

# DCHO: A Decomposition–Composition Framework for Predicting Higher-Order Brain Connectivity to Enhance Diverse Downstream Applications

Weibin Li\*, Wendu Li\*, Quanying Liu<sup>†</sup>,

Department of Biomedical Engineering, Southern University of Science and Technology, China  
liwb2023@mail.sustech.edu.cn, 12432838@mail.sustech.edu.cn, liuqy@sustech.edu.cn

## Abstract

Higher-order brain connectivity (HOBC), which captures interactions among three or more brain regions, provides richer organizational information than traditional pairwise functional connectivity (FC). Recent studies have begun to infer latent HOBC from noninvasive imaging data, but they mainly focus on static analyses, limiting their applicability in dynamic prediction tasks. To address this gap, we propose DCHO, a unified approach for modeling and forecasting the temporal evolution of HOBC based on a **Decomposition–Composition** framework, which is applicable to both non-predictive tasks (state classification) and predictive tasks (brain dynamics forecasting). DCHO adopts a decomposition–composition strategy that reformulates the prediction task into two manageable subproblems: HOBC inference and latent trajectory prediction. In the inference stage, we propose a dual-view encoder to extract multiscale topological features and a latent combinatorial learner to capture high-level HOBC information. In the forecasting stage, we introduce a latent-space prediction loss to enhance the modeling of temporal trajectories. Extensive experiments on multiple neuroimaging datasets demonstrate that DCHO achieves superior performance in both non-predictive tasks (state classification) and predictive tasks (brain dynamics forecasting), significantly outperforming existing methods.

## 1 Introduction

Functional connectivity (FC) (Yeo et al. 2011; Hutchison et al. 2013), the most widely used framework, models pairwise statistical dependencies between brain regions based on fMRI data. FC is limited to pairwise interactions and may overlook higher-level coordination. Recent studies highlight the importance of higher-order brain connectivity (HOBC) that involves three or more regions, supported by evidence across scales (Battiston et al. 2020; Battiston and Petri 2022; Bianconi 2021; Jun et al. 2017; Steinmetz et al. 2021; Paulk et al. 2022; Chelaru et al. 2021; Tadić, Chutani, and Gupte 2022). In human neuroimaging, direct observation of HOBC remains challenging due to the limitations of noninvasive techniques. A recent study (Santoro et al. 2024, 2023) was the first to infer latent HOBC from fMRI signals and applied

\*These authors contributed equally.

<sup>†</sup>Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

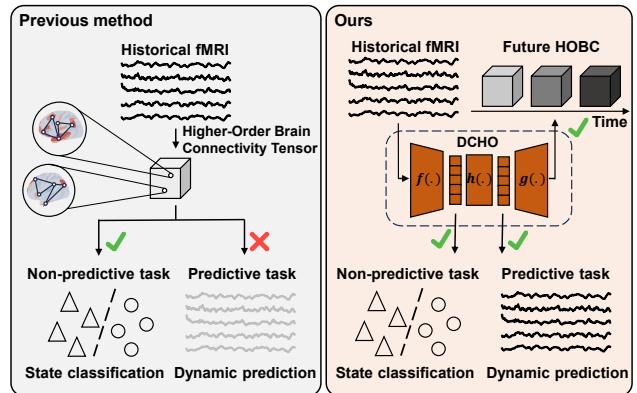


Figure 1: **Motivation.** (Left): Previous methods analyze HOBC in static windows, overlooking their temporal evolution and limiting predictive applications. (Right): DCHO overcomes this limitation by forecasting the dynamic trajectories of HOBC.

them to downstream tasks such as classification. However, they focus on analyzing HOBC within static or predefined temporal windows, without modeling their temporal evolution. This limits the potential of HOBC in predictive applications, such as forecasting brain dynamics. Addressing this gap requires models that can anticipate the dynamic trajectory of the HOBC over time.

Predicting the evolution of the HOBC in complex brain systems is challenging. First, the combinatorial explosion in the number of possible higher-order interactions poses a major obstacle to direct modeling (Bassett and Sporns 2017; Breakspear 2017; Roebroek, Formisano, and Goebel 2011; Herzog et al. 2022). The curse of dimensionality poses a major challenge to inferring the HOBC. Second, brain dynamics are inherently nonlinear, nonstationary, and stochastic, governed by latent neural processes that evolve across multiple spatial and temporal scales. Capturing the temporal evolution of HOBC thus requires not only modeling individual node trajectories, but also learning coordinated multi-region dependencies that reflect complex integration, segregation, and recurrent feedback mechanisms. These inter-

actions are not additive but shaped by nonlinear dynamical couplings that defy simple compositional rules (Millán et al. 2025; Battiston et al. 2021; Zhang, Lucas, and Battiston 2023; Alvarez-Rodriguez et al. 2021).

To address the above limitations, we propose DCHO, a model designed to capture and predict the temporal evolution of higher-order brain connectivity, while leveraging its high-level representations to support both non-predictive and predictive tasks. DCHO introduces an innovative decomposition–composition framework that reformulates the complex prediction task into two more tractable subtasks: HOBC tensor inference and latent trajectory prediction. In the inference stage, DCHO employs a dual-view encoder to capture multi-scale topological features and incorporates a latent composition learner to further encode higher-order interaction patterns, effectively mitigating the curse of dimensionality and enhancing representational capacity. In the forecasting stage, DCHO introduces a latent-space prediction loss to model the dynamic evolution of HOBC in an abstract information space. With the pre-trained encoder and predictor, DCHO is broadly applicable to both non-predictive tasks (e.g., state classification) and predictive tasks (e.g., brain dynamics forecasting), and consistently outperforms existing state-of-the-art methods across multiple datasets (Figure 1). Our main contributions are summarized as follows:

- **Decomposition–Composition Framework:** DCHO decomposes prediction into HOBC inference and latent forecasting, supported by theoretical analysis and enhanced by a latent-space prediction loss for abstract trajectory modeling.
- **Dual-view Encoder and Latent Combinatorial Learner:** We propose the dual-view encoder and the latent learner to capture multiscale topologies and high-level HOBC information.
- **Effective Across Predictive and Non-predictive Applications:** DCHO leverages pretrained encoder and predictor for both non-predictive (state classification) and predictive tasks (brain dynamics forecasting), outperforming state-of-the-art methods on multiple datasets.

