 individuals with HIV-associated asymptomatic neurocognitive impairment (ANI) and 69 healthy controls (HCs). DTI is processed using PANDA toolbox [7] with a standard pipeline: brain extraction, bias field correction, eddy-current and motion correction, diffusion gradient strength adjustment, and co-registration of the anatomical image to diffusion space with AAL atlas to parcellate the brain into 116 regions-of-interest (ROIs). Three metrics are extracted: fiber number (FN), fractional anisotropy (FA), and fiber length (FL). We construct an SC graph for each subject, with each node represented by concatenated FN, FA, and FL features, and each edge denoted as the sum of the three metrics. The fMRI data is preprocessed using DPARSF [36], including removing the first 10 volumes for magnetization equilibrium, motion correction, bandpass filtering (0.01–0.1 Hz), nuisance regression, and MNI normalization. Using the AAL atlas, the mean time series for each ROI is extracted, and the FC graph is then constructed based on Pearson correlation (PC) coefficients between regional time series, with both node features and the adjacency matrix derived from the PC matrix.

Problem Formulation. We formulate DTI and fMRI feature fusion as a cross-modality graph embedding learning problem. This is achieved by integrating SC graphs derived from DTI and FC graphs from fMRI to capture local-to-global interactions among ROIs, with each ROI corresponding to a specific node. Denote an SC graph as $G_S = (\boldsymbol{A}_S, \boldsymbol{X}_S)$ and an FC graph as $G_F = (\boldsymbol{A}_F, \boldsymbol{X}_F)$, where $\boldsymbol{X}_S \in \mathbb{R}^{116\times 348}$ and $\boldsymbol{X}_F \in \mathbb{R}^{116\times 116}$ are feature matrices, and $\boldsymbol{A}_S \in \mathbb{R}^{116\times 116}$ and $\boldsymbol{A}_F \in \mathbb{R}^{116\times 116}$ are adjacency matrices. We aim to train a model that learns their cross-modality interactions via an *SC-FC coupling graph* $G_C = (\boldsymbol{A}_C, \boldsymbol{X}_C)$ to generate fused graph embeddings, followed by a predictor for disease identification.

Proposed Method. As shown in Fig. 1, the proposed HKC framework consists of three key components: (1) feature extraction using HKGCN, which leverages hyperbolic kernels to learn hierarchical representations of brain SC and FC

graphs; (2) SC-FC coupling for modeling global interactions between SC and FC graphs, thereby improving biological interpretability and multimodal data fusion; and (3) prediction via a new hyperbolic neural network (HNN).

(1) ***Feature Extraction with HKGCN***. To effectively model the hierarchical structures of DTI and fMRI data, we design a novel HKGCN for multimodality feature learning. The HKGCN innovatively integrates GCN with *hyperbolic kernels* to encode DTI and fMRI data in hyperbolic space. The Poincaré ball defined as $\mathbb{B}_c^n = \{\boldsymbol{x} \in \mathbb{R}^n : c\|\boldsymbol{x}\|^2 < 1, c > 0\}$, is widely used in hyperbolic models, where $-c$ denotes the negative curvature [9, 12]. While the Poincaré ball can encode hierarchical structures, conventional hyperbolic GCN (HGCN) methods rely on Möbius operations that are computationally intensive [3]. To address this issue, we introduce hyperbolic kernels based on the tangent space (Euclidean space) of the Poincaré ball using logarithmic map [9, 12], defined as:

$$\log_{\mathbf{0}}^{c}(\boldsymbol{x}) = \tanh^{-1}(\sqrt{c}\,\|\boldsymbol{x}\|)\,\frac{\boldsymbol{x}}{\sqrt{c}\,\|\boldsymbol{x}\|}. \tag{1}$$

Using this map, our HKGCN is capable of effectively reducing computational complexity by avoiding Möbius transformations of the Poincaré ball.

