

Learning Temporal Higher-order Patterns to Detect Anomalous Brain Activity

Ali Behrouz

Cornell University, NY, USA

AB2947@CORNELL.EDU

Farnoosh Hashemi

Cornell University, NY, USA

SH2574@CORNELL.EDU

Abstract

Due to recent advances in machine learning on graphs, representing the connections of the human brain as a network has become one of the most pervasive analytical paradigms. However, most existing graph machine learning-based methods suffer from a subset of five critical limitations: They are (1) designed for simple pair-wise interactions while recent studies on the human brain show the existence of higher-order dependencies of brain regions, (2) designed to perform on pre-constructed networks from time-series data, which limits their generalizability, (3) designed for classifying brain networks, limiting their ability to reveal underlying patterns that might cause the symptoms of a disease or disorder, (4) designed for learning of static patterns, missing the dynamics of human brain activity, and (5) designed in supervised setting, relying their performance on the existence of labeled data. To address these limitations, we present HADiB, an end-to-end anomaly detection model that automatically learns the structure of the hypergraph representation of the brain from neuroimage data. HADiB uses a tetra-stage message-passing mechanism along with an attention mechanism that learns the importance of higher-order dependencies of brain regions. We further present a new adaptive hypergraph pooling to obtain brain-level representation, enabling HADiB to detect the neuroimage of people living with a specific disease or disorder. Our experiments on Parkinson’s Disease, Attention Deficit Hyperactivity Disorder, and Autism Spectrum Disorder show the efficiency and effectiveness of our approaches in detecting anomalous brain activity.

Keywords: Brain Networks, Temporal Hypergraphs, Anomaly Detection

1. Introduction

Recent advancements in neuroscience and neuroimaging have led researchers to shift from examining isolated brain regions to exploring network models (Bassett and Sporns, 2017). This shift is largely attributed to the rapid progress in technologies like functional Magnetic Resonance Imaging (fMRI) and structural Magnetic Resonance Imaging (sMRI) (Misić and Sporns, 2016), which can provide large-scale neuroimaging data with better quality. In network models of the brain, brain regions of interest (ROIs) are represented as nodes, and the similarities between these regions form edges (Finn et al., 2015). Brain network models have demonstrated their effectiveness in enhancing our understanding of brain diseases and disorders (Chatterjee et al., 2021; Preti et al., 2017). As a result, empirical data on brain networks has substantially increased in size and complexity, leading to a strong demand for appropriate tools and methods to model and analyze this data (Preti et al., 2017)

Simultaneously, machine learning techniques for analyzing graph-structured data have gained attention across various fields, including drug discovery, neuroscience, and biology (Xiong et al., 2019; Abrate and Bonchi, 2021; Gao et al., 2023). Although numerous studies have confirmed the effectiveness of machine learning for human brain network analysis, most have concentrated on graph or node classification tasks (Kan et al., 2022b; Cui et al., 2022). These tasks usually aim at disease detection (Zhu et al., 2022), biological feature prediction (Kan et al., 2022b), or functional system identification (Behrouz and Hashemi, 2022). However, the identification of abnormal brain activity, especially in those with neurological disorders, remains a critical focus for researchers. This is essential for understanding the mechanisms behind symptoms, enabling early detection, and facilitating the development of medical treatments. Most exist-

ing studies consider pairwise interaction among brain regions, neglecting the effect of non-pairwise interactions on the emerging dynamics. Recently, several studies have discussed the importance of higher-order correlation of brain regions. [Rosenthal et al. \(2018\)](#) show that the network context might not be directly accessible at the level of individual regions, and [Santoro et al. \(2023\)](#) show the signatures of higher-order patterns in brain functional activity. Temporal hypergraph representation of the brain from neuroimaging data can address this limitation by capturing both higher-order interactions of brain regions as well as higher-order patterns that are correlated to emerging dynamics of brain activity. Hypergraphs are powerful paradigms to model higher-order interactions, where each connection, also known as hyperedge, can connect a group of nodes at once.

Limitation of Previous Methods. The process of brain network analysis generally unfolds in two sequential steps. Initially, brain networks are derived from individual neuroimage data (e.g., fMRI). This typically involves choosing a brain atlas, which identifies specific ROIs to serve as nodes, and edges that show association between ROIs. For example in fMRI data, from each designated brain region, fMRI blood-oxygen-level-dependent (BOLD) signal sequences are then extracted, and the subsequent phase of edge generation involves calculating pairwise connectivity between these nodes, often based on Pearson correlation and/or mutual information. Finally, the connectivity measures established between node pairs are utilized for the downstream tasks. Although this process along with diverse machine learning methods is widely used in the existing literature ([Kan et al., 2022b, 2021](#); [Behrouz and Seltzer, 2023a](#)), both the pipeline and existing methods suffer from a subset of the following limitations, which makes directly applying them in real-world scenarios challenging and/or impractical:

① Capturing linear correlation and ignoring temporal order: Most existing studies assume that the true dependency structure between brain regions is known prior to model training. That is, existing methods often use Pearson correlation and/or mutual information between the signals of brain regions (e.g, BOLD signals), while there is no statistical measure of dependency for truly capturing functional connectivity ([Mohanty et al., 2020](#)). Moreover, these statistical correlations focus on capturing linear correlation and ignore temporal order, which means shuffling the time in each time window does not change the results. In

order to mitigate the above limitation and to ensure that the model can effectively learn meaningful network representations for use in downstream tasks it is crucial to establish an approach for constructing the dependency structure of the network that accurately reflects the underlying neuroimage data.

② Missing higher-order dependencies: Existing studies assume that the dependencies of brain regions can be captured by pair-wise interactions, while the signatures of higher-order patterns in brain functional activity have been seen in recent studies ([Santoro et al., 2023](#)). Moreover, the network context might not be directly accessible at the level of individual regions and requires considering higher-order interactions between a group of regions ([Rosenthal et al., 2018](#); [Zu et al., 2016](#)). Accordingly, ignoring the higher-order dependencies of ROIs might result in suboptimal performance.

③ Missing hierarchical structure of the brain: The human brain is comprised of functional systems (FS) ([van den Heuvel and Hulshoff Pol, 2010](#)), which are groups of ROIs that perform similar functions as an integrated network ([Smith et al., 2013](#)). These communities are essential in understanding the functional organization of the brain ([van den Heuvel et al., 2009](#)) as ROIs in the same functional modules tend to have similar behaviors and clustered representations ([Smith et al., 2013](#)). Most existing studies neglect the hierarchical structure of the brain, leading to missing the functional dependencies of ROIs in the same FS.

④ Missing the dynamics of the network: Some existing studies neglect the fact that the functional connectivity of the human brain dynamically changes over time, even in resting-state neuroimaging data ([Calhoun et al., 2014](#)). In task-dependent neuroimage data, subjects are asked to perform different tasks in different time windows, and the dynamics of the brain activity during these tasks play an important role in understanding neurological disease/disorder ([Hernandez et al., 2015](#)). Moreover, [Liégeois et al. \(2019\)](#) shows that brain dynamics at different timescales capture distinct aspects. Accordingly, mitigating the above limitation requires a model that not only captures the dynamics of the brain activity but it also can capture its dynamics at different timescales.

⑤ Designed for classification tasks: Most existing studies have focused on semi-supervised/supervised classification tasks, e.g., detecting diseases ([Zhu et al., 2022](#)), predicting biological features ([Kan et al., 2022b](#)), or identifying functional systems ([Behrouz and Hashemi, 2022](#)). This setting not only relies their

performance on the existence of labeled data, but it also limits their ability to reveal underlying patterns that might cause the symptoms of a disease or disorder. That is, a black-box classification method only predicts the label for a given brain network while understanding the cause of the brain disease/disorder (i.e., detecting abnormal brain activity) is a crucial step in facilitating early detection, and developing medical treatments.