## 2 Related Work

**Higher-order Brain Connectivity:** Although FC (Yeo et al. 2011; Hutchison et al. 2013) has been the standard for brain network modeling, it fails to capture higher-order interactions among multiple brain regions. HOBC describes complex coordination among three or more regions and reveals richer organizational patterns. However, due to the limitations of non-invasive neuroimaging techniques, directly observing HOBC in humans remains challenging. Recent work (Santoro et al. 2024, 2023) has attempted to infer HOBC from fMRI and apply it to static classification tasks. However, the study overlooked its temporal dynamics, limiting their use in predictive settings. To bridge this gap, we propose the DCHO, which models and forecasts the temporal evolution of HOBC from neural signals, enabling both non-predictive tasks (task state classification) and more demanding predictive tasks (brain dynamics forecasting).

**Brain Dynamics Forecasting:** In recent years, various architectures have been proposed for modeling neural dynamics. RNN-based methods (LtrRNN (Pellegrino, Gajic, and Chadwick 2023) and LSTM (Graves and Graves 2012)) are effective in capturing short-term temporal dependencies but are limited in expressing long-term dynamics. Transformer-based models (NetFormer (Lu et al. 2025), STNDT (Le and Shlizerman 2022) and Transformer (Vaswani et al. 2017)) leverage self-attention mechanisms to better model global temporal dependencies, yet often overlook the underlying brain connectivity structure. GNN-based methods (AMAG (Li et al. 2023)) are naturally suited for encoding spatial topological structures between brain regions and have been applied to model static or pairwise functional connectivity, but typically neglect global topological information. In contrast, the proposed DCHO framework integrates a dual-graph encoder and a latent learner to capture high-level information representations of HOBC and explicitly models the underlying dynamic evolution in the latent space, substantially enhancing modeling capacity and robustness in long-term prediction.

## 3 Preliminaries

### 3.1 The Definition of Higher-order Brain Connectivity

We first introduce the definition of HOBC (Santoro et al. 2024, 2023). Let  $\mathbf{x}_i \in \mathbb{R}^T = [x_i^1, x_i^2, \dots, x_i^T]$  denotes the original time series of region  $i$ . We first  $z$ -score each region’s time series:

$$\tilde{\mathbf{x}}_i = \frac{\mathbf{x}_i - \mu[\mathbf{x}_i]}{\sigma[\mathbf{x}_i]}, \quad (1)$$

where  $\mu[\cdot]$  and  $\sigma[\cdot]$  denote the time-averaged mean and standard deviation, respectively. For a  $(k+1)$ -node simplex, we define the  $k$ -order  $z$ -scored co-fluctuation signal at time  $t$  as:

$$\xi_{0\dots k}^t = \frac{\prod_{p=0}^k \tilde{x}_p^t - \mu \left[ \prod_{p=0}^k \tilde{\mathbf{x}}_p \right]}{\sigma \left[ \prod_{p=0}^k \tilde{\mathbf{x}}_p \right]}. \quad (2)$$

To distinguish concordant from discordant interactions within a  $k$ -order product, we assign positive values to fully concordant signs, whereas discordant combinations are mapped to negative values. Namely,

$$\text{sign} [\xi_{0\dots k}^t] := (-1)^{\text{sgn}((k+1) - |\sum_0^k \text{sgn}[\tilde{x}_i^t]|)}, \quad (3)$$

where  $\text{sgn}[\cdot]$  is the signum function of a real number. The final weighted co-fluctuation signal is defined as:

$$w_{0\dots k}^t = \text{sign} [\xi_{0\dots k}^t] \cdot |\xi_{0\dots k}^t|. \quad (4)$$

If all  $k$ -order products are computed, this yields a total of  $\binom{N}{k+1}$  distinct co-fluctuation time series for each order  $k$ , where  $N$  denotes the number of brain regions.

At each time point  $t$ , we organize the resulting  $k$ -order co-fluctuations into a weighted simplicial complex  $\mathcal{K}^t$ . For simplicity, in this work, we only consider co-fluctuations of dimension up to  $k = 2$ , so that triangles represent the higher-order connectivity in the weighted simplicial complex  $\mathcal{K}$ . We

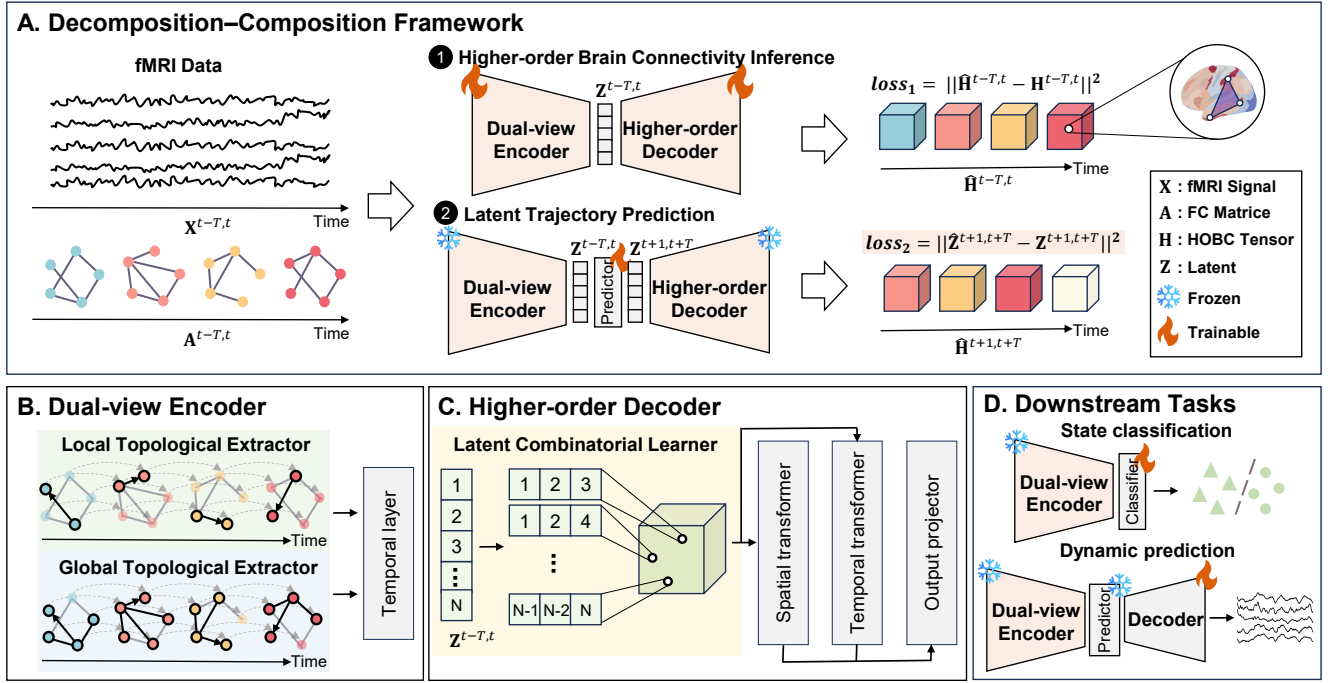


Figure 2: **Overview of DCHO** (A) **Decomposition-Composition Framework**: DCHO framework decomposes prediction into HOBC inference and latent trajectory prediction, using a **latent-space prediction loss** to model high-level temporal dynamics. (B) **Dual-view Encoder**: DCHO applies two parallel GNN branches that extract local and global topological features. (C) **Higher-order Decoder**: DCHO proposes a **latent combinatorial learner** to capture high-level HOBC information. (D) DCHO leverages the pretrained encoder and predictor to support both non-predictive and predictive tasks.