Inspired by classic kernels [6, 27], we adopt two hyperbolic kernels, *i.e.*, hyperbolic radial basis function (HRBF) kernel and hyperbolic arc-cos (HAC) kernel, to capture local and global geometric structures. Let $p(\boldsymbol{w})$ represent the probability density function. These two kernels are defined as:

$$k^{\mathrm{HRBF}}(\boldsymbol{x}_i, \boldsymbol{x}_j) = \int e^{j\boldsymbol{w}^\top(\log_{\mathbf{0}}^{c}(\boldsymbol{x}_i) - \log_{\mathbf{0}}^{c}(\boldsymbol{x}_j))} p(\boldsymbol{w}) d\boldsymbol{w}, \tag{2}$$

$$k^{\mathrm{HAC}}(\boldsymbol{x}_i, \boldsymbol{x}_j) = 2\int \phi^{\mathrm{HAC}}(\boldsymbol{x}_i, \boldsymbol{w})\phi^{\mathrm{HAC}}(\boldsymbol{x}_j, \boldsymbol{w}) p(\boldsymbol{w}) d\boldsymbol{w}. \tag{3}$$

Here, the feature mapping function is defined as:

$$\phi^{\mathrm{HAC}}(\boldsymbol{x}, \boldsymbol{w}) = \Theta(\boldsymbol{w}^\top \log_{\mathbf{0}}^{c}(\boldsymbol{x}) + b)(\boldsymbol{w}^\top \log_{\mathbf{0}}^{c}(\boldsymbol{x}) + b), \tag{4}$$

with $\Theta(\cdot) = \frac{1}{2}(1 + \mathrm{sign}(.))$. By substituting Θ into Eq. (4), we can obtain $\phi^{\mathrm{HAC}}(\boldsymbol{x}, \boldsymbol{w}) = \mathrm{ReLU}(\boldsymbol{w}^\top \log_{\mathbf{0}}^{c}(\boldsymbol{x}) + b)$. Utilizing random Fourier features [27], the two kernels can be approximated as:

$$\Phi^{\mathrm{HRBF}}(\boldsymbol{x}) = \sqrt{\frac{2}{D}}\cos(\boldsymbol{W}^\top \log_{\mathbf{0}}^{c}(\boldsymbol{x}) + \boldsymbol{b}),\ \Phi^{\mathrm{HAC}}(\boldsymbol{x}) = \sqrt{\frac{2}{D}}\mathrm{ReLU}(\boldsymbol{W}^\top \log_{\mathbf{0}}^{c}(\boldsymbol{x}) + \boldsymbol{b}), \tag{5}$$

where $\boldsymbol{W} = [\boldsymbol{w}_1, \ldots, \boldsymbol{w}_D]$ and $\boldsymbol{b}$ are trainable parameters, and D is kernel mapping dimension. Given a graph $G{=}(\boldsymbol{A}, \boldsymbol{X})$ in Euclidean space, we integrate hyperbolic kernels into GCN [18] for feature learning, yielding a single-layer HKGCN:

$$\tilde{\boldsymbol{X}} = \log_{\mathbf{0}}^{c}(P(\boldsymbol{X})), \quad \Phi(\boldsymbol{A}, \boldsymbol{X}) = \mathrm{ReLU}(\hat{\boldsymbol{A}}\tilde{\boldsymbol{X}}\boldsymbol{W} + \boldsymbol{b}) + \alpha\cos(\hat{\boldsymbol{A}}\tilde{\boldsymbol{X}}\boldsymbol{W}' + \boldsymbol{b}'), \tag{6}$$

where $\boldsymbol{W}$ and $\boldsymbol{W}'$ are trainable weight matrices, $\boldsymbol{b}$ and $\boldsymbol{b}'$ are trainable bias terms, and $\hat{\boldsymbol{A}}$ is the normalized adjacency matrix. The projection function $P(\boldsymbol{X}) = \frac{\boldsymbol{X}}{\sqrt{c}\|\boldsymbol{X}\|} - \epsilon$ if $\sqrt{c}\|\boldsymbol{X}\| \geq 1$, otherwise $P(\boldsymbol{X}) = \boldsymbol{X}$, ensures that the data is inside the Poincaré ball. Here, $\mathrm{ReLU}(\cdot)$ and $\cos(\cdot)$ are applied to aggregate