Contributions. To address all the above limitations, we design HADiB (**H**ypergraph **A**nomaly **D**etection **i**n **B**rain), an end-to-end unsupervised anomaly detection method that can detect anomalous patterns at different level of granularity. HADiB first uses a novel hierarchical multivariate time-series encoder that can capture both cross-time and cross-ROI dependencies of ROIs’ signals at different time scales. Next, to mitigate ① and ②, HADiB employs a temporal hypergraph constructor that learns the temporal higher-order dependencies between ROIs’ signals. To address the ③ and ④, HADiB uses a tetra-stage message passing at different levels of granularity to learn the ROI, functional system, and brain encodings, taking advantage of the hierarchical structure of the brain. Experimental evaluation supports the need for considering higher-order patterns and shows the superior performance of HADiB over baselines, as well as the importance of HADiB’s critical components. Finally, using real-world case studies we show how HADiB can be used to detect abnormal brain regions or functional systems in a control group with brain disease or disorder.

2. Related Work

Anomaly Detection in Hypergraphs and Time Series. The problem of anomaly detection, which aims to detect abnormal nodes, edges, or subgraphs within a graph, has been extensively studied for both static and temporal graphs (Ma et al., 2021; Akoglu et al., 2015). Surprisingly, hypergraph anomaly detection is relatively unexplored. Park et al. (2009) uses scan statistics to detect anomalous nodes. Leontjeva et al. (2012) suggest taking advantage of the structural features of the nodes to detect anomalies. Recently, Lee et al. (2022) designed a fast and effective algorithm based on a proximity matrix to detect abnormal nodes.

Moreover, several studies use hypergraphs to learn higher-order patterns in time series forecasting (Park

et al., 2009; Sawhney et al., 2021). However, all these methods ① are designed for general applications and cannot take advantage of special properties of the brain, ② cannot detect anomalies at different levels of granularity (e.g., hyperedge, node, etc.).

Machine Learning and Anomaly Detection in Brain Networks. Several studies have recently analyzed brain networks to differentiate between healthy and diseased human brains (Jie et al., 2016; Chen et al., 2011; Wee et al., 2011). With the success of graph neural networks in graph data analysis, deep learning models have been developed to predict brain diseases by studying brain network structures (Behrouz et al., 2023; Kan et al., 2021; Cui et al., 2021; Kan et al., 2022a; Zhu et al., 2022; Cui et al., 2022). However, these models focus on graph or node classification and aren’t directly suited for anomaly detection. There are also methods aimed at detecting anomalies in brain regions or subgraphs, which could signal disease presence (Chatterjee et al., 2021; Zhang et al., 2016; Liu et al., 2020). These studies adopt non-learning approaches and rely on pre-defined patterns or rules for detecting anomalies. Due to the complex nature of brain activity, it is unrealistic to assume that all abnormal brain activities follow the same pre-define pattern or rule. To address this limitation and to learn the abnormal patterns in the human brain from data, recently, Behrouz and Seltzer (2023b) proposed ADMIRE, an unsupervised anomaly detection method that uses multiplex random walks to extract networks motif and to detect anomalous patterns in the brains of people living with a disease or disorder, accordingly.

All these methods target anomaly detection in brain networks with pairwise interactions, missing temporal higher-order patterns that play important roles in understanding the network context and its dynamics over time (Rosenthal et al., 2018; Santoro et al., 2023). Moreover, they do not take advantage of the hierarchical structure of the brain, missing the dependencies of ROIs’ activity in the same functional systems (Smith et al., 2013; van den Heuvel et al., 2009).

3. Methods

Each neuroimage data (e.g., fMRI or EEG) can be represented as a multivariate timeseries $\mathbf{X}_{1:T} = (\mathbf{x}_1, \dots, \mathbf{x}_T) \in \mathbb{R}^{n \times T}$, where n is the number of ROIs. Next, we formally define temporal hypergraphs:




Figure 1: **Schematic of the HADiB.** HADiB consists of three stages: (1) Feature Extraction, (2) Hypergraph Construction, and (3) Hypergraph Learning.

DEFINITION 1 (TEMPORAL HYPERGRAPHS):

A temporal hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X})$, can be represented as a sequence of snapshots $\mathcal{G} = \{\mathcal{G}^{(t)}\}_{t=1}^T$ that arrive over time, where $\mathcal{G}^{(t)} = (\mathcal{V}, \mathcal{E}^{(t)}, \mathcal{X}^{(t)})$ is the t -th snapshot, \mathcal{V} is the set of nodes, $e_i \in \mathcal{E}^{(t)} \subseteq 2^{\mathcal{V}}$ are hyperedges, and $\mathcal{X}^{(t)} \in \mathbb{R}^{|\mathcal{V}| \times f}$ is a matrix that encodes node attribute information for nodes in \mathcal{V} . Note that we treat each hyperedge e_i as the set of all vertices connected by e_i .

Given a hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X})$, we represent the set of hyperedges attached to a node u before time t as $\mathcal{E}^t(u) = \{(e, t') | t' < t, u \in e\}$. We say two hyperedges e and e' are adjacent if $e \cap e' \neq \emptyset$ and use $\mathcal{E}^t(e) = \{(e', t') | t' < t, e' \cap e \neq \emptyset\}$ to represent the set of hyperedges adjacent to e before time t . Each hypergraph snapshot $\mathcal{G}^{(t)}$ can be represented by an adjacency matrix $\mathcal{A}^{(t)} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{E}^{(t)}|}$ such that $\mathcal{A}_{u,e}^{(t)} = 1$ if $u \in e$ and $\mathcal{A}_{u,e}^{(t)} = 0$ otherwise.

We focus on the problem of anomaly detection in the human brain at three levels of granularity: ① Hyperedge anomaly detection, ② ROI anomaly detection, and ③ FS/Brain anomaly detection.

3.1. HADiB: An Unsupervised Anomaly Detection Method in Human Brain

In this section, we design HADiB, an unsupervised anomaly detection method for the human brain, and discuss each of its components in detail. Figure 1 illustrates the overview of HADiB architecture.

Hierarchical Feature Extraction from Time Series. The Time Mixer module aims to encode the multivariate time series data at different time scales.

To this end, we first map the input sequence $\mathbf{X}_{1:T}$ into a hierarchical series $\mathbb{X} = \{\mathbf{X}^{(s)}\}_{s \in \mathcal{S}}$ such that $\mathbf{X}_{1:T_s}^{(s)} = (\mathbf{x}_1^{(s)}, \dots, \mathbf{x}_{T_s}^{(s)}) \in \mathbb{R}^{n \times T_s}$, \mathcal{S} is the set of all scales, and each $s \in \mathcal{S}$ is a time window of $\mathbf{X}^{(s)}$ (i.e., $T_s = \lceil \frac{T}{s} \rceil$). At each time scale $s \in \mathcal{S}$, we need to capture both cross-ROI and cross-time dependencies. Accordingly, inspired by MLP-MIXER (Tolstikhin et al., 2021), we use a Time-Mixer and ROI-Mixer modules to capture cross-time and cross-ROI dependencies, respectively. In the Time-Mixer module, we have:

$$\mathbf{H}_{\text{TIME}}^{(s)} = \mathbf{X}^{(s)} + \mathbf{W}_{\text{TIME}}^{(1)(s)} \sigma \left(\mathbf{W}_{\text{TIME}}^{(2)(s)} \text{LayerNorm} \left(\mathbf{X}^{(s)} \right) \right),$$

where $\mathbf{W}_{\text{TIME}}^{(1)(s)}$ and $\mathbf{W}_{\text{TIME}}^{(2)(s)}$ are learnable metrics, $\text{LayerNorm}(\cdot)$ is layer-normalization Ba et al. (2016), and $\sigma(\cdot)$ is a non-linearity (e.g., GELU (Hendrycks and Gimpel, 2020)). To capture cross-ROI dependencies, we use ROI-Mixer as follows:

$$\begin{aligned} \mathbf{H}_{\text{ROI}}^{(s)} &= \mathbf{H}_{\text{TIME}}^{(s)} \\ &+ \sigma \left(\text{LayerNorm} \left(\mathbf{H}_{\text{TIME}}^{(s)} \right) \mathbf{W}_{\text{ROI}}^{(1)(s)} \right) \mathbf{W}_{\text{ROI}}^{(2)(s)}, \end{aligned} \quad (1)$$

where $\mathbf{W}_{\text{ROI}}^{(1)(s)}$ and $\mathbf{W}_{\text{ROI}}^{(2)(s)}$ are learnable metrics. Note that, while the feature extractor consists of simple all-MLP Time- and ROI-Mixer modules, its architecture (i.e., capturing cross-time and cross-ROI dependencies) make it a powerful model (see § 4).