represent the higher-order connectivity  $\mathcal{K}^t$  as a third-order tensor  $H^t \in \mathbb{R}^{N \times N \times N}$ , where each element  $h_{ijk}^t$  denotes the weighted co-fluctuation signal  $w_{ijk}^t$  among regions  $i, j, k$  at time  $t$ .

## 4 Method

### 4.1 Decomposition-Composition Framework

Let  $\mathbf{X}^{t-T,t} = [X^{t-T}; \dots; X^t] \in \mathbb{R}^{T \times N}$  be the original fMRI signals and  $\mathbf{A}^{t-T,t} = [A^{t-T}; \dots; A^t] \in \mathbb{R}^{T \times N \times N}$  be the corresponding functional connectivity matrices. Our goal is to predict the future value of HOBC  $\hat{\mathbf{H}}^{t+1,t+T} = [\hat{H}^{t+1}; \dots; \hat{H}^{t+T}] \in \mathbb{R}^{T \times N \times N \times N}$  based on  $\mathbf{A}^{t-T,t}$  and  $\mathbf{X}^{t-T,t}$ . In this study, we present a decomposition-composition framework that breaks the overall objective into two tractable subtasks: (1) higher-order brain connectivity inference and (2) latent trajectory prediction. The overview of DCHO is summarized in Figure 2 and Algorithm 1. Algorithm 1 is provided in the Appendix.

The first subtask aims to infer a sequence of the HOBC tensors  $\hat{\mathbf{H}}^{t-T,t}$  based on  $\mathbf{X}^{t-T,t}$  and  $\mathbf{A}^{t-T,t}$ . A dual-view encoder first transforms these inputs into latent spatio-temporal representations  $\mathbf{Z}^{t-T,t}$  by applying two parallel GNN branches that extract local and global topological features. These latent representations are then fed into a higher-order decoder to infer the HOBC tensors. The decoder consists of a latent combinatorial learner and a dual-stream

Transformer module, where the learner further encourages the latent representation  $\mathbf{Z}^{t-T,t}$  to capture higher-order interaction patterns. The loss of the first subtask is:

$$loss_1 = \|\hat{\mathbf{H}}^{t-T,t} - \mathbf{H}^{t-T,t}\|^2. \quad (5)$$

The second subtask focuses on predicting future latent dynamics. We apply a LSTM-based predictor to model temporal dependencies in the latent space and predict future embeddings  $\hat{\mathbf{Z}}^{t+1,t+T}$  from  $\mathbf{Z}^{t-T,t}$ . This latent trajectory prediction forms the basis for estimating future HOBC tensors. The loss of the second subtask is:

$$loss_2 = \|\hat{\mathbf{Z}}^{t+1,t+T} - \mathbf{Z}^{t+1,t+T}\|^2. \quad (6)$$

The training process is conducted in two stages: we first train the higher-order inference subtask, and subsequently freeze both the encoder and decoder to optimize the latent trajectory prediction subtask. The advantages of the proposed framework are theoretically justified and analyzed in the following.

**Theorem 1.** Let  $f_{\text{enc}} : \mathbb{R}^{T \times N \times N} \rightarrow \mathbb{R}^{T \times N \times F}$ ,  $f_{\text{dyn}} : \mathbb{R}^{T \times N \times F} \rightarrow \mathbb{R}^{T \times N \times F}$ , and  $f_{\text{dec}} : \mathbb{R}^{T \times N \times F} \rightarrow \mathbb{R}^{T \times N \times N \times N}$  be measurable mappings that are Lipschitz continuous, with Lipschitz constants  $L_{\text{enc}}, L_{\text{dyn}}, L_{\text{dec}}$ , respectively. The prediction of a trained model is defined as:

$$\hat{\mathbf{H}}^{t+1,t+T} = f_{\text{dec}}(f_{\text{dyn}}(f_{\text{enc}}(\mathbf{X}^{t-T,t}, \mathbf{A}^{t-T,t}))). \quad (7)$$

Metrics				MAE			
Dataset	Emotion	Gambling	Language	Motor	Rest	Lorenz	HR
MLP	0.3660 ± 0.0005	0.3008 ± 0.0009	0.3311 ± 0.0008	0.2907 ± 0.0008	0.5875 ± 0.0003	0.1155 ± 0.0017	0.0833 ± 0.0028
Trans.	0.3618 ± 0.0020	0.2605 ± 0.0018	0.2532 ± 0.0010	0.2185 ± 0.0023	0.4959 ± 0.0019	<u>0.1250</u> ± 0.0024	0.0559 ± 0.0019
LSTM	<u>0.2887</u> ± 0.0004	<u>0.2216</u> ± 0.0009	<u>0.2352</u> ± 0.0005	<u>0.1971</u> ± 0.0009	<u>0.3865</u> ± 0.0002	0.2608 ± 0.0085	<u>0.0416</u> ± 0.0009
DCHO	<b>0.1744</b> ± 0.0004	<b>0.0720</b> ± 0.0003	<b>0.1156</b> ± 0.0005	<b>0.0778</b> ± 0.0004	<b>0.1704</b> ± 0.0004	<b>0.0780</b> ± 0.0001	<b>0.0249</b> ± 0.0001
Metrics				RMSE			
Dataset	Emotion	Gambling	Language	Motor	Rest	Lorenz	HR
MLP	0.5911 ± 0.0022	0.4790 ± 0.0015	0.5402 ± 0.0014	0.4589 ± 0.0024	1.1158 ± 0.0067	0.2017 ± 0.0070	0.1772 ± 0.0055
Trans.	0.6033 ± 0.0062	0.4466 ± 0.0020	0.4117 ± 0.0008	0.3446 ± 0.0041	0.9793 ± 0.0030	<u>0.1923</u> ± 0.0037	0.1083 ± 0.0070
LSTM	<u>0.4709</u> ± 0.0008	<u>0.3858</u> ± 0.0016	<u>0.3874</u> ± 0.0010	<u>0.3170</u> ± 0.0012	<u>0.5240</u> ± 0.0030	0.3779 ± 0.0269	<u>0.0904</u> ± 0.0031
DCHO	<b>0.2672</b> ± 0.0011	<b>0.1161</b> ± 0.0009	<b>0.1852</b> ± 0.0009	<b>0.1239</b> ± 0.0009	<b>0.2738</b> ± 0.0042	<b>0.1388</b> ± 0.0288	<b>0.0412</b> ± 0.0186

Table 1: Prediction performance comparison (MAE and RMSE) of DCHO and baseline models on multiple datasets for HOBC forecasting with a 10-step horizon.