neighborhood features in SC or FC graphs. ReLU(·) approximates the HAC kernel, capturing global angular similarity in the tangent space and enabling the model to learn *global nonlinear relationships*. And cos(·) approximates the HRBF kernel, which decays rapidly with geodesic distance and captures *local interactions*. The hyperparameter α balances the contributions of the two kernels. We can construct a multi-layer HKGCN Ψ by stacking multiple layers. And HKGCN-encoded SC and FC graph embeddings are denoted as $\boldsymbol{X}'_S = \Psi(\boldsymbol{A}_S, \boldsymbol{X}_S)$ and $\boldsymbol{X}'_F = \Psi(\boldsymbol{A}_F, \boldsymbol{X}_F)$. By integrating hyperbolic kernels, our HKGCN leverages kernel-based feature transformation to aggregate both local and global information from neighboring nodes at each layer. Meanwhile, its multi-layer architecture progressively preserves local-to-global interactions in brain networks, while hyperbolic kernels naturally model their intrinsic hierarchical organization.

(2) ***SC-FC Coupling***. To model global interactions between SC and FC graphs, we create a new *SC-FC coupling graph* G_C through an inner-product operation between normalized SC and FC graph embeddings, defined as $\boldsymbol{A}_C = \frac{\boldsymbol{X}'_F}{\|\boldsymbol{X}'_F\|} \cdot \frac{\boldsymbol{X}'^{\top}_S}{\|\boldsymbol{X}'_S\|}$, where $\boldsymbol{A}_C \in \mathbb{R}^{116\times116}$ is the adjacency matrix of the SC-FC coupling graph, encoding global SC-FC interactions across ROIs. To represent each node/ROI in this coupling graph, we concatenate the learned SC and FC features for each node to obtain the new node feature $\boldsymbol{X}_C = [\boldsymbol{X}_S, \boldsymbol{X}_F]$. We then apply multi-layer HKGCN Ψ to obtain fused representation $\boldsymbol{X}_{\mathrm{CP}} = \Psi(\boldsymbol{A}_C, \boldsymbol{X}_C)$. By constructing the SC-FC coupling graph, we can explicitly capture SC-FC interactions. Compared to traditional statistical approaches (*e.g.*, Pearson correlation [16]) that rely on predefined assumptions, we construct a data-driven SC-FC coupling graph, while the integration of HKGCN and SC-FC coupling allows for effective modeling of hierarchical structures of brain networks [8, 15].

(3) ***Prediction with HNN***. Given the feature $\boldsymbol{X}_{CP} \in \mathbb{R}^{116\times64}$ extracted from the SC-FC coupling module, we average it into a vector $\boldsymbol{X}_H \in \mathbb{R}^{64}$. To align with the hyperbolic nature of HKGCN-extracted representations, we design HNN for classification. Unlike conventional predictors that apply fully connected layers directly to GCN-extracted features, HNN operates in the tangent space to maintain geometric consistency and hierarchical properties [9] learned by HKGCN. In each layer, we apply the following transformation:

$$\boldsymbol{X}_{\mathrm{HF}} = \delta(\boldsymbol{W}^{\top} \log_{\mathbf{0}}^{c}(P(\boldsymbol{X}_H)) + \boldsymbol{b}), \tag{7}$$

where $\delta(\cdot)$ is the activation function (ReLU), and $P(.)$ ensures the input remains in the Poincaré ball. We stack two layers of this transformation. The final representation $\boldsymbol{X}_{\mathrm{HF}}$ is then fed into a softmax layer for classification with a cross-entropy loss. Compared to existing hyperbolic networks [12], our HNN decreases the computational cost by avoiding complex transformations in hyperbolic space.

Implementation. In the proposed HKC, the HKGCN consists of two hyperbolic kernel layers, both equipped with a hidden dimension of 64 and a rescaling factor $\alpha = 0.01$. The HNN contains two layers (each with a hidden dimension of 32). Both HKGCN and HNN utilize the curvature parameter $c = 10$. The Adam optimizer is used with a learning rate of 0.0002 and a weight decay of 1×10^{-4}. The batch size is set to 128, with a training epoch of 50.

Table 1. Results (%) of different methods in ANI vs. HC classification. '*' denotes that the HKC and a competing method are significantly different ($p < 0.05$ via t-test).

