Effective and Interpretable fMRI Analysis via Functional Brain Network Generation

Xuan Kan [1]   Hejie Cui [1]   Ying Guo [2]   Carl Yang [1 †]

Abstract

Recent studies in neuroscience show great potential of functional brain networks constructed from fMRI data for popularity modeling and clinical predictions. However, existing functional brain networks are noisy and unaware of downstream prediction tasks, while also incompatible with recent powerful machine learning models of GNNs. In this work, we develop an end-to-end trainable pipeline to extract prominent fMRI features, generate brain networks, and make predictions with GNNs, all under the guidance of downstream prediction tasks. Preliminary experiments on the PNC fMRI data show the superior effectiveness and unique interpretability of our framework.

1. Introduction

In recent year, network-oriented analysis has become increasingly important in neuroimaging studies in order to understand human brain organizations in healthy as well as diseased individuals (Satterthwaite et al., 2015; Wang et al., 2016; Wang & Guo, 2019; Bullmore & Sporns, 2009; Deco et al., 2011). There are abundant findings in neuroscience research showing that neural circuits are the key for understanding the differences in brain functioning between populations and the disruptions in neural circuits largely cause and define brain disorders (Insel & Cuthbert, 2015; Williams, 2016). Functional magnetic resonance imaging (fMRI) is one of the most commonly used imaging modalities to investigate brain functions and organizations (Ganis & Kosslyn, 2002; Lindquist, 2008; Smith, 2012). There is a strong interest in neuroimaging community to predict clinical outcomes or classify individuals based on brain networks derived from fMRI images (Kawahara et al., 2017; Yahata et al., 2016).

[1]Department of Computer Science, Emory University [2]Department of Biostatistics and Bioinformatics, Emory University. Correspondence to: Carl Yang <j.carlyang@emory.edu>.

Current network analysis typically takes the following approach (Smith et al., 2011; Simpson et al., 2013; Wang et al., 2016). First, functional brain networks are estimated based on individuals' fMRI data. This is usually done by selecting a brain atlas or a set of regions of interests and extracting fMRI blood-oxygen-level-dependent (BOLD) signal series from each node or brain regions. Then, pairwise connectivity is calculated between node pairs using measures such as Pearson correlation and partial correlations. The calculated brain connectivity measures between all the node pairs are then used in the subsequent classification or prediction analyses to classify individuals or predict their clinical outcomes. However, the original BOLD signal series are often high-dimensional and noisy, and the brain networks constructed in this way are not customized towards specific downstream clinical predictions.

Recently, there is a growing trend in applying graph neural networks (GNNs) on brain connectivity matrices from fMRI data (Yan et al., 2019; Anirudh & Thiagarajan, 2019; Li et al., 2019; 2020). GNNs are the state-of-the-art deep learning methods on various graph-based tasks which can combine graph structures and node features for various graph-related predictions (Kipf & Welling, 2017; Xu et al., 2019; Veličković et al., 2019; Zhang & Chen, 2018; Yang et al., 2020). Since the mechanism of most GNNs (i.e., message passing) is not compatible with existing functional brain networks which possess both positive and negative weighted edges but no proper node features, recent works also studied how to generate graphs or apply GNNs without pre-computed brain networks (Franceschi et al., 2019; Shang et al., 2021; Yang et al., 2019). Interestingly, (Zhang & Huang, 2019) shows that applying GNNs on randomly generated graphs with proper node features can also provide prediction power, but such random graphs do not capture useful brain connectivity and do not support interpretable clinical analysis and prediction.

In this work, in order to unleash the power of GNNs in network-based fMRI analysis while providing valuable interpretability regarding brain region connectivity, we propose to generate functional brain networks that are compatible with GNNs and customized towards downstream clinical predictions from fMRI data. Specifically, we develop an end-to-end differentiable pipeline from BOLD signal series
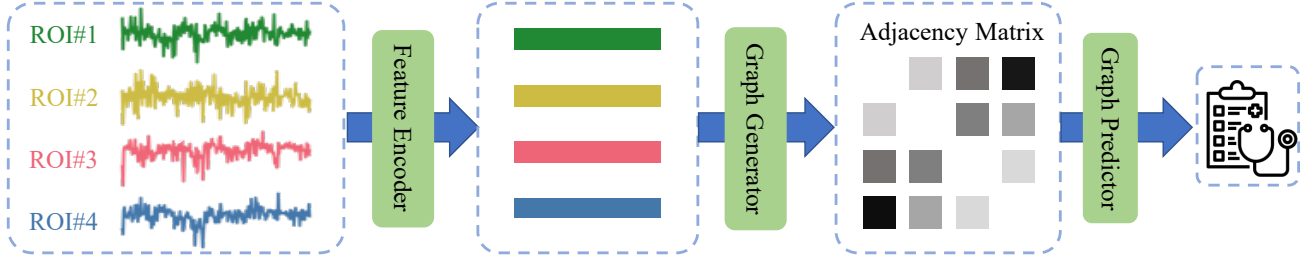
Figure 1: Overall framework of our end-to-end fMRI analysis pipeline with functional brain network generation.