Finally, we project $\mathbf{H}_{\text{ROI}}^{(s)}$ to have the same size encoding for different time scales:

$$\mathbf{H}_{\text{PROJ}}^{(s)} = \text{MLP} \left(\mathbf{H}_{\text{ROI}}^{(s)} \right). \quad (2)$$

The above procedure encodes the signals at different time scales separately. Accordingly, it does not take advantage the complementary information provided by different time scales (Liégeois et al., 2019). To address this challenge, we concatenate the $\mathbf{H}_{\text{PROJ}}^{(s)}$ for different $s \in \mathcal{S}$ in the order of the time scale, and use an MLP-MIXER module (Tolstikhin et al., 2021) to combine the encodings:

$$\mathbf{H}_F = \text{MLP-MIXER} \left(\bigoplus_{s \in \mathcal{S}} \mathbf{H}_{\text{PROJ}}^{(s)} \right), \quad (3)$$

where \bigoplus denotes concatenation along the ROI dimension (i.e., $\mathbf{H}_F \in \mathbb{R}^{n \times d}$).

Temporal Hypergraph Construction. To capture dependencies between brain regions, while considering the extracted temporal features, we use a self-attention mechanism to learn temporal hypergraph representation of the brain, where each node is a ROI and each hyperedge represents the association of a group of nodes. As discussed in §1, the connectivity of the human brain dynamically changes over time (Calhoun et al., 2014). For example in task-dependent fMRI, subjects are asked to perform a task in each time window. To this end, we use $s_{\max} = \max\{s \in \mathcal{S}\}$, as the time window to construct different snapshots of the network. Accordingly, using zero-padding, we use the above feature extraction procedure for $\mathbf{X}_{1:(t+1)s_{\max}}$ to obtain $\mathbf{H}_F^{(t)}$ for all $0 \leq t \leq \lfloor \frac{T}{s_{\max}} \rfloor$. Next, to construct the t -th snapshot of the hypergraph, we use modified self-attention mechanism (Vaswani et al., 2017) as follows:

$$\mathcal{A}^{(t)} = \text{SIGMOID} \left(\frac{Q^{(t)}K^{(t)}}{\sqrt{K^{(t)^\top}}}, \right), \quad (4)$$

where $Q^{(t)} = \mathbf{H}_F^{(t)} \mathbf{W}_{\text{ATTN}}^{(1)}$, $K^{(t)} = \mathbf{H}_F^{(t)\top} \mathbf{W}_{\text{ATTN}}^{(2)}$, and $\mathbf{W}_{\text{ATTN}}^{(1)}$ and $\mathbf{W}_{\text{ATTN}}^{(2)}$ are learnable matrices. We interpret $\mathcal{A}^{(t)}$ as the adjacency matrix of the hypergraph representation of the brain in the t -th snapshot. Inspired by results reported by Said et al. (2023), we use $\mathcal{X}^{(t)} = \mathbf{H}_F^{(t)}$ as the node features.

Hypergraph Message-Passing. After constructing the temporal hypergraph, to learn the dependencies between temporal patterns of different scales, we propose a tetra-stage message passing mechanism, which contains ① ROI-ROI, ② ROI-Hyperedge, ③ Hyperedge-FS, and ④ FS-FS phases.

① ROI-ROI phase: To learn the local dependencies of ROIs, we iteratively aggregate messages from the

local neighborhood of ROIs. Given t -th snapshot, if u and v are in a hyperedge $e \in \mathcal{E}^{(t)}$, then we update ROI encodings in ℓ -th layer of neural network as follows:

$$\begin{aligned} m_{u \rightarrow v}^{(\ell)} &= \mathbf{W}_{\text{ROI}}^{(\ell)} \text{CONCAT} \left(\hat{\mathbf{h}}_u^{(t)(\ell-1)}, \hat{\mathbf{h}}_v^{(t)(\ell-1)}, \zeta_e^{(t)(\ell-1)} \right), \\ \hat{\mathbf{h}}_u^{(t)(\ell)} &= \text{SUM} \left(\left\{ m_{v \rightarrow u}^{(\ell)} | v \in \mathcal{N}^{(t)}(u) \right\} \right) + \hat{\mathbf{h}}_u^{(t)(\ell-1)}, \end{aligned} \quad (5)$$

where $\mathbf{W}_{\text{ROI}}^{(\ell)}$ is a learnable matrix, $\zeta_e^{(t)(\ell-1)}$ is the encoding (strength) of hyperedge e (we discuss and compute it in ②), and $\mathcal{N}^{(t)}(u)$ is the set of nodes that are connected to u by at least a hyperedge in t -th snapshot. We initialize the $\hat{\mathbf{h}}_u^{(t)(0)} = \mathbf{h}_F^{(t)}(u)$, where $\mathbf{h}_F^{(t)}(u)$ is the corresponding row of matrix $\mathbf{H}_F^{(t)}$ for u .
 ② ROI-Hyperedge phase: To obtain the encoding of a hyperedge, one can simply consider the summation of all ROI's encoding connected by the hyperedge. However, the strength of each dependencies for each ROI is different. Accordingly, we use the attention mechanism in Equation 4 to learn the importance of each hyperedge for each node (ROI) in the hypergraph. Let $\zeta_e^{(t)}$ be the encoding of hyperedge $e \in \mathcal{E}^{(t)}$:

$$\zeta_e^{(t)(\ell)} = \sum_{u \in e} \mathcal{A}_{u,e}^{(t)} \hat{\mathbf{h}}_u^{(t)(\ell)}, \quad (6)$$

where $\mathcal{A}_{u,e}^{(t)}$ is computed by Equation 4.

③ Hyperedge-FS phase: We consider each FS as a set of ROIs. Accordingly, to encode an FS from its ROIs' encoding one might suggest aggregating the encoding of its ROIs. However, this aggregation misses the dependencies of ROIs in each FS. To this end, we propose a pooling function $\text{POOL}(\cdot)$ that aggregates the encodings of hyperedge within each FS. This approach not only incorporates the information from different ROIs, but also considers their dependencies and their roles in the dynamics of the system. Let \mathcal{F} be the set of all functional systems, for $f \in \mathcal{F}$ we use \mathcal{E}_f to denote the set of hyperedges within the f and matrix $\mathbf{Z}_f^{(t)(\ell)} = \bigoplus_{e \in \mathcal{E}_f} \zeta_e^{(t)(\ell)}$. We compute the FS encoding as:

$$\begin{aligned} \hat{\mathbf{Z}}_f^{(t)(\ell)} &= \text{GRU} \left(\mathbf{Z}_f^{(t)(\ell)}, \mathbf{Z}_f^{(t-1)(\ell)} \right) \\ \tilde{\mathbf{Z}}_f^{(t)(\ell)} &= \hat{\mathbf{Z}}_f^{(t)(\ell)} + \sigma \left(\text{Softmax} \left(\text{LayerNorm} \left(\hat{\mathbf{Z}}_f^{(t)(\ell)\top} \right) \right) \right)^\top \\ \Psi_f^{(t)(\ell)} &= \tilde{\mathbf{Z}}_f^{(t)(\ell)} + \sigma \left(\text{LayerNorm} \left(\tilde{\mathbf{Z}}_f^{(t)(\ell)} \right) \mathbf{W}_P^{(1)} \right) \mathbf{W}_P^{(2)} \\ \psi_f^{(t)(\ell)} &= \text{MEAN} \left(\Psi_f^{(t)(\ell)} \right), \end{aligned} \quad (7)$$

where $\text{MEAN}(\cdot)$ is mean function along the hyperedge dimension. The above pooling method uses a simi-

lar architecture as the feature extractor to capture both cross-hyperedge and cross-feature dependencies. However, there are two main differences: First, to reduce the number of parameters, we bind features in a non-parametric manner using a `Softmax(.)` function. Second, the first line, the GRU cell (Chung et al., 2014), is used to update the encodings over time, capturing the dynamics of the functional systems.