The inference error is defined as:

$$\epsilon_{\text{inf}} = \left\| f_{\text{dec}}(f_{\text{enc}}(\mathbf{X}^{t+1,t+T}, \mathbf{A}^{t+1,t+T})) - \mathbf{H}^{t+1,t+T} \right\|. \quad (8)$$

Let  $\mathbf{Z}^{t-T,t} = f_{\text{enc}}(\mathbf{X}^{t-T,t}, \mathbf{A}^{t-T,t})$ . The latent prediction error is defined as:

$$\epsilon_{\text{dyn}} = \left\| f_{\text{dyn}}(\mathbf{Z}^{t-T,t}) - \mathbf{Z}^{t+1,t+T} \right\|. \quad (9)$$

Then the prediction error satisfies:

$$\left\| \hat{\mathbf{H}}^{t+1,t+T} - \mathbf{H}^{t+1,t+T} \right\| \leq \epsilon_{\text{inf}} + L_{\text{dec}} \cdot \epsilon_{\text{dyn}}. \quad (10)$$

The proof of Theorem 1 is provided in the Appendix.

**Remark 1.** By decoupling the sources of error, each component can be independently analyzed and optimized, making the total prediction error more controllable and theoretically better bounded than in end-to-end models with entangled errors. Benefiting from the dual-view encoder, latent combinatorial learner, and latent-space prediction loss, DCHO effectively reduces both the inference error  $\epsilon_{\text{inf}}$  and the latent dynamics error  $\epsilon_{\text{dyn}}$ , enabling fine-grained modeling of the two subtasks. As shown in the Appendix, ablation results further validate the effectiveness of the decomposition–composition framework.

## 4.2 The Dual-view Encoder

The dual-view encoder consists of two parallel GNN branches—a local and a global topological extractor—that process spatiotemporal inputs  $\mathbf{X}^{t-T,t}$  and  $\mathbf{A}^{t-T,t}$  into latent representations  $\mathbf{Z}^{t-T,t}$  enriched with structural and temporal information. Details are as follows:

(i) **The local topological extractor** adaptively encodes pairwise interactions using a GNN-based mechanism (Kipf and Welling 2017; Xu et al. 2019). First, interaction scores are computed between the central node and its neighbors, and these scores are used to weight the neighbors’ embeddings from the previous layer. Mathematically, for the embedding of the  $i_t$  node  $v_i^t$  at the  $s$ -th layer:  $e_i^{t,(s)}$ , the interaction scores between the  $i_t$  node and its neighbor  $j_t$  are derived from the adjacency matrix and the node features in

the temporal graph:

$$sc^{(s)}(v_i^t, v_j^t) = A(i_t, j_t) \cos\left(W_q e_i^{t,(s)}, W_k e_j^{t,(s)}\right), \quad (11)$$

where the cosine similarity:  $\cos(\cdot, \cdot)$  quantifies the similarity between two vectors.  $W_q$  and  $W_k$  are transformation matrices employed for similarity calculations. Then, the node representation at the  $(s+1)$ -th layer is updated as follows:

$$e_i^{t,(s+1)} = e_i^{t,(s)} + \sigma\left(\sum_{v_j^t \in N_{v_i^t}} sc^{(s)}(v_i^t, v_j^t) W_v e_j^{t,(s)}\right), \quad (12)$$

where  $\sigma(\cdot)$  is a non-linear activation function,  $W_v$  is a learnable matrix, and  $N_{v_i^t}$  denote as the neighbors of  $v_i^t$ . After  $S$  layers, the final node representation:  $e_i^t = e_i^{t,(S)}$  is adaptively determined.

(ii) **The global topological extractor** is designed based on spectral graph convolution, which serves as a complementary approach to uncover HOBC information hidden in the spectral domain. Specifically, first, we introduce the necessary mathematical definitions. The Chebyshev polynomials are defined as  $T_0(x) = 1$ ,  $T_1(x) = x$ , and  $T_m(x) = 2xT_{m-1}(x) - T_{m-2}(x)$ . The normalized graph Laplacian  $L$  is defined as  $L = I - \tilde{D}^{-\frac{1}{2}} A \tilde{D}^{-\frac{1}{2}}$ , where  $D$  is the degree matrix and  $I$  is the identity matrix. Define the matrix  $R^{(0)} = X$ , which stacks the features of all nodes. Then, based on the second-order Chebyshev graph convolution (Defferrard, Bresson, and Vandergheynst 2016), the node representation matrix  $R^{(s)}$  at the  $s$ -th layer is calculated as:

$$R^{(s)} = \sum_{m=0}^2 T_m(\tilde{L}) R^{(s-1)} W_m^{(s)} \quad (13)$$

where  $\tilde{L} = \frac{2L}{\lambda_{\max}} - I$ ,  $\lambda_{\max}$  denotes the maximum eigenvalue of  $L$ , and  $W_m^{(s)} \in \mathbb{R}^{d \times d}$  is a learnable weight matrix at the  $s$ -th layer. By stacking  $S$  layers, we derive each representation vector  $r_i^t$  from  $R^{(S)} \in \mathbb{R}^{NT \times d}$ , capturing higher order non-local information from a spectral perspective.