to clinical predictions, which includes a feature extractor for denoising and reducing the dimension of raw time-series data, a graph generator for generating individual brain networks from the extracted features, and a graph predictor of GNN for clinical predictions from the generated brain networks (c.f. Figure 1).

We conduct preliminary experiments on the real-world fMRI dataset of PNC and focus on the downstream task of gender prediction. Our significant performance improvements over state-of-the-art time-series models and GNNs on existing functional brain networks demonstrate the plausible effectiveness of our method, whereas our in-depth analysis with visualizations on the generated brain networks showcase the unique interpretability of our method.

2. Proposed Framework

2.1. Overview

In this section, we elaborate the design of our end-to-end pipeline and its three main components as shown in Figure 1. Specifically, the input $X \in \mathbb{R}^{n \times v \times t}$ denotes the BOLD time-series for regions of interest (ROIs) as the input, where n is the sample size, v is the number of ROIs and t is the length of time-series. Each $x \in \mathbb{R}^{v \times t}$ represents a sample (individual). The target output is the prediction label $Y \in \mathbb{R}^{n \times |\mathcal{C}|}$, where $\mathcal{C}$ is the class set of Y and $|\mathcal{C}|$ is the number of classes. As an intermediate output of the pipeline, we also generate a functional brain network $A \in \mathbb{R}^{v \times v}$ for each sample $x \in \mathbb{R}^{v \times t}$, which represents the brain connectivity matrix between ROIs.

2.2. Feature Encoder

In the feature encoder component, when EXT is set as bi-GRU with window size τ, the process of generating $h_e \in \mathbb{R}^{v \times o}$ can be decomposed as two steps, where o is the feature size of each ROI,

$$z = 1, \cdots, \lfloor \frac{t}{\tau} \rfloor,$$

$$h_r = \text{biRNN}([x^{(z\tau - \tau):z\tau}]), h_r \in \mathbb{R}^{v \times 2\tau},$$

$$h_e = \text{MLP}(h_r),$$

Similarly, when EXT is set as u-layer 1D-CNN, the process can be decomposed like

$$h^u = \text{CONV}_u(h^{u-1}),$$

$$h_e = \text{MLP}(\text{MAXPOOL}(h^u)),$$

where $h^0 = x$ is the original data sample.

2.3. Graph Generator

In the middle stage between encoder and predictor, a feature-based and prediction-driven functional brain network can be generated from h_e, formulated as the connectivity matrix A, which has ROIs along rows and columns and stores the pair-wise connectivity strengths. Unlike the basic (existing) approach for functional brain network construction that calculates the pairwise Pearson correlations between raw time-series of ROIs in x (Smith et al., 2011), we generate A from feature extracted from its time-series data as

$$h_A = \text{softmax}(h_e),$$

$$A = h_A h_A^T, \tag{1}$$

which is end-to-end learnable based on both raw time-series features and downstream prediction tasks. The softmax operation highlights the strong ROI connections, which generates skewed positive edge weights that are compatible with GNNs and valuable for interpretation.

In order to facilitate the learning of brain networks beyond the sheer supervision of graph-based classification, we further apply three group-based and structure-based constraints during training, named as group inner loss, group intra loss and sparsity loss, respectively. Given a class $c \in \mathcal{C}$, $\mathcal{S}^c = \{i \mid Y_{i,c} = 1\}$ is the set containing all sample's index whose label is c. In order to construct group-based regularizers, we need to first obtain μ_c and σ_c^2 as the mean and variance of all samples' A with label c,

$$\mu_c = \sum_{k \in \mathcal{S}^c} \frac{A^k}{|\mathcal{S}^c|},$$

$$\sigma_c^2 = \sum_{k \in \mathcal{S}^c} \frac{\left\| A^k - \mu_c \right\|_2^2}{|\mathcal{S}^c|}.$$