④ FS-FS phase: Finally, we perform structural learning at the level of functional systems to capture the dependencies of brain activity at a higher level and use it to obtain brain-level encoding. To this end, we assume that all functional systems are connected and use a self-attention mechanism (Vaswani et al., 2017) to learn the strength of the dependencies of functional systems. Given $\psi_f^{(t)(\ell)}$ for each $f \in \mathcal{F}$, let $\Psi^{(t)(\ell)} \in \mathbb{R}^{|\mathcal{F}| \times d}$ be the matrix whose rows are $\psi_f^{(t)(\ell)}$. We use

$$\mathcal{A}_{\mathcal{F}}^{(t)} = \text{SIGMOID} \left(\frac{Q_{\mathcal{F}}^{(t)} K_{\mathcal{F}}^{(t)}}{\sqrt{K_{\mathcal{F}}^{(T)\top}}} \right), \quad (8)$$

where $Q_{\mathcal{F}}^{(t)} = \Psi^{(t)(\ell)} \mathbf{W}_{\mathcal{F}}^{(1)}$, $K_{\mathcal{F}}^{(t)} = \Psi^{(t)(\ell)} \mathbf{W}_{\mathcal{F}}^{(2)}$, and $\mathbf{W}_{\mathcal{F}}^{(1)}$ and $\mathbf{W}_{\mathcal{F}}^{(2)}$ are learnable matrices. Accordingly, we can obtain the brain-level encoding as the weighted aggregation of functional system encodings:

$$\Upsilon^{(t)(\ell)} = \sum_{f \in \mathcal{F}} \mathcal{A}_{f,f}^{(t)} \psi_f^{(t)(\ell)}. \quad (9)$$

To make the final prediction at each level of granularity, we use an MLP on the encodings.

3.2. Joint Training

In the training phase, we want the model to learn to detect abnormal brain activities at the level of higher-order dependencies (i.e., hyperedges), ROIs (i.e., nodes), and function system or Brain. Accordingly, we use contrastive learning to train the model in an unsupervised manner. We generate negative samples at the level of hyperedges and nodes in the hypergraph. We adopt the commonly used negative sample generation method (Ma et al., 2021; Akoglu et al., 2015) to generate negative hyperedges. That is, for each hyperedge $e = \{u_1, u_2, \dots, u_k\}$ we choose $\alpha\%$ of nodes connected by e and change them to a randomly selected set of vertices. We use binary cross-entropy loss, $\mathcal{L}_{\text{ENTROPY}}$, to learn hyperedge encodings. To learn the time series encoding (i.e., nodes in the

hypergraph), we follow existing studies (Yue et al., 2022; Woo et al., 2022) and replace a brain signal in the time window with another signal that is randomly selected from the batch. We also follow these studies and use the contrastive loss, $\mathcal{L}_{\text{CONTRAST}}$ proposed by Woo et al. (2022).

Loss Function. As we discussed earlier, ROIs in the same functional systems have similar patterns, and accordingly, it is expected to have similar encodings. To this end, inspired by DGI (Veličković et al., 2019), we maximize the mutual information between ROIs encodings and their corresponding functional system encoding. We refer to this loss function as \mathcal{L}_{MI} . Accordingly, we aim to minimize the following loss function: ($\theta_1, \theta_2 \in [0, 1]$ are hyperparameters.)

$$\mathcal{L} = \theta_1 \mathcal{L}_{\text{ENTROPY}} + \theta_2 \mathcal{L}_{\text{CONTRAST}} - (1 - \theta_1 - \theta_2) \mathcal{L}_{\text{MI}}.$$

4. Experiments

Datasets. We use five real-world datasets: ① PD (Day et al., 2019) consists of the functional MRI images of 25 participants with and 21 participants without PD, who do the ANT task (Fan et al., 2005). ② ADHD (Brown et al., 2012) contains data for 100 subjects in the ADHD group and 100 subjects in the typically developed (TD) control group. ③ The Seizure detection TUH-EEG dataset (Shah et al., 2018) consists of EEG data (31 channels) of 642 subjects. ④ ASD (Craddock et al., 2013a) contains data for 45 subjects in the ASD group and 45 subjects in the TD control group. ⑤ ABIDE (Craddock et al., 2013b) consists of resting-state functional MRI of 1009 subjects (516 with ASD) parcellated by Craddock 200 atlas. For the first part of the experiment, we follow the methodology used in existing studies (Akoglu et al., 2015; Behrouz and Seltzer, 2022, 2023b) and inject 1% and 5% anomalous edges into the brain networks in the control group.

Baselines. We compare our HADIb with state-of-the-art time series learning as well as graph and hypergraph anomaly detection and learning. In all tasks, we use ① Graph-based methods: GOutlier (Aggarwal et al., 2011), NETWALK (Yu et al., 2018), BRAINGNN (Li et al., 2021), BRAINNETCNN (Kawahara et al., 2017), ADMIRE (Behrouz and Seltzer, 2023b), and BNTRANSFORMER (Kan et al., 2022b). ② Hypergraph-based methods: HYPER-SAGCN (Zhang et al., 2020), NHP (Yadati et al.,

Table 1: Performance on anomaly detection: Mean AUC (%) \pm standard deviation. Boldfaced letters shaded blue indicate the best result. N/A: The method has numerical precision or computational issues.