Metrics				MAE			
Dataset	Emotion	Gambling	Language	Motor	Rest	Lorenz	HR
MLP	0.7649 ± 0.0008	0.7554 ± 0.0024	0.7101 ± 0.0019	0.7469 ± 0.0024	0.7834 ± 0.0001	0.1381 ± 0.0025	0.4927 ± 0.0060
Trans.	0.2651 ± 0.0018	0.3176 ± 0.0022	0.2781 ± 0.0009	0.2525 ± 0.0016	0.8444 ± 0.0018	0.1017 ± 0.0021	0.6998 ± 0.0096
LSTM	0.2361 ± 0.0001	0.2831 ± 0.0002	<u>0.2625</u> ± 0.0003	<u>0.2347</u> ± 0.0003	0.9874 ± 0.0002	<u>0.1033</u> ± 0.0008	0.3726 ± 0.0484
STNDT	0.6624 ± 0.0052	0.7134 ± 0.0038	<u>0.7802</u> ± 0.0033	<u>0.6964</u> ± 0.0023	0.7637 ± 0.0005	<u>0.3146</u> ± 0.0105	0.2961 ± 0.0454
AMAG	<u>0.2103</u> ± 0.0020	<u>0.1973</u> ± 0.0019	0.3571 ± 0.0012	0.4734 ± 0.0015	<u>0.5558</u> ± 0.0009	0.5547 ± 0.0107	<u>0.1670</u> ± 0.0785
LtrRNN	0.6665 ± 0.0077	0.7559 ± 0.0026	0.8094 ± 0.0033	0.7811 ± 0.0036	0.7926 ± 0.0007	0.7222 ± 0.0068	0.8560 ± 0.0588
NetFormer	0.5873 ± 0.0026	0.6776 ± 0.0030	0.7121 ± 0.0034	0.6367 ± 0.0016	0.7014 ± 0.0004	0.2596 ± 0.0090	0.5355 ± 0.0860
DCHO	<b>0.1582</b> ± 0.0011	<b>0.0750</b> ± 0.0011	<b>0.1382</b> ± 0.0006	<b>0.0936</b> ± 0.0008	<b>0.2515</b> ± 0.0007	<b>0.0613</b> ± 0.0012	<b>0.0545</b> ± 0.0008

Metrics				RMSE			
Dataset	Emotion	Gambling	Language	Motor	Rest	Lorenz	HR
MLP	0.9500 ± 0.0011	0.9545 ± 0.0030	0.8928 ± 0.0029	0.9482 ± 0.0031	0.9851 ± 0.0007	0.2444 ± 0.0052	0.7160 ± 0.0156
Trans.	0.3377 ± 0.0021	0.4051 ± 0.0031	0.3561 ± 0.0015	0.3218 ± 0.0019	1.0676 ± 0.0023	<u>0.1492</u> ± 0.0040	0.7762 ± 0.0220
LSTM	0.3161 ± 0.0002	0.3740 ± 0.0004	<u>0.3494</u> ± 0.0005	<u>0.3116</u> ± 0.0004	1.2426 ± 0.0006	0.1525 ± 0.0030	0.5054 ± 0.0600
STNDT	0.8332 ± 0.0057	0.8882 ± 0.0043	<u>0.9672</u> ± 0.0036	<u>0.8701</u> ± 0.0025	0.9653 ± 0.0007	0.5762 ± 0.0229	0.4945 ± 0.0665
AMAG	<u>0.3035</u> ± 0.0025	<u>0.2024</u> ± 0.0022	0.4140 ± 0.0015	0.5103 ± 0.0025	<u>0.6104</u> ± 0.0012	0.8349 ± 0.0163	<u>0.2438</u> ± 0.0873
LtrRNN	0.8405 ± 0.0099	0.9497 ± 0.0043	1.0004 ± 0.0042	0.9325 ± 0.0042	0.9966 ± 0.0011	1.0935 ± 0.0099	0.9721 ± 0.0901
NetFormer	0.7420 ± 0.0032	0.8548 ± 0.0037	0.9034 ± 0.0051	0.8169 ± 0.0018	0.8913 ± 0.0005	0.5857 ± 0.0095	0.7124 ± 0.1034
DCHO	<b>0.2172</b> ± 0.0017	<b>0.1064</b> ± 0.0026	<b>0.1891</b> ± 0.0007	<b>0.1316</b> ± 0.0012	<b>0.3212</b> ± 0.0009	<b>0.1254</b> ± 0.0010	<b>0.0721</b> ± 0.0003

Table 2: Prediction performance comparison (MAE and RMSE) of DCHO and baseline models on multiple datasets for raw fMRI signals forecasting with a 10-step horizon.

(iii) The above two extractors have captured the representations of pairwise interactions and higher-order interactions, respectively. We then employ an attention mechanism to aggregate these temporal representations into a latent representation  $\mathbf{Z}_i^t$  with advanced spatio-temporal features. Specifically, first, we merge the two representations from the last layer with the temporal embeddings  $TE(t)$ , and apply a MLP  $\delta(\cdot)$  to calculate the representation  $c_i^t$  for each node  $v_i^t$ . Mathematically:

$$c_i^t = \delta([e_i^t, r_i^t] + TE(t)). \quad (14)$$

We then refine the node representation through an attention-based transformation that integrates contextual dependencies at each timestamp. Mathematically:

$$\zeta_i^t = \psi(c_i^t + f_{att}(c_i^t)), \quad (15)$$

$$\mathbf{Z}_i^t = \psi(\zeta_i^t + f_{fn}(\zeta_i^t)). \quad (16)$$

where  $f_{att}(\cdot)$  denotes an attention refinement operator implemented with multi-head self-attention,  $f_{fn}(\cdot)$  denotes a feed-forward network, and  $\psi(\cdot)$  denotes a Layer Normalization operation.

### 4.3 Higher-order Decoder

The higher-order decoder infers a sequence of HOBC tensors  $\hat{\mathbf{H}}^{t-T,t}$  from the latent representation  $\hat{\mathbf{Z}}^{t-T,t}$ , and is mainly composed of a latent combinatorial learner and a dual-stream transformer module. The details are presented below:

(i) **The latent combinatorial learner:** For each region embedding  $Z_i^{t-T,t}, Z_j^{t-T,t}, Z_k^{t-T,t} \in \mathbb{R}^{T \times F}$  from  $\mathbf{Z}^{t-T,t}$ , we apply three linear projections:

$$h_i = W_i Z_i^{t-T,t}, \quad h_j = W_j Z_j^{t-T,t}, \quad h_k = W_k Z_k^{t-T,t}, \quad (17)$$

where  $W_i, W_j, W_k \in \mathbb{R}^{F \times F}$  are learnable linear projections. These are concatenated and projected to obtain unified triplet tokens:

$$o_{ijk}^{t-T,t} = \text{Linear}([h_i, h_j, h_k]) \in \mathbb{R}^{T, 3 \times F}. \quad (18)$$

All  $M$  triplet-level representations  $o_{ijk}^{t-T,t} \in \mathbb{R}^{T \times 3F}$  are stacked to form the unified tensor  $o_M^{t-T,t} \in \mathbb{R}^{T \times M \times 3F}$ , where  $M = \binom{N}{3} = \frac{N(N-1)(N-2)}{6}$  denotes the total number of triplets. This tensor is then passed through a linear layer to obtain the final combinatorial representation  $\tilde{o}_M^{t-T,t} \in \mathbb{R}^{T, M, F}$ .