Method	AUC	ACC	F1	BAC	SEN	SPE	PRE	p-value
SVM [26]	$54.55^{*}_{4.04}$	$56.66_{2.86}$	$56.94_{2.56}$	$57.66_{2.56}$	$58.03_{4.19}$	$57.21_{6.74}$	$58.93_{3.39}$	< 0.001
RF [1]	$54.92^{*}_{3.19}$	$50.64_{2.47}$	$51.03_{5.15}$	$53.52_{2.62}$	$55.45_{7.31}$	$51.60_{6.21}$	$53.58_{4.64}$	< 0.001
XGBoost [4]	$61.06^{*}_{3.19}$	$58.09_{1.72}$	$57.56_{2.61}$	$58.84_{1.11}$	$58.42_{3.88}$	$59.27_{4.23}$	$59.54_{2.12}$	< 0.001
LSTM [14]	$64.68^{*}_{1.61}$	$64.67_{1.95}$	$62.49_{5.11}$	$65.40_{1.52}$	$63.98_{8.25}$	$66.81_{5.67}$	$68.03_{1.88}$	0.002
STGCN [37]	$53.94^{*}_{3.74}$	$52.50_{2.08}$	$52.18_{0.86}$	$53.20_{2.17}$	$55.23_{3.19}$	$51.18_{7.33}$	$52.56_{2.13}$	< 0.001
GCN [18]	$62.94^{*}_{2.58}$	$62.06_{3.23}$	$60.83_{1.98}$	$62.27_{3.53}$	$61.24_{4.66}$	$63.30_{9.48}$	$63.92_{4.14}$	< 0.001
GAT [31]	$67.95^{*}_{1.36}$	$62.06_{3.23}$	$60.83_{1.98}$	$62.27_{3.53}$	$61.24_{4.66}$	$63.30_{9.48}$	$63.92_{4.14}$	0.030
GraphSAGE [13]	$64.26^{*}_{2.89}$	$57.99_{4.64}$	$56.50_{4.33}$	$58.64_{3.46}$	$58.64_{12.77}$	$58.63_{18.32}$	$60.68_{7.48}$	< 0.001
GIN [35]	$66.30_{4.47}$	$59.41_{4.21}$	$59.88_{3.92}$	$60.33_{1.92}$	$63.07_{6.06}$	$57.59_{9.80}$	$61.86_{5.10}$	0.011
HGCN [3]	$63.71^{*}_{2.77}$	$65.85_{1.06}$	$62.49_{5.57}$	$65.21_{0.76}$	$60.70_{11.78}$	$69.73_{12.68}$	$69.13_{7.47}$	< 0.001
BrainNetCNN [17]	$63.34^{*}_{1.70}$	$66.43_{4.46}$	$66.16_{5.14}$	$65.75_{4.61}$	$69.76_{7.64}$	$61.74_{7.09}$	$68.62_{4.34}$	< 0.001
BrainGNN [19]	$56.43^{*}_{4.22}$	$51.41_{4.13}$	$52.07_{4.20}$	$52.43_{4.00}$	$54.69_{4.67}$	$50.17_{6.12}$	$52.87_{4.92}$	< 0.001
HKC (Ours)	$\mathbf{70.56}_{1.61}$	$\mathbf{68.65}_{2.07}$	$\mathbf{66.71}_{2.87}$	$\mathbf{68.46}_{1.55}$	$\mathbf{66.65}_{6.67}$	$\mathbf{70.28}_{7.76}$	$\mathbf{71.15}_{4.50}$	–

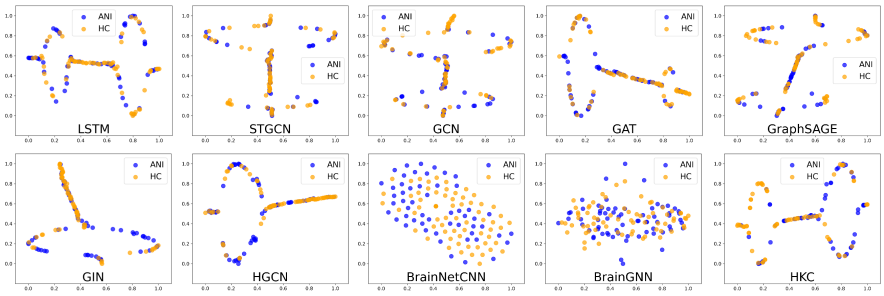

Fig. 2. Visualization of features generated by nine deep learning methods and the proposed HKC through t-SNE [21] in ANI vs. HC classification.