Methods	PD		ABIDE		ADHD		TUH-EEG		ASD		
	Anomaly %	1%	5%	1%	5%	1%	5%	1%	5%	1%	5%
Hyperedge Anomaly Detection											
Graph-based Methods											
GOUTLIER	61.42 \pm 1.04	59.98 \pm 2.21	64.08 \pm 1.58	63.32 \pm 1.48	65.37 \pm 0.93	64.70 \pm 2.09	65.61 \pm 1.82	64.12 \pm 0.97	60.85 \pm 0.97	59.13 \pm 1.86	
NETWALK	69.71 \pm 1.99	69.02 \pm 2.83	67.85 \pm 1.72	67.12 \pm 1.23	70.29 \pm 2.15	69.86 \pm 2.58	71.14 \pm 1.36	70.27 \pm 1.42	69.07 \pm 2.20	68.52 \pm 2.55	
BRAINGNN	73.48 \pm 1.75	72.84 \pm 1.52	72.04 \pm 1.27	71.47 \pm 1.59	79.02 \pm 1.85	78.64 \pm 1.43	72.96 \pm 1.58	71.73 \pm 1.14	72.14 \pm 1.25	71.82 \pm 1.73	
BRAINNETCNN	75.04 \pm 1.52	74.63 \pm 0.94	72.74 \pm 0.86	72.31 \pm 1.20	80.58 \pm 1.62	79.95 \pm 2.01	73.06 \pm 1.74	72.87 \pm 1.31	72.68 \pm 2.12	72.01 \pm 1.45	
BN-TRANSFORMER	78.64 \pm 1.01	77.28 \pm 1.24	76.64 \pm 1.28	75.49 \pm 1.37	85.83 \pm 1.97	85.14 \pm 1.67	75.91 \pm 1.72	75.24 \pm 1.53	74.92 \pm 1.18	74.11 \pm 1.37	
ADMIRE	81.36 \pm 1.73	80.98 \pm 2.05	77.26 \pm 1.43	76.58 \pm 1.19	86.23 \pm 1.74	85.18 \pm 2.21	76.68 \pm 1.82	75.14 \pm 1.67	86.52 \pm 1.72	85.44 \pm 1.49	
Hypergraph-based Methods											
NHP	73.24 \pm 0.84	72.65 \pm 1.06	69.32 \pm 1.81	68.79 \pm 1.12	82.61 \pm 0.64	81.22 \pm 1.34	73.07 \pm 1.49	72.24 \pm 0.67	67.98 \pm 0.37	67.17 \pm 1.41	
HYPERSAGCN	75.82 \pm 1.64	74.10 \pm 1.28	72.52 \pm 0.73	71.06 \pm 1.18	84.22 \pm 1.61	83.96 \pm 1.47	73.99 \pm 0.83	72.65 \pm 0.97	73.26 \pm 1.08	73.18 \pm 0.92	
HADIb	82.25\pm1.05	81.12\pm1.14	79.01\pm0.88	78.99\pm1.07	92.23\pm1.33	90.87\pm1.16	78.92\pm1.36	78.04\pm1.05	88.57\pm0.91	89.45\pm1.08	
Time Series-based Methods											
USAD	64.22 \pm 1.07	63.37 \pm 1.28	64.71 \pm 1.75	63.01 \pm 1.61	72.79 \pm 1.48	72.19 \pm 0.94	72.81 \pm 1.42	71.36 \pm 1.03	66.28 \pm 1.16	65.17 \pm 1.15	
MVTS	N/A	N/A	N/A	N/A	N/A	N/A	80.99 \pm 1.36	80.27 \pm 1.49	N/A	N/A	
Graph-based Methods											
GOUTLIER	60.18 \pm 0.54	57.14 \pm 0.89	62.84 \pm 1.15	61.53 \pm 1.08	68.97 \pm 1.16	67.12 \pm 0.93	65.18 \pm 1.09	65.01 \pm 1.57	59.67 \pm 1.42	58.49 \pm 1.35	
NETWALK	68.62 \pm 0.84	67.12 \pm 1.19	67.02 \pm 1.21	66.17 \pm 1.07	75.16 \pm 1.23	74.73 \pm 1.01	72.21 \pm 0.91	71.62 \pm 1.46	71.28 \pm 1.17	71.02 \pm 1.49	
BRAINGNN	71.42 \pm 0.92	71.02 \pm 1.14	70.78 \pm 1.22	70.11 \pm 1.06	81.14 \pm 1.05	80.83 \pm 0.87	73.06 \pm 1.14	72.74 \pm 0.86	72.54 \pm 1.38	71.12 \pm 1.19	
BRAINNETCNN	72.51 \pm 1.16	72.06 \pm 1.27	73.08 \pm 1.38	72.59 \pm 1.03	82.79 \pm 1.08	81.12 \pm 1.16	73.98 \pm 1.24	73.01 \pm 1.08	73.18 \pm 0.95	72.88 \pm 1.04	
BN-TRANSFORMER	73.08 \pm 1.07	72.53 \pm 1.12	76.72 \pm 1.13	76.04 \pm 1.06	86.13 \pm 1.21	86.11 \pm 1.82	77.96 \pm 1.32	77.08 \pm 1.06	76.05 \pm 1.52	75.72 \pm 1.18	
Hypergraph-based Methods											
NHP	70.73 \pm 1.08	69.37 \pm 1.17	70.88 \pm 1.14	70.21 \pm 0.96	83.01 \pm 1.03	82.14 \pm 1.10	73.62 \pm 1.14	73.01 \pm 0.92	72.25 \pm 0.88	71.07 \pm 1.06	
HYPERWALK	72.81 \pm 1.15	71.97 \pm 1.08	73.36 \pm 1.27	72.65 \pm 1.16	83.94 \pm 1.13	83.01 \pm 0.92	75.62 \pm 1.12	74.83 \pm 0.78	74.93 \pm 1.47	74.15 \pm 1.19	
HASHNWALK	72.48 \pm 1.01	72.03 \pm 1.47	71.98 \pm 1.07	70.96 \pm 1.12	83.45 \pm 1.03	83.07 \pm 0.91	74.48 \pm 1.15	73.49 \pm 0.96	73.98 \pm 0.73	73.14 \pm 1.04	
HADIb	79.14\pm1.12	78.57\pm0.97	80.26\pm1.62	79.87\pm1.83	92.34\pm0.97	92.01\pm1.56	82.98\pm1.14	82.10\pm1.09	81.72\pm1.05	80.84\pm1.71	
Time Series-based Methods											
USAD	63.27 \pm 1.66	62.75 \pm 1.38	66.29 \pm 1.24	66.08 \pm 1.13	82.87 \pm 2.03	80.52 \pm 1.84	72.03 \pm 1.17	71.48 \pm 1.05	71.62 \pm 1.58	70.98 \pm 1.41	
MVTS	N/A	N/A	N/A	N/A	N/A	N/A	83.53 \pm 1.91	82.41 \pm 1.02	N/A	N/A	
Graph-based Methods											
NETWALK	62.12 \pm 2.10	61.54 \pm 1.49	66.01 \pm 1.26	65.12 \pm 1.74	83.11 \pm 1.02	82.81 \pm 1.61	71.06 \pm 1.05	69.94 \pm 1.12	72.85 \pm 1.17	72.21 \pm 1.34	
BRAINGNN	67.69 \pm 1.14	67.02 \pm 1.45	69.62 \pm 6.04	69.03 \pm 5.96	84.59 \pm 1.26	83.72 \pm 1.35	72.41 \pm 1.38	71.55 \pm 1.16	75.12 \pm 1.33	74.57 \pm 1.52	
BRAINNETCNN	68.32 \pm 0.86	67.24 \pm 1.08	73.02 \pm 5.04	72.95 \pm 5.47	85.84 \pm 0.96	85.07 \pm 1.52	73.92 \pm 0.97	73.07 \pm 1.51	75.96 \pm 1.66	75.03 \pm 1.28	
BN-TRANSFORMER	70.22 \pm 0.97	68.14 \pm 1.52	76.01 \pm 3.22	75.68 \pm 3.02	87.54 \pm 1.04	86.92 \pm 1.48	79.36 \pm 1.71	78.08 \pm 1.16	77.19 \pm 2.01	76.58 \pm 1.73	
Hypergraph-based Methods											
NHP	67.74 \pm 1.13	67.55 \pm 0.91	70.86 \pm 1.42	69.13 \pm 0.99	84.12 \pm 1.27	84.01 \pm 1.35	74.42 \pm 1.04	74.06 \pm 1.22	75.05 \pm 1.16	74.63 \pm 1.43	
HYPERWALK	68.97 \pm 1.16	68.04 \pm 0.96	74.06 \pm 1.77	73.46 \pm 1.48	86.94 \pm 1.63	86.17 \pm 1.49	75.31 \pm 0.85	74.79 \pm 1.09	76.72 \pm 1.32	75.81 \pm 1.58	
HADIb	76.14\pm1.58	75.82\pm1.37	81.08\pm0.87	80.02\pm1.07	93.52\pm0.98	92.45\pm1.10	83.65\pm1.87	82.54\pm1.07	82.41\pm1.92	81.76\pm1.83	

2020), HASHNWALK (Lee et al., 2022). ③ Time-series-based methods: USAD (Audibert et al., 2020) and MVTS (Potter et al., 2022). We may exclude some baselines in some tasks as they cannot be applied in that setting. For the sake of fair comparison, we use the same training procedure for all the baselines and HADIb.