Group Inner Loss aims to minimize the difference of connectivity matrices within a class. Through the following derivation, this loss can be effectively calculated in $\mathcal{O}(n)$ time as

$$L_{inner} = \sum_{c \in \mathcal{C}} \sum_{i \in \mathcal{S}^c} \frac{\left\| \boldsymbol{A}^i - \mu_c \right\|_2^2}{|\mathcal{S}^c|}$$
$$= \sum_{c \in \mathcal{C}} \sigma_c^2.$$

Group Intra Loss aims to maximize the difference of connectivity matrices across different classes, while keeping those within the same class similar. With proper derivation, this loss can also be calculated in $\mathcal{O}(n)$ time as

$$L_{intra} = \sum_{a,b \in \mathcal{C}} \left(\sigma_a^2 + \sigma_b^2 - \frac{\sum_{i \in \mathcal{S}^a} \sum_{j \in \mathcal{S}^b} \left\| \boldsymbol{A}^i - \boldsymbol{A}^j \right\|_2^2}{|\mathcal{S}^a||\mathcal{S}^b|} \right)$$
$$= - \sum_{a,b \in \mathcal{C}} \left\| \mu_a - \mu_b \right\|_2^2.$$

Sparsity Loss aims to enforce the sparsity of brain networks, so as to highlight the task-specific ROI connections, which is formulated as

$$L_{sparsity} = \frac{1}{nvv} \sum_{i=1}^{n} \left\| \text{vec}(\boldsymbol{A}^i) \right\|_1.$$

2.4. Graph Predictor

With the initial node features $\boldsymbol{F} \in \mathbb{R}^{v \times f}$ of ROIs and the learnable connectivity matrix $\boldsymbol{A}$ from graph generator, we can apply a k-layer graph convolutional network (Kipf & Welling, 2017) to compute the node embeddings through

$$\boldsymbol{h}^k = \text{ReLU} \left(\boldsymbol{A} \boldsymbol{h}^{k-1} W^k \right),$$

where W^k represents learnable parameters in convolutional layers and $\boldsymbol{h}^0 = \boldsymbol{F}$. A SUM function is then used to generate the graph-level embedding from all node embeddings. A BatchNorm1D is used to deal with these explored summed values. Finally, another MLP function is employed for prediction,

$$\hat{y} = \text{MLP} \left(\text{BatchNorm1D} \left(\text{SUM} \left(\boldsymbol{h}^k \right) \right) \right).$$

The supervised objective for prediction is the cross-entropy loss L_{ce}. Overall, our total loss function is composed of four terms:

$$L = L_{ce} + \alpha L_{inner} + \beta L_{intra} + \gamma L_{sparsity},$$

where α, β, γ are tunable hyper-parameters representing the weights of different regularizers.

3. Experiments

In this section, we evaluate the effectiveness and interpretability of our proposed framework. The better performance shows our framework's effectiveness and the highlight of task-related brain regions in generated brain networks demonstrates the interpretability of our method.

The purpose of empirical studies on the effectiveness is to answer the following three questions.

- **RQ1.** How does our proposed framework outperform other baseline methods that directly model the time-series features?
- **RQ2.** Do the graphs generated from our framework better facilitate GNNs compared with existing functional brain networks?
- **RQ3.** What influences do the regularizers bring to the performance of our framework?

For interpretability study, our purpose is to investigate the advantages of our learnable graph and the consistency of its patterns with existing neuroscience discoveries.

3.1. Experimental Settings

Dataset. In this paper, we use the PNC dataset (Satterthwaite et al., 2014) for the evaluation. PNC is a collaborative project from the Brain Behavior Laboratory at the University of Pennsylvania and the Children's Hospital of Philadelphia. It includes a population-based sample of individuals aged 8–21 years. Prior to analysis, we perform quality control on the rs-fMRI data such as displacement analysis to remove images with excessive motion (Satterthwaite et al., 2015; Wang et al., 2016). 503 subjects participants' rs-fMRI and DTI data met the quality control criterion and are used in our following analysis. Among these subjects, 289 (57.46%) are female and the mean age is 15.28 years (SD = 3.11). In our paper, we adapt the 264-node cortical parcellation system defined by (Power et al., 2011) for connectivity analysis. The nodes are grouped into 10 functional modules that correspond to the major resting state networks (Smith et al., 2009). Standard pre-processing procedures are applied to the rs-fMRI data. For rs-fMRI, the preprocessing include despiking, slice timing correction, motion correction, registration to MNI 2mm standard space, normalization to percent signal change, removal of linear trend, regressing out CSF, WM, and 6 movement parameters, bandpass filtering (0.009–0.08), and spatial smoothing with a 6mm FWHM Gaussian kernel. Finally, each sample contains 264 nodes and 120 time steps.