(ii) **Spatial and Temporal stream:** We add the positional encoding  $\text{PosEnc}(M)$  for each time step  $t$ , and the temporal encoding  $\text{TimeEnc}(T)$  for each triplet  $(i, j, k)$ :

$$\tilde{o}_s = \text{Transformer}_s \left( \tilde{o}_M^{t-T,t} + \text{PosEnc}(M) \right), \quad (19)$$

$$\tilde{o}_t = \text{Transformer}_t \left( \tilde{o}_M^{t-T,t} + \text{TimeEnc}(T) \right). \quad (20)$$

The spatial and temporal streams are then concatenated and passed through an output layer to obtain the final inference.

$$\hat{\mathbf{H}}^{t-T,t} = \text{MLP}([\tilde{o}_s, \tilde{o}_t]) \in \mathbb{R}^{T \times N \times N \times N}. \quad (21)$$

This dual-stream architecture jointly captures higher-order spatial topologies and their dynamic evolution.

### 4.4 LSTM-based predictor

We employ a multi-layer LSTM network to model the temporal evolution of latent representations. Given the latent

states  $\mathbf{Z}^{t-T,t}$  from the encoder, the predictor forecasts future trajectories as:

$$\hat{\mathbf{Z}}^{t+1,t+T} = \text{LayerNorm}(\text{Linear}(\text{LSTM}(\mathbf{Z}^{t-T,t}))). \quad (22)$$

This architecture captures long-range dependencies in HOBC dynamics, providing informative latent features for future HOBC prediction.

## 5 Results

### 5.1 Dataset and Preprocessing

To evaluate the performance of our proposed method, we conducted experiments on both synthetic data and real-world fMRI datasets. The synthetic data were generated from two nonlinear dynamical systems: the Lorenz system (Wang et al. 2020) and the Hindmarsh–Rose (HR) neuronal model (Hindmarsh and Rose 1984). In both cases, we simulated networks consisting of 20 coupled nodes. For real-world evaluation, we utilized the Human Connectome Project (HCP) dataset (Van Essen et al. 2012), which includes both resting-state and task-based fMRI recordings. Specifically, we employed seven task-based datasets: Emotion, Gambling, Language, Motor, Relational, Social, and Working Memory (WM) and a resting-state dataset (Rest). All HCP data had been preprocessed using the standard minimal preprocessing pipeline. We parcellated the cortex into 17 regions based on the Yeo 17-network atlas. It is important to note that the dimensionality of the higher-order brain connectivity tensor  $H^t \in \mathbb{R}^{N \times N \times N}$  increases cubically with the number of brain regions  $N$ , indicating that the HOBC derived from the Yeo-17 parcellation is already highly complex and high-dimensional. For each subject and session, we extracted the average BOLD signal within each region and applied temporal standardization (z-scoring) independently for each region. In terms of data segmentation, all time series were divided into fixed-length temporal segments using a sliding window, with the window length set equal to the prediction length and a stride of 2. Subjects were split into two groups, with 70% used for training and 30% for testing. See Appendix for further details.

### 5.2 Evaluation of Higher-Order Brain Connectivity Tensor Prediction

By predicting HOBC, DCHO encourages its latent representations to capture higher-order connectivity information. This facilitates both non-predictive and predictive downstream tasks. Before exploring these applications, we first evaluate DCHO’s performance in predicting HOBC.

**Experimental Setup:** Since HOBC prediction is a novel task, no existing methods have been specifically designed for this challenge. We therefore compare DCHO with commonly used baselines, including MLP (Rosenblatt 1958), LSTM (Graves and Graves 2012), and Transformer (Vaswani et al. 2017), to provide a meaningful reference. Detailed descriptions of these methods are provided in the Appendix. We evaluate model performance using two common metrics (MAE and RMSE) across ten datasets (Emotion, Gambling, Language, Motor, Relational, Social, WM, Rest, Lorenz, HR).

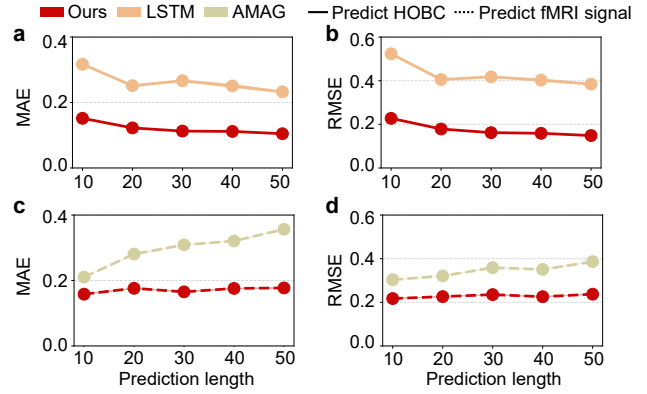


Figure 3: (a–b) Multi-step HOBC prediction of DCHO and LSTM, and (c–d) multi-step raw fMRI prediction of DCHO and AMAG on the Emotion dataset. (Metrics: MAE (left) and RMSE (right), Prediction lengths: 10 to 50 steps)

**Performance Analysis:** We systematically evaluate the model’s performance from both short-term and long-term prediction perspectives. For short-term prediction, we conduct comparative experiments on ten datasets, with the baseline results summarized in Table 1 and Appendix. The results show that DCHO consistently outperforms other baseline models in both MAE and RMSE, demonstrating superior modeling capability. For long-term prediction, as shown in Figure 3 a–b, DCHO consistently maintains leading performance as the prediction horizon extends from 10 to 50 steps, demonstrating strong robustness. We attribute this to two key factors: (1) In the HOBC inference stage, DCHO leverages the dual-view encoder and the latent composition learner to effectively encode the latent information of HOBC; (2) In the latent trajectory prediction stage, DCHO introduces a latent-space prediction loss to model the dynamic evolution of HOBC in an abstract information space.

### 5.3 Non-predictive Downstream Task

Building upon its capacity to predict HOBC, DCHO learns rich information representations of HOBC. We used these higher-order representations to perform state classification across the seven HCP task datasets.

**Experimental Setup:** To validate the effectiveness of the representations extracted by the dual-view encoder, we compare them with three representative baselines: raw fMRI signals, FC, and HOBC. For each type of feature, we train the lightweight MLP classifiers separately, and evaluate performance using five standard metrics: Accuracy, Precision, Recall, F1 Score, and AUROC across seven datasets (Emotion, Gambling, Language, Motor, Relational, Social, WM).