Quantitative Results. In the first experiment, we compare the performance of HADIb with baselines in hyperedge, ROI, and graph anomaly detection tasks. Table 1 reports the the AUC of HADIb and baselines. HADIb outperforms all the baselines in all three tasks. The reason is fourfold: ① Graph-based methods can only capture the pair-wise dependencies of ROIs, missing higher-order dependencies. ②

Hypergraph-based methods also miss the temporal properties as well as the dynamics of the functional connectivity. ③ HADIb’s training strategy and encoding of the brain at different level of granularity can provide complementary information. ④ HADIb is an end-to-end model that simultaneously learns the hypergraph structure and its element (e.g., node, hyperedge) encodings. However, all the baselines use pre-computed functional connectivity, which limits their generalizability.

Ablation Study. We conduct ablation studies to validate the effectiveness of HADIb’s critical components and to better understand the reason for HADIb superior performance. Table 2 shows the AUC for the hyperedge anomaly detection task. The first row

Figure 2: The distribution of found anomalies by HADI_B in condition groups.Table 2: Ablation study on HADI_B. AUC scores on hyperedge anomaly detection.

Methods	PD	ADHD	ASD
HADI _B	81.12 \pm 1.14	90.87 \pm 1.16	89.45 \pm 1.08
① Remove Feature Extraction	76.52 \pm 2.98	84.80 \pm 1.53	70.09 \pm 0.94
② Remove Time Scales	79.31 \pm 0.85	85.29 \pm 0.97	71.23 \pm 1.19
③ Remove Self-attention	77.41 \pm 1.05	85.98 \pm 1.26	70.68 \pm 1.47
④ Remove GRU	80.79 \pm 0.43	88.23 \pm 1.07	76.08 \pm 1.25
⑤ Remove $\mathcal{L}_{\text{ENTROPY}}$	58.12 \pm 5.89	57.31 \pm 8.03	58.18 \pm 0.98
⑥ Remove $\mathcal{L}_{\text{CONTRAST}}$	59.50 \pm 7.93	72.53 \pm 1.16	68.44 \pm 0.96
⑦ Remove \mathcal{L}_{MI}	65.17 \pm 3.97	80.74 \pm 1.08	69.93 \pm 1.27
⑧ Use Simple Message-passing	77.67 \pm 1.09	84.32 \pm 1.18	74.97 \pm 1.13
⑨ Use Pair-wise Graph	77.64 \pm 0.38	86.94 \pm 1.20	74.13 \pm 1.12

reports the performance of the complete HADI_B implementation. Each subsequent row shows results for HADI_B with one module modification: row 1 replaces feature extraction with Pearson’s correlation, row 2 considers only one time scale, row 3 removes the self-attention mechanism, row 4 removes GRU cell in the pooling, rows 5-7 removes one of the loss functions at each time, rows 8 replaces the tetra stage messaging with a simple message passing, and finally, the last row converts the hypergraph to a graph. These results show that each component is critical for achieving HADI_B’s superior performance. The greatest contribution comes from the training process and loss functions, followed by Feature Extraction modules.

5. Case Studies

In the following case studies, we train our model on the healthy control group and then test it on the condition group to find anomalous brain activities of people living with a specific brain disease/disorder. To this end, we report how anomalous brain activities found by HADI_B are distributed in the brains of people living with PD/ADHD/ASD.

Parkinson’s Disease. Figure 2 reports the distribution of anomalous ROIs within the PD group. A majority (95%) of the identified anomalies by HADI_B

involve ROIs located in one of the Left Thalamus, Supramarginal Gyrus, Superior Parietal, Medial Orbitofrontal, or Pars Opercularis. Interestingly, these ROIs are correlated with some PD symptoms (e.g., affected motor skills). Also, these results are consistent with previous studies using resting-state fMRI (Boutet et al., 2021; Tahmasian et al., 2017).

Attention Deficit Hyperactivity Disorder. Figure 2 reports the distribution of anomalous ROIs within the brain networks of the ADHD group. Most found abnormal ROIs (80% of all found anomalies) by HADI_B are located in the Left Temporal Pole, Frontal Pole, Left Lateral Occipital, and Lingual Gyrus. These findings are consistent with previous studies on ADHD by using diffusion tensor imaging (Lei et al., 2014) and Forman–Ricci curvature changes (Chatterjee et al., 2021).

Autism Spectrum Disorder. The distribution of anomalous ROIs within the brain networks of the ASD group is visualized in Figure 2. More than 80% of all found anomalies by HADI_B involve ROIs located in the Right Cerebellum Cortex, Left Cerebellum Cortex, Right Cerebellum Cortex, Frontal Pole, Left Lateral Occipital, and Right Superior Temporal Gyrus. Our results on findings of abnormal activity in the cerebellum cortex are consistent with previous studies (Rogers et al., 2013).

6. Conclusion

We present HADiB, an end-to-end unsupervised learning method on brain time series data to detect abnormal brain activity at different levels of granularity that might suggest a brain disease or disorder. HADiB uses a hierarchical multivariate time-series encoder to capture both cross-time and cross-ROI dependencies of ROIs' signals at different time scales and employs a temporal hypergraph constructor module on top of that to learn the temporal higher-order dependencies between ROIs. Using a novel hierarchical tetra-stage message passing on the constructed hypergraph, HADiB first learns ROI-level, hyperedge-level, and brain-level encodings and then leverages them to detect abnormal brain activity at different levels. Our experimental results show the superior performance of HADiB against baselines and the potential of HADiB in detecting abnormal brain activity.

Code and Datasets

All codes and implementations are available in [here](#).

References

- Carlo Abrate and Francesco Bonchi. Counterfactual graphs for explainable classification of brain networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery; Data Mining*, KDD '21, page 2495–2504, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383325. doi: 10.1145/3447548.3467154. URL <https://doi.org/10.1145/3447548.3467154>.
- Charu C. Aggarwal, Yuchen Zhao, and Philip S. Yu. Outlier detection in graph streams. In *2011 IEEE 27th International Conference on Data Engineering*, pages 399–409, 2011. doi: 10.1109/ICDE.2011.5767885.
- Leman Akoglu, Hanghang Tong, and Danai Koutra. Graph based anomaly detection and description: a survey. *Data Mining and Knowledge Discovery*, 29(3):626–688, May 2015. ISSN 1573-756X. doi: 10.1007/s10618-014-0365-y. URL <https://doi.org/10.1007/s10618-014-0365-y>.
- Julien Audibert, Pietro Michiardi, Frédéric Guyard, Sébastien Marti, and Maria A Zuluaga. Usad: Unsupervised anomaly detection on multivariate time series. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 3395–3404, 2020.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016.
- Danielle S. Bassett and Olaf Sporns. Network neuroscience. *Nature Neuroscience*, 20(3):353–364, Mar 2017. ISSN 1546-1726. doi: 10.1038/nn.4502. URL <https://doi.org/10.1038/nn.4502>.
- Ali Behrouz and Farnoosh Hashemi. Cs-mlgcn: Multiplex graph convolutional networks for community search in multiplex networks. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management, CIKM '22*, page 3828–3832, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392365. doi: 10.1145/3511808.3557572. URL <https://doi.org/10.1145/3511808.3557572>.
- Ali Behrouz and Margo Seltzer. Anomaly detection in multiplex dynamic networks: from blockchain security to brain disease prediction. In *NeurIPS 2022 Temporal Graph Learning Workshop*, 2022. URL <https://openreview.net/forum?id=UDGZDfwmay>.
- Ali Behrouz and Margo Seltzer. Anomaly detection in human brain via inductive learning on temporal multiplex networks. In *Machine Learning for Healthcare Conference*, volume 219. PMLR, 2023a.
- Ali Behrouz and Margo Seltzer. ADMIRE++: Explainable anomaly detection in the human brain via inductive learning on temporal multiplex networks. In *ICML 3rd Workshop on Interpretable Machine Learning in Healthcare (IMLH)*, 2023b. URL <https://openreview.net/forum?id=t4H8acYudJ>.
- Ali Behrouz, Parsa Delavari, and Farnoosh Hashemi. Unsupervised representation learning of brain activity via bridging voxel activity and functional connectivity. In *NeurIPS 2023 AI for Science Workshop*, 2023. URL <https://openreview.net/forum?id=HSvg7qFFd2>.
- Alexandre Boutet, Radhika Madhavan, Gavin J. B. Elias, Suresh E. Joel, Robert Gramer, Manish Ranjan, Vijayashankar Paramanandam, David Xu, Jürgen Germann, Aaron Loh, Sunil K. Kalia, Mogan Hodaie, Bryan Li, Sreeram Prasad, Ailish