3 Experiment

Experimental Settings. We compare our HKC with three conventional methods (*i.e.*, **SVM** [26], **RF** [1], and **XGBoost** [4]) with 1,168-dimensional node- and graph-level features from SC and FC graphs as input, seven deep learning methods (*i.e.*, **LSTM** [14], **STGCN** [37], **GCN** [18], **GAT** [31], **GraphSAGE** [13]), **GIN** [35], and **HGCN** [3]), and two SOTA methods designed for brain network analysis (*i.e.*, **BrainNetCNN** [17] and **BrainGNN** [19]). The two methods (*i.e.*, LSTM and STGCN) take BOLD signals derived from fMRI as input. The remaining deep methods share the same graph input as our HKC. For the competing deep learning methods, we fine-tune their hyperparameters to ensure that their training parameters are comparable to ours. All methods are evaluated using a 5-fold cross-validation strategy, with mean and standard deviation results reported. Seven evaluation metrics are used: area under the ROC curve (AUC), accuracy (ACC), F1 score (F1), balanced accuracy (BAC), sensitivity (SEN), specificity (SPE), and precision (PRE).

Results and Analysis. Results of 13 methods in ANI vs. HC classification are given in Table 1. We can see that multi-modal methods outperform single-modal methods (LSTM and STGCN), demonstrating the benefit of SC and

Table 2. Results (%) of HKC and its three variants in ANI vs. HC classification.

Method	AUC	ACC	F1	BAC	SEN	SPE	PRE
HKC-G	$68.11_{1.68}$	$67.05_{2.48}$	$64.01_{3.45}$	$67.41_{1.18}$	$63.31_{7.21}$	$71.50_{7.88}$	$73.83_{5.83}$
HKC-K	$64.73_{3.60}$	$63.55_{2.42}$	$61.41_{3.41}$	$63.84_{3.07}$	$61.63_{5.34}$	$66.04_{8.01}$	$68.24_{3.01}$
HKCw/oC	$64.32_{3.59}$	$64.25_{2.45}$	$57.63_{6.68}$	$64.22_{1.83}$	$53.16_{9.60}$	$75.29_{6.76}$	$72.94_{5.08}$
HKCw/oH	$69.72_{2.75}$	$68.06_{2.99}$	$63.77_{2.29}$	$67.97_{1.95}$	$60.63_{4.05}$	$\mathbf{75.30}_{6.13}$	$\mathbf{74.57}_{5.88}$
HKC (Ours)	$\mathbf{70.56}_{1.61}$	$\mathbf{68.65}_{2.07}$	$\mathbf{66.71}_{2.87}$	$\mathbf{68.46}_{1.55}$	$\mathbf{66.65}_{6.67}$	$70.28_{7.76}$	$71.15_{4.50}$

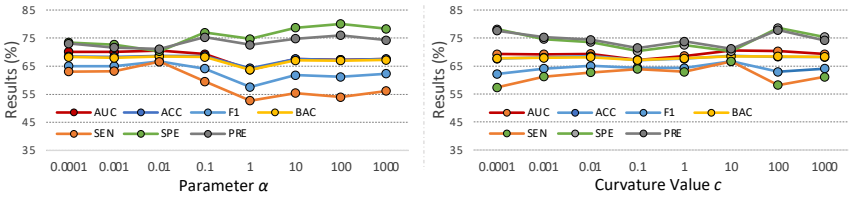

Fig. 3. Influence of the two parameters (*i.e.*, α and c) on the performance of HKC.