**Performance Analysis:** As shown in Figure 4, our method achieves the best performance across all evaluation metrics on seven downstream tasks. In terms of average performance across all datasets, the representations extracted by DCHO outperform baseline methods, improving accuracy, precision, recall, F1 score, and AUROC by 6.24%,

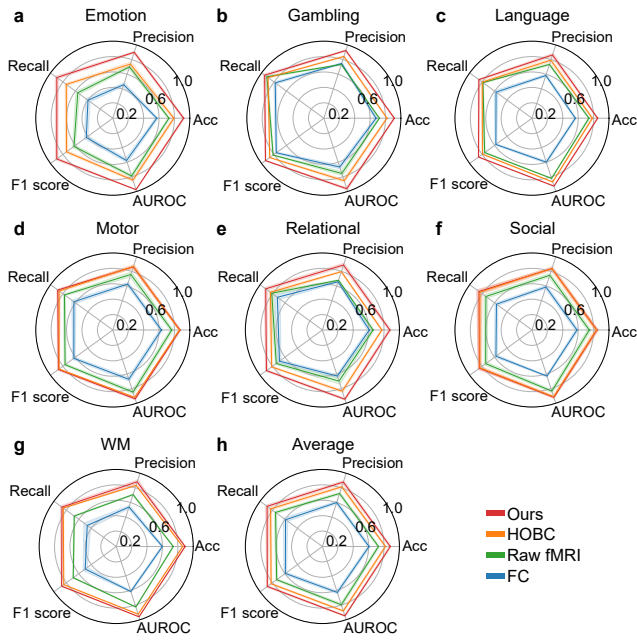


Figure 4: Task-state classification performance across raw fMRI, FC, HOBC, and DCHO representations. DCHO consistently outperforms the baselines across seven cognitive tasks and five evaluation metrics, as illustrated in (a–g), with averaged results in (h).

6.40%, 5.58%, 6.00%, and 6.77%, respectively. These results demonstrate the effectiveness of our learned representations.

#### 5.4 Predictive Downstream Task

Compared to the study (Santoro et al. 2024, 2023) that focus primarily on analyzing HOBC within static or predefined temporal windows, DCHO captures the underlying dynamic evolution of HOBC, enabling more sophisticated predictive tasks.

**Benchmark Methods and Evaluation Metrics:** To evaluate the effectiveness of the pretrained dual-graph encoder and predictor, we compare DCHO against several state-of-the-art methods, including RNN-based models (LtrRNN (Pellegrino, Gajic, and Chadwick 2023), LSTM (Graves and Graves 2012)), Transformer-based models (Transformer (Vaswani et al. 2017), STNDT (Le and Shlizerman 2022) and NetFormer (Lu et al. 2025)), GNN-based models (AMAG (Li et al. 2023)), and a standard MLP (Rosenblatt 1958) baseline. Detailed descriptions of these methods are provided in Appendix. We evaluate model performance using two common metrics (MAE and RMSE) across ten datasets (Emotion, Gambling, Language, Motor, Relational, Social, WM, Rest, Lorenz, HR).

**Performance Analysis:** We systematically evaluate the model’s performance in raw fMRI signal prediction from both short-term and long-term prediction perspectives. For short-term prediction (with a forecast horizon of 10 steps),

we conduct comparative experiments on ten datasets, with baseline results summarized in Table 2 and Appendix. The results demonstrate that DCHO consistently outperforms existing models in terms of MAE and RMSE, showcasing its strong modeling capability. In the long-term prediction task, as the prediction horizon extends from 10 to 50 steps, DCHO maintains its leading performance, exhibiting notable robustness (Figure 3 c-d). We attribute this advantage to two key factors: (1) Compared to baseline methods, the dual-graph encoder in DCHO effectively captures the latent information of HOBC, which inherently reflects nonlinear interactions across brain regions and enhances the modeling of brain dynamics; (2) In the latent space, the predictor captures the underlying evolutionary mechanisms of HOBC, enabling the modeling of deeper and more complex dynamic patterns, which facilitates more accurate long-range forecasting.

#### 5.5 Ablation Study

We conduct an ablation study to evaluate the contributions of the proposed framework and its key components, with evaluations performed on both the task-based (Emotion) and resting-state datasets from HCP dataset. Specifically, we introduce the following model variants: DCHO.a: An end-to-end baseline that directly predicts the future HOBC tensor  $\hat{\mathbf{H}}^{t,t+T}$  from  $\mathbf{X}^{t-T,t}$  and  $\mathbf{A}^{t-T,t}$ , without explicit structural decomposition. DCHO.b: A variant that removes the local topological extractor from the dual-view encoder. DCHO.c: A variant that removes the global topological extractor from the dual-view encoder. DCHO.d: A variant that excludes the latent composition learner. DCHO.e: A variant that replaces the latent-space prediction loss with the original HOBC prediction loss. As shown in the Appendix, experimental results indicate that DCHO consistently outperforms all ablated variants, highlighting the effectiveness of the decomposition–composition framework, the dual-view encoder, and the latent composition learner in modeling HOBC dynamics.

### 6 Conclusion

In this study, we propose DCHO, a novel framework for modeling and forecasting the temporal dynamics of higher-order brain connectivity (HOBC). By decomposing the complex prediction task into two manageable subtasks: HOBC tensor inference and latent trajectory prediction, DCHO effectively addresses the challenges of combinatorial explosion and high-dimensional modeling inherent to HOBC. The framework designs a dual-view encoder and a latent combinatorial learner to capture multiscale topological information and higher-order interaction patterns. Furthermore, it introduces a latent-space prediction loss to enable abstract and robust modeling of HOBC’s temporal evolution. Extensive experiments on multiple datasets demonstrate that DCHO significantly outperforms existing state-of-the-art methods in both predictive tasks and non-predictive tasks, highlighting its strong performance and potential for neurocognitive modeling in real-world applications. Limitations and future work are provided in the Appendix.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (62472206), National Key R&D Program of China (2025YFC3410000), Shenzhen Science and Technology Innovation Committee (RCYX20231211090405003, KJZD20230923115221044), Guangdong Basic and Applied Basic Research Foundation (2025A1515011645 to ZC.L.), Shenzhen Doctoral Startup Project (RCBS20231211090748082 to XK.S.), Guangdong Provincial Key Laboratory of Advanced Biomaterials (2022B1212010003), and the open research fund of the Guangdong Provincial Key Laboratory of Mathematical and Neural Dynamical Systems, the Center for Computational Science and Engineering at Southern University of Science and Technology.