- Coblentz, Renato P. Munhoz, Jeffrey Ashe, Walter Kucharczyk, Alfonso Fasano, and Andres M. Lozano. Predicting optimal deep brain stimulation parameters for parkinson’s disease using functional mri and machine learning. *Nature Communications*, 12(1):3043, May 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-23311-9. URL <https://doi.org/10.1038/s41467-021-23311-9>.
- Jesse A Brown, Jeffrey D Rudie, Anita Bandrowski, John D Van Horn, and Susan Y Bookheimer. The UCLA multimodal connectivity database: a web-based platform for brain connectivity matrix sharing and analysis. *Front Neuroinform*, 6:28, November 2012.
- Vince D. Calhoun, Robyn Miller, Godfrey Pearlson, and Tulay Adalı. The chronnectome: Time-varying connectivity networks as the next frontier in fmri data discovery. *Neuron*, 84(2):262–274, 2014. ISSN 0896-6273. doi: <https://doi.org/10.1016/j.neuron.2014.10.015>. URL <https://www.sciencedirect.com/science/article/pii/S0896627314009131>.
- Tanima Chatterjee, Réka Albert, Stuti Thapliyal, Nazanin Azarhooshang, and Bhaskar DasGupta. Detecting network anomalies using formannic curvature and a case study for human brain networks. *Scientific Reports*, 11(1):8121, Apr 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-87587-z. URL <https://doi.org/10.1038/s41598-021-87587-z>.
- Gang Chen, B Douglas Ward, Chunming Xie, Wenjun Li, Zhilin Wu, Jennifer L Jones, Malgorzata Franczak, Piero Antuono, and Shi-Jiang Li. Classification of alzheimer disease, mild cognitive impairment, and normal cognitive status with large-scale network analysis based on resting-state functional mr imaging. *Radiology*, 259(1):213, 2011.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Cameron Craddock, Yassine Benhajali, Carlton Chu, Francois Chouinard, Alan Evans, András Jakab, Budhachandra Singh Khundrakpam, John David Lewis, Qingyang Li, Michael Milham, et al. The neuro bureau preprocessing initiative: open sharing of preprocessed neuroimaging data and derivatives. *Frontiers in Neuroinformatics*, 7:27, 2013a.
- Cameron Craddock, Yassine Benhajali, Carlton Chu, Francois Chouinard, Alan Evans, András Jakab, Budhachandra Singh Khundrakpam, John David Lewis, Qingyang Li, Michael Milham, et al. The neuro bureau preprocessing initiative: open sharing of preprocessed neuroimaging data and derivatives. *Frontiers in Neuroinformatics*, 7(27):5, 2013b.
- Hejie Cui, Wei Dai, Yanqiao Zhu, Xiaoxiao Li, Lifang He, and Carl Yang. Brainnexplainer: An interpretable graph neural network framework for brain network based disease analysis. *arXiv preprint arXiv:2107.05097*, 2021.
- 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. Springer, 2022.
- Trevor K M Day, Tara M Madhyastha, Mary K Askren, Peter Boord, Thomas J Montine, and Thomas J Grabowski. Attention network test fMRI data for participants with parkinson’s disease and healthy elderly. *F1000Res*, 8:780, June 2019.
- Jin Fan, Bruce D McCandliss, John Fossella, Jonathan I Flombaum, and Michael I Posner. The activation of attentional networks. *Neuroimage*, 26(2):471–479, 2005.
- Emily S. Finn, Xilin Shen, Dustin Scheinost, Monica D. Rosenberg, Jessica Huang, Marvin M. Chun, Xenophon Papademetris, and R. Todd Constable. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, 18(11):1664–1671, Nov 2015. ISSN 1546-1726. doi: 10.1038/nn.4135. URL <https://doi.org/10.1038/nn.4135>.
- Ziqi Gao, Chenran Jiang, Jiawen Zhang, Xiaosen Jiang, Lanqing Li, Peilin Zhao, Huanming Yang, Yong Huang, and Jia Li. Hierarchical graph learning for protein–protein interaction. *Nature Communications*, 14(1):1093, 2023.
- Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus), 2020.
- Leanna M Hernandez, Jeffrey D Rudie, Shulamite A Green, Susan Bookheimer, and Mirella Dapretto. Neural signatures of autism spectrum disorders: insights into brain network dynamics. *Neuropsychopharmacology*, 40(1):171–189, 2015.

- Biao Jie, Mingxia Liu, Xi Jiang, and Daoqiang Zhang. Sub-network based kernels for brain network classification. In *Proceedings of the 7th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 622–629, 2016.
- Xuan Kan, Hejie Cui, Ying Guo, and Carl Yang. Effective and interpretable fmri analysis via functional brain network generation. *arXiv preprint arXiv:2107.11247*, 2021.
- Xuan Kan, Hejie Cui, Joshua Lukemire, Ying Guo, and Carl Yang. Fbnetgen: Task-aware gnn-based fmri analysis via functional brain network generation. *arXiv preprint arXiv:2205.12465*, 2022a.
- Xuan Kan, Wei Dai, Hejie Cui, Zilong Zhang, Ying Guo, and Carl Yang. Brain network transformer. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022b. URL <https://openreview.net/forum?id=1cJ1cbA6NLN>.
- 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.
- Geon Lee, Minyoung Choe, and Kijung Shin. Hash-nwalk: Hash and random walk based anomaly detection in hyperedge streams. In Lud De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 2129–2137. International Joint Conferences on Artificial Intelligence Organization, 7 2022. doi: 10.24963/ijcai.2022/296. URL <https://doi.org/10.24963/ijcai.2022/296>. Main Track.
- Du Lei, Jun Ma, Xiaoxia Du, Guohua Shen, Xingming Jin, and Qiyong Gong. Microstructural abnormalities in the combined and inattentive subtypes of attention deficit hyperactivity disorder: a diffusion tensor imaging study. *Scientific reports*, 4(1):6875, 2014.
- Anna Leontjeva, Konstantin Tretyakov, Jaak Vilo, and Taavi Tamkivi. Fraud detection: Methods of analysis for hypergraph data. In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 1060–1064. IEEE, 2012.
- Xiaoxiao Li, Yuan Zhou, Nicha Dvornek, Muhan Zhang, Siyuan Gao, Juntang Zhuang, Dustin Scheinost, Lawrence H Staib, Pamela Ventola, and James S Duncan. Brainngn: Interpretable brain graph neural network for fmri analysis. *Medical Image Analysis*, 74:102233, 2021.
- Raphaël Liégeois, Jingwei Li, Ru Kong, Csaba Orban, Dimitri Van De Ville, Tian Ge, Mert R. Sabuncu, and B. T. Thomas Yeo. Resting brain dynamics at different timescales capture distinct aspects of human behavior. *Nature Communications*, 10(1):2317, May 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-10317-7. URL <https://doi.org/10.1038/s41467-019-10317-7>.
- Jiaxin Liu, Wei Zhao, Ye Hong, Sheng Gao, Xi Huang, Yingjie Zhou, Alzheimer’s Disease Neuroimaging Initiative, et al. Learning features of brain network for anomaly detection. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 900–905. IEEE, 2020.
- Xiaoxiao Ma, Jia Wu, Shan Xue, Jian Yang, Chuan Zhou, Quan Z. Sheng, Hui Xiong, and Leman Akoglu. A comprehensive survey on graph anomaly detection with deep learning. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–1, 2021. doi: 10.1109/TKDE.2021.3118815.
- Bratislav Mišić and Olaf Sporns. From regions to connections and networks: new bridges between brain and behavior. *Current Opinion in Neurobiology*, 40:1–7, 2016. ISSN 0959-4388. doi: <https://doi.org/10.1016/j.conb.2016.05.003>. URL <https://www.sciencedirect.com/science/article/pii/S095943881630054X>. Systems neuroscience.
- Rosaleena Mohanty, William A. Sethares, Veena A. Nair, and Vivek Prabhakaran. Rethinking measures of functional connectivity via feature extraction. *Scientific Reports*, 10(1):1298, Jan 2020. ISSN 2045-2322. doi: 10.1038/s41598-020-57915-w. URL <https://doi.org/10.1038/s41598-020-57915-w>.
- Youngser Park, C Priebe, D Marchette, and Abdou Youssef. Anomaly detection using scan statistics on time series hypergraphs. In *Link Analysis, Counterterrorism and Security (LACTS) Conference*, page 9. SIAM Pennsylvania, 2009.