FC feature integration. Our HKC achieves the best AUC (70.56%) and ACC (68.65%), suggesting its effectiveness in capturing cross-modality local-to-global interactions and hierarchical dependencies among ROIs. Additionally, HKC outperforms SOTA BrainNetCNN and BrainGNN. The possible reason is that HKC employs an SC-FC coupling module for feature fusion, while these two methods rely on SC and FC feature concatenation. We conduct t-test on results achieved by HKC and each competing method to assess their performance difference significance, with significant results ($p<0.05$) marked as '*' in Table 1. Table 1 suggests that HKC significantly outperforms most competing methods. Following [5], we also visualize features extracted from the last layer of each model using t-SNE [21] in Fig. 2. This figure shows that HKC yields more separable clusters and clearer class-level distinctions, compared with the competing methods.

Ablation Study. To assess the impact of HKC's key components, we compare it to its variants: (1) **HKC-G** that uses GCN instead of HKGCN as backbone; (2) **HKC-K**, which replaces HKGCN with HGCN; (3) **HKCw/oC**, which removes SC-FC coupling and concatenates HKGCN-extracted fMRI and DTI features for prediction; and (4) **HKCw/oH** that uses multilayer perceptron (MLP) instead of HNN as the predictor. As shown in Table 2, HKC achieves the overall best performance compared to all variants. HKC-G and HKC-K exhibit moderate performance degradation, with AUC dropping to 68.11% and 64.73%, respectively. This implies that our method effectively models the hierarchical dependencies among ROIs compared to HKC-G, and better captures cross-modality local-to-global interactions than HKC-K. The performance drop of HKCw/oC highlights the importance of explicit SC-FC coupling in feature fusion. Additionally, the slight performance decline observed in HKCw/oH suggests that HNN better aligns with the hyperbolic geometry compared to a standard MLP.

Parameter Analysis. We analyze the impact of two hyperparameters in HKC: curvature value c which defines the geometry of hyperbolic space, and α which

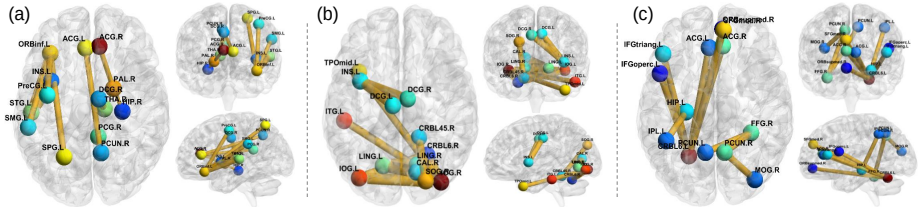

Fig. 4. Top 10 discriminative connectivities identified by the proposed HKC from (a) SC, (b) FC, and (c) SC-FC coupling graphs in ANI vs. HC classification.

scales the cosine function. The experimental results achieved by HKC with different parameter values are reported in Fig. 3. This figure indicates that varying c has minimal impact, as AUC fluctuates slightly within the range of 67.17% to 70.56%, demonstrating the robustness of our method. In contrast, α influences performance to some extent but remains stable within a certain range.

Visualization of Discriminative Connectivity. We visualize the discriminative connectivities identified by our HKC from SC and FC, and SC-FC coupling graph embeddings through the BrainNet Viewer [34]. The SC and FC features are extracted using the HKGCN, while SC-FC coupling features are derived from our SC-FC coupling module. We compute PC for these three view embeddings to generate FC, SC, and SC-FC coupling graphs, use t-test (significant level: 0.05) to identify significant connectivity differences between the two groups, and visualize the top 10 discriminative connections in Fig. 4.

From Fig. 4, we can observe significant alterations in SC, FC, and SC-FC coupling graphs across several key ROIs. *First*, SC differences are particularly evident in frontal-subcortical circuits (ACG.R-PAL.R, ACG.R-THA.R, DCG.R-THA.R) [25] and the central executive network (CEN) (ORBinf.L-SMG.L, PreCG.L-ORBinf.L, ACG.L-PCUN.R) [39]. Structural changes in these regions may be associated with cognitive impairment [11, 29]. *Second*, FC abnormalities are primarily observed in the occipital lobe and lingual gyrus within the visual network (LING.L-LING.R, IOG.L-IOG.R, CAL.R-LING.L), which are related to HIV-related cognitive impairment in attention and cognitive flexibility [28]. Meanwhile, core regions of the salience network (SN) are also altered, including the insula and cingulate gyrus (INS-DCG), which is closely associated with impairments in executive function and attention [2]. *Third*, SC-FC coupling strength is altered in hippocampal-parietal (HIP-IPL) circuits, suggesting impaired coordination between brain structure and function [22]. Additionally, reduced fronto-cerebellar SC-FC coupling highlights the critical role of the cerebellum in ANI analysis, as it not only influences motor coordination and executive function but also affects cognitive and emotional regulation [32]. We further verify the robustness of the identified discriminative connections by testing their consistency across different data splits and weight initialization, and find that key connections and involved regions (*e.g.*, ACG, PCUN) remain stable across