## References

- Alvarez-Rodriguez, U.; Battiston, F.; de Arruda, G. F.; Moreno, Y.; Perc, M.; and Latora, V. 2021. Evolutionary dynamics of higher-order interactions in social networks. *Nature Human Behaviour*, 5(5): 586–595.
- Bassett, D. S.; and Sporns, O. 2017. Network neuroscience. *Nature neuroscience*, 20(3): 353–364.
- Battiston, F.; Amico, E.; Barrat, A.; Bianconi, G.; Ferraz de Arruda, G.; Franceschiello, B.; Iacopini, I.; Kéfi, S.; Latora, V.; Moreno, Y.; et al. 2021. The physics of higher-order interactions in complex systems. *Nature Physics*, 17(10): 1093–1098.
- Battiston, F.; Cencetti, G.; Iacopini, I.; Latora, V.; Lucas, M.; Patania, A.; Young, J.-G.; and Petri, G. 2020. Networks beyond pairwise interactions: Structure and dynamics. *Physics Reports*, 874: 1–92.
- Battiston, F.; and Petri, G. 2022. *Higher-order systems*. Springer.
- Bianconi, G. 2021. *Higher-Order Networks. An Introduction to Simplicial Complexes*. Cambridge University Press.
- Breakspear, M. 2017. Dynamic models of large-scale brain activity. *Nature neuroscience*, 20(3): 340–352.
- Chelaru, M. I.; Eagleman, S.; Andrei, A. R.; Milton, R.; Kharas, N.; and Dragoi, V. 2021. High-order interactions explain the collective behavior of cortical populations in executive but not sensory areas. *Neuron*, 109(24): 3954–3961.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *NeurIPS*, 29.
- Graves, A.; and Graves, A. 2012. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, 37–45.
- Herzog, R.; Rosas, F. E.; Whelan, R.; Fittipaldi, S.; Santamaria-Garcia, H.; Cruzat, J.; Birba, A.; Moguilner, S.; Tagliazucchi, E.; Prado, P.; and Ibanez, A. 2022. Genuine high-order interactions in brain networks and neurodegeneration. *Neurobiology of Disease*, 175: 105918.
- Hindmarsh, J. L.; and Rose, R. 1984. A model of neuronal bursting using three coupled first order differential equations. *Proceedings of the Royal Society of London. Series B. Biological sciences*, 221(1222): 87–102.
- Hutchison, R. M.; Womelsdorf, T.; Allen, E. A.; Bandettini, P. A.; Calhoun, V. D.; Corbetta, M.; Della Penna, S.; Duyn, J. H.; Glover, G. H.; Gonzalez-Castillo, J.; et al. 2013. Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage*, 80: 360–378.
- Jun, J. J.; Steinmetz, N. A.; Siegle, J. H.; Denman, D. J.; Bauza, M.; Barbarits, B.; Lee, A. K.; Anastassiou, C. A.; Andrei, A.; Aydın, Ç.; et al. 2017. Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679): 232–236.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- Le, T.; and Shlizerman, E. 2022. Stndt: Modeling neural population activity with spatiotemporal transformers. *Advances in Neural Information Processing Systems*, 35: 17926–17939.
- Li, J.; Scholl, L.; Le, T.; Rajeswaran, P.; Orsborn, A.; and Shlizerman, E. 2023. Amag: Additive, multiplicative and adaptive graph neural network for forecasting neuron activity. *Advances in Neural Information Processing Systems*, 36: 8988–9014.
- Lu, Z.; Zhang, W.; Le, T.; Wang, H.; Sümbül, U.; SheaBrown, E. T.; and Mi, L. 2025. NetFormer: An interpretable model for recovering dynamical connectivity in neuronal population dynamics. In *The Thirteenth International Conference on Learning Representations*.
- Millán, A. P.; Hanlin, S.; Giambagli, L.; Muolo, R.; Carletti, T.; Torres, J.; Radicchi, F.; Kurths, J.; and Bianconi, G. 2025. Topology shapes dynamics of higher-order networks. *Nature Physics*, 21: 353–361.
- Paulk, A. C.; Kfir, Y.; Khanna, A. R.; Mustroph, M. L.; Trautmann, E. M.; Soper, D. J.; Stavisky, S. D.; Welkenhuysen, M.; Dutta, B.; Shenoy, K. V.; et al. 2022. Large-scale neural recordings with single neuron resolution using Neuropixels probes in human cortex. *Nature neuroscience*, 25(2): 252–263.
- Pellegrino, A.; Gajic, N. A. C.; and Chadwick, A. 2023. Low Tensor Rank Learning of Neural Dynamics. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Roebroek, A.; Formisano, E.; and Goebel, R. 2011. The identification of interacting networks in the brain using fMRI: Model selection, causality and deconvolution. *NeuroImage*, 58(2): 296–302.
- Rosenblatt, F. 1958. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6): 386.
- Santoro, A.; Battiston, F.; Lucas, M.; Petri, G.; and Amico, E. 2024. Higher-order connectomics of human brain function reveals local topological signatures of task decoding, individual identification, and behavior. *Nature Communications*, 15(1): 10244.
- Santoro, A.; Battiston, F.; Petri, G.; and Amico, E. 2023. Higher-order organization of multivariate time series. *Nature Physics*, 19(2): 221–229.

Steinmetz, N. A.; Aydin, C.; Lebedeva, A.; Okun, M.; Pachitariu, M.; Bauza, M.; Beau, M.; Bhagat, J.; Böhm, C.; Broux, M.; et al. 2021. Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, 372(6539): eabf4588.

Tadić, B.; Chutani, M.; and Gupte, N. 2022. Multiscale fractality in partial phase synchronisation on simplicial complexes around brain hubs. *Chaos, Solitons & Fractals*, 160: 112201.

Van Essen, D. C.; Ugurbil, K.; Auerbach, E.; Barch, D.; Behrens, T. E.; Bucholz, R.; Chang, A.; Chen, L.; Corbetta, M.; Curtiss, S. W.; et al. 2012. The Human Connectome Project: a data acquisition perspective. *Neuroimage*, 62(4): 2222–2231.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, H.; Yue, H.; Liu, S.; and Li, T. 2020. Adaptive fixed-time control for Lorenz systems. *Nonlinear Dynamics*, 102(4): 2617–2625.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How powerful are graph neural networks? *In ICLR*.

Yeo, B. T.; Krienen, F. M.; Sepulcre, J.; Sabuncu, M. R.; Lashkari, D.; Hollinshead, M.; Roffman, J. L.; Smoller, J. W.; Zöllei, L.; Polimeni, J. R.; et al. 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*.

Zhang, Y.; Lucas, M.; and Battiston, F. 2023. Higher-order interactions shape collective dynamics differently in hypergraphs and simplicial complexes. *Nature communications*, 14(1): 1605.