- İlkay Yıldız Potter, George Zerveas, Carsten Eickhoff, and Dominique Duncan. Unsupervised multivariate time-series transformers for seizure identification on eeg. In *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 1304–1311. IEEE, 2022.
- Maria Giulia Preti, Thomas AW Bolton, and Dimitri Van De Ville. The dynamic functional connectome: State-of-the-art and perspectives. *Neuroimage*, 160: 41–54, 2017.
- Tiffany D Rogers, Eric McKimm, Price E Dickson, Dan Goldowitz, Charles D Blaha, and Guy Mittleman. Is autism a disease of the cerebellum? an integration of clinical and pre-clinical research. *Front Syst Neurosci*, 7:15, May 2013.
- Gideon Rosenthal, František Váša, Alessandra Griffa, Patric Hagmann, Enrico Amico, Joaquín Goñi, Galia Avidan, and Olaf Sporns. Mapping higher-order relations between brain structure and function with embedded vector representations of connectomes. *Nature Communications*, 9(1):2178, Jun 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-04614-w. URL <https://doi.org/10.1038/s41467-018-04614-w>.
- Anwar Said, Roza G. Bayrak, Tyler Derr, Mudassir Shabbir, Daniel Moyer, Catie Chang, and Xenofon Koutsoukos. Neurograph: Benchmarks for graph machine learning in brain connectomics, 2023.
- Andrea Santoro, Federico Battiston, Giovanni Petri, and Enrico Amico. Higher-order organization of multivariate time series. *Nature Physics*, 19(2): 221–229, Feb 2023. ISSN 1745-2481. doi: 10.1038/s41567-022-01852-0. URL <https://doi.org/10.1038/s41567-022-01852-0>.
- Ramit Sawhney, Shivam Agarwal, Arnav Wadhwa, Tyler Derr, and Rajiv Ratn Shah. Stock selection via spatiotemporal hypergraph attention network: A learning to rank approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 497–504, 2021.
- Vinit Shah, Eva Von Weltin, Silvia Lopez, James Riley McHugh, Lillian Veloso, Meysam Golmohammadi, Iyad Obeid, and Joseph Picone. The temple university hospital seizure detection corpus. *Frontiers in neuroinformatics*, 12:83, 2018.
- Stephen M Smith, Diego Vidaurre, Christian F Beckmann, Matthew F Glasser, Mark Jenkinson, Karla L Miller, Thomas E Nichols, Emma C Robinson, Gholamreza Salimi-Khorshidi, Mark W Woolrich, Deanna M Barch, Kamil Uğurbil, and David C Van Essen. Functional connectomics from resting-state fMRI. *Trends Cogn Sci*, 17(12):666–682, November 2013.
- Masoud Tahmasian, Simon B. Eickhoff, Kathrin Giehl, Frank Schwartz, Damian M. Herz, Alexander Drzezga, Thilo van Eimeren, Angela R. Laird, Peter T. Fox, Habibolah Khazaie, Mojtaba Zarei, Carsten Eggers, and Claudia R. Eickhoff. Resting-state functional reorganization in parkinson’s disease: An activation likelihood estimation meta-analysis. *Cortex*, 92:119–138, 2017. ISSN 0010-9452. doi: <https://doi.org/10.1016/j.cortex.2017.03.016>. URL <https://www.sciencedirect.com/science/article/pii/S0010945217300965>.
- Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Peter Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. MLP-mixer: An all-MLP architecture for vision. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=EI2K0XKdnP>.
- Martijn P. van den Heuvel and Hilleke E. Hulshoff Pol. Exploring the brain network: A review on resting-state fmri functional connectivity. *European Neuropsychopharmacology*, 20(8):519–534, 2010. ISSN 0924-977X. doi: <https://doi.org/10.1016/j.euroneuro.2010.03.008>. URL <https://www.sciencedirect.com/science/article/pii/S0924977X10000684>.
- Martijn P van den Heuvel, René C W Mandl, René S Kahn, and Hilleke E Hulshoff Pol. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Hum Brain Mapp*, 30(10):3127–3141, October 2009.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

- Petar Veličković, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rklz9iAcKQ>.
- Chong-Yaw Wee, Pew-Thian Yap, Wenbin Li, Kevin Denny, Jeffrey N Browndyke, Guy G Potter, Kathleen A Welsh-Bohmer, Lihong Wang, and Dinggang Shen. Enriched white matter connectivity networks for accurate identification of mci patients. *Neuroimage*, 54(3):1812–1822, 2011.
- Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, and Steven Hoi. CoST: Contrastive learning of disentangled seasonal-trend representations for time series forecasting. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=PilZY3omXV2>.
- Zhaoping Xiong, Dingyan Wang, Xiaohong Liu, Feisheng Zhong, Xiaozhe Wan, Xutong Li, Zhaojun Li, Xiaomin Luo, Kaixian Chen, Hualiang Jiang, et al. Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *Journal of medicinal chemistry*, 63(16): 8749–8760, 2019.
- Naganand Yadati, Vikram Nitin, Madhav Nimishakavi, Prateek Yadav, Anand Louis, and Partha Talukdar. Nhp: Neural hypergraph link prediction. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 1705–1714, 2020.
- Wenchao Yu, Wei Cheng, Charu C. Aggarwal, Kai Zhang, Haifeng Chen, and Wei Wang. Network: A flexible deep embedding approach for anomaly detection in dynamic networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, page 2672–2681, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355520. doi: 10.1145/3219819.3220024. URL <https://doi.org/10.1145/3219819.3220024>.
- Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yunhai Tong, and Bixiong Xu. Ts2vec: Towards universal representation of time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8980–8987, 2022.
- Jingyuan Zhang, Bokai Cao, Sihong Xie, Chun-Ta Lu, Philip S Yu, and Ann B Ragin. Identifying connectivity patterns for brain diseases via multi-side-view guided deep architectures. In *Proceedings of the 2016 SIAM International Conference on Data Mining*, pages 36–44. SIAM, 2016.
- Ruochi Zhang, Yuesong Zou, and Jian Ma. Hyper-sagnn: a self-attention based graph neural network for hypergraphs. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ryeHuJBtPH>.
- 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 *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 272–276. IEEE, 2022.
- Chen Zu, Yue Gao, Brent Munsell, Minjeong Kim, Ziwen Peng, Yingying Zhu, Wei Gao, Daoqiang Zhang, Dinggang Shen, and Guorong Wu. Identifying high order brain connectome biomarkers via learning on hypergraph. In *Machine Learning in Medical Imaging: 7th International Workshop, MLMI 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 17, 2016, Proceedings 7*, pages 1–9. Springer, 2016.