splits and weight initialization. These imaging biomarkers identified by our HKC could be potentially used for automated ANI detection in clinical practice.

4 Conclusion and Future Work

This work presents a novel framework, HKC, that integrates HKGCN with SC-FC Coupling for brain connectivity analysis. It leverages hyperbolic kernels to capture the hierarchical structure of brain networks, while SC-FC coupling models global interactions between SC and FC. Experimental results suggest the superiority of HKC over SOTA methods. Future efforts will extend HKC to other disorders to assess its generalizability across multi-site multi-modal cohorts. Additionally, the limited sample size poses a challenge for the effective training of deep learning models. We will address this limitation by leveraging pretrained models [10, 33] based on auxiliary data to improve model robustness.

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

References

1. Breiman, L.: Random forests. Machine Learning **45**, 5–32 (2001)
2. Chaganti, J., Heinecke, A., Gates, T., Moffat, K., Brew, B.: Functional connectivity in virally suppressed patients with HIV-associated neurocognitive disorder: A resting-state analysis. American Journal of Neuroradiology **38**(8), 1623–1629 (2017)
3. Chami, I., Ying, R., Ré, C., Leskovec, J.: Hyperbolic graph convolutional neural networks. In: NeurIPS. vol. 32 (2019)
4. Chen, T., Guestrin, C.: XGBoost: A scalable tree boosting system. In: KDD. pp. 785–794 (2016)
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML. pp. 1597–1607. PmLR (2020)
6. Cho, Y., Saul, L.: Kernel methods for deep learning. In: NeurIPS. vol. 22 (2009)
7. Cui, Z., Zhong, S., Xu, P., He, Y., Gong, G.: PANDA: a pipeline toolbox for analyzing brain diffusion images. Frontiers in Human Neuroscience **7**, 42 (2013)
8. Deco, G., Ponce-Alvarez, A., Hagmann, P., Romani, G.L., Mantini, D., Corbetta, M.: How local excitation–inhibition ratio impacts the whole brain dynamics. Journal of Neuroscience **34**(23), 7886–7898 (2014)
9. Fang, P., Harandi, M., Lan, Z., Petersson, L.: Poincaré kernels for hyperbolic representations. International Journal of Computer Vision pp. 1–23 (2023)
10. Fang, Y., Zhang, J., Wang, L., Wang, Q., Liu, M.: ACTION: Augmentation and computation toolbox for brain network analysis with functional MRI. NeuroImage **305**, 120967 (2025)
11. Ferrarese, C., Aliprandi, A., Tremolizzo, L., Stanzani, L., De Micheli, A., Dolara, A., Frattola, L.: Increased glutamate in CSF and plasma of patients with HIV dementia. Neurology **57**(4), 671–675 (2001)
12. Ganea, O., Bécigneul, G., Hofmann, T.: Hyperbolic neural networks. In: NeurIPS. vol. 31, pp. 5345–5355 (2018)

13. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. NeurIPS **30** (2017)
14. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation **9**(8), 1735–1780 (1997)
15. Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P.: Predicting human resting-state functional connectivity from structural connectivity. PNAS **106**(6), 2035–2040 (2009)
16. Iyer, S.P., Shafran, I., Grayson, D., Gates, K., Nigg, J.T., Fair, D.A.: Inferring functional connectivity in MRI using Bayesian network structure learning with a modified PC algorithm. NeuroImage **75**, 165–175 (2013)
17. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. NeuroImage **146**, 1038–1049 (2017)
18. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017)
19. Li, X., Zhou, Y., Dvornek, N.C., Zhang, M., Gao, S., Zhuang, J., Scheinost, D., Staib, L.H., Ventola, P., Duncan, J.S.: BrainGNN: Interpretable brain graph neural network for fMRI analysis. In: CVPR. pp. 9260–9270 (2021)
20. Luo, X., Wu, J., Yang, J., Xue, S., Beheshti, A., Sheng, Q.Z., McAlpine, D., Sowman, P.F., Giral, A., Yu, P.S.: Graph neural networks for brain graph learning: A survey. In: IJCAI. pp. 8170–8178 (2024)
21. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. Journal of Machine Learning Research **9**(11) (2008)
22. Maki, P., Cohen, M., Weber, K., Little, D., Fornelli, D., Rubin, L., Perschler, P., Gould, F., Martin, E.: Impairments in memory and hippocampal function in HIV-positive vs HIV-negative women: A preliminary study. Neurology **72**(19), 1661–1668 (2009)
23. Moreno-Dominguez, D., Anwander, A., Knösche, T.R.: A hierarchical method for whole-brain connectivity-based parcellation. Human Brain Mapping **35**(10), 5000–5025 (2014)
24. Mulkern, R.V., Davis, P.E., Haker, S.J., Estepar, R.S.J., Panych, L.P., Maier, S.E., Rivkin, M.J.: Complementary aspects of diffusion imaging and fMRI; I: structure and function. Magnetic Resonance Imaging **24**(4), 463–474 (2006)
25. Ortega, M., Brier, M.R., Ances, B.M.: Effects of HIV and combination antiretroviral therapy on cortico-striatal functional connectivity. Aids **29**(6), 703–712 (2015)
26. Pisner, D.A., Schnyer, D.M.: Support Vector Machine. In: Machine Learning, pp. 101–121. Elsevier (2020)
27. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. In: NeurIPS. vol. 20 (2007)
28. Saylor, D., Dickens, A.M., Sacktor, N., Haughey, N., Slusher, B., Pletnikov, M., Mankowski, J.L., Brown, A., Volsky, D.J., McArthur, J.C.: HIV-associated neurocognitive disorder—pathogenesis and prospects for treatment. Nature Reviews Neurology **12**(4), 234–248 (2016)
29. Schrock, J.M.: Accelerated aging in people living with hiv: The neuroimmune feedback model. Brain, Behavior, & Immunity-Health **36**, 100737 (2024)
30. Thomas Yeo, B., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., et al.: The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of neurophysiology **106**(3), 1125–1165 (2011)

31. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017)
32. Wang, H., Li, R., Zhou, Y., Wang, Y., Cui, J., Nguchu, B.A., Qiu, B., Wang, X., Li, H.: Altered cerebro-cerebellum resting-state functional connectivity in HIV-infected male patients. Journal of NeuroVirology **24**, 587–596 (2018)
33. Wang, X., Fang, Y., Wang, Q., Yap, P.T., Zhu, H., Liu, M.: Self-supervised graph contrastive learning with diffusion augmentation for functional MRI analysis and brain disorder detection. Medical Image Analysis **101**, 103403 (2025)
34. Xia, M., Wang, J., He, Y.: BrainNet Viewer: A network visualization tool for human brain connectomics. PloS One **8**(7), e68910 (2013)
35. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? In: ICLR (2019)
36. Yan, C., Zang, Y.: DPARSF: A MATLAB toolbox for "pipeline" data analysis of resting-state fMRI. Frontiers in Systems Neuroscience **4**, 1377 (2010)
37. Yu, B., Yin, H., Zhu, Z.: Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In: IJCAI. pp. 3634–3640 (2018)
38. Zhang, L., Na, S., Liu, T., Zhu, D., Huang, J.: Multimodal deep fusion in hyperbolic space for mild cognitive impairment study. In: MICCAI. pp. 674–684. Springer (2023)
39. Zhou, Z., Wang, W., Li, H., Shi, Y., Zhao, L., Lu, Y., Wei, X., Li, H.: Decoding HIV-associated neurocognitive disorders: a new perspective from multimodal connectomics. Frontiers in Neurology **16**, 1467175 (2025)