Metrics. We perform gender prediction as the evaluation task. Since it is a binary classification problem and our dataset is balanced between classes, AUROC is the most comprehensive performance metric and is adopted here for performance comparison.
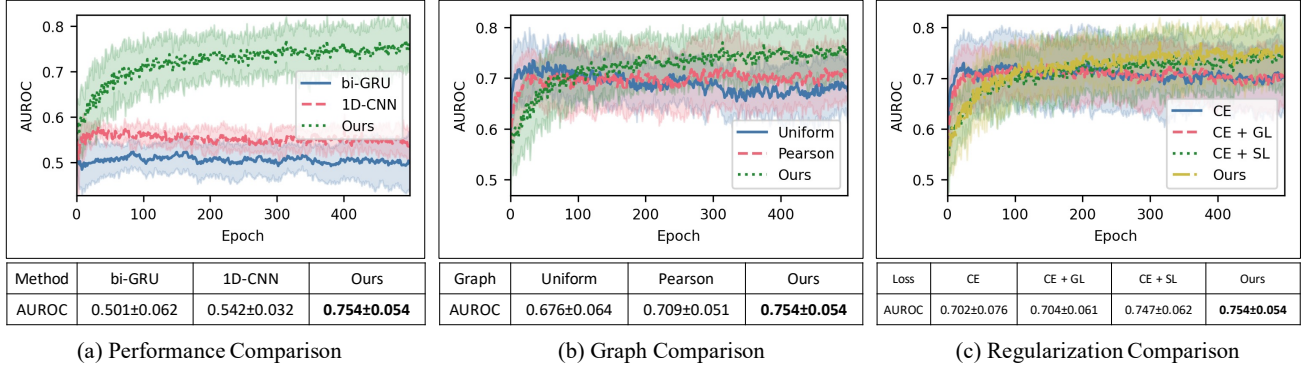
Figure 2: Performances of compared methods. All of the reported scores are evaluated on the test set and are based on 5-fold cross-validation with 3 repetitions. The results shown in the tables are based on the models trained after 500 epochs. (a) Time-series model comparison shows the necessity of intermediate brain network generation. (b) Graph comparison shows that our learnable graph is more compatible with GNN and can achieve better performance than the two others with a large gap. (c) Ablation study: GL represents Group Loss, SL represents Sparsity Loss and CE is Cross Entropy Loss.

Implement Details. The batch size is set as 16 and the learning rate is set as 10^{-4}. α, β, γ are repectively set as $10^{-3}, 10^{-3}$ and 10^{-4}. In the feature encoder, when using 1D-CNN, the number of CNN layer u is 3 and the kernel size is 16. Besides, the kernel size of MAXPOOL is 4. When using GRU as the feature encoder, the window size τ is 16 and the layer number of the GRU is 3. The feature dimension o of h_e is 8. In the graph predictor, we set its layer number as 3 and construct the node feature F for nodes as a vector of Pearson correlation scores between its time series with all nodes contained in the graph. The overall framework is trained with the Adam optimizer with the weight decay of 10^{-4}.

3.2. Comparison with time-series models (RQ1)

We compare our model with two commonly used models, 1D-CNN and bi-GRU, to handle temporal sequence. 1D-CNN and bi-GRU are applied to encode BOLD time series, then both of them use a multilayer perceptron to directly make predictions based on the encoded features without building brain networks. The overall performance is presented in Figure 2(a). To ensure fairness, all hyper-parameters are shared across two baselines and our model. It can be seen that our model outperforms both baseline by significant margins (up to 20% absolute improvements), which demonstrates the necessity of intermediate brain network generation. The comparison between two baselines also provide insights on the choice of feature encoder. As is shown by the curves, 1D-CNN consistently shows a better performance than bi-GRU. Therefore, 1D-CNN is adopted as the feature extractor in our end-to-end framework.

3.3. Comparison with other graphs (RQ2)

We compare our learnable graphs A calculated from Equation 1 with existing brain networks (Smith et al., 2011) and

uniform graphs. For existing brain networks $\mathbf{A}^P$, each entry is calculated as the the Pearson correlation coefficient between two raw time-series,

$$\mathbf{A}^P_{p,q} = |\text{cov}(\boldsymbol{x}_p, \boldsymbol{x}_q)|,$$

where cov represents the co-variance function. To isolate the influence of node features, we also use the uniform graphs $\mathbf{A}^U$ as a control group, which correspond to adjacency matrices with all 1's,

$$\mathbf{A}^U = \{1\}^{v \times v}.$$

All these three types of graphs are paired with the same node features F and then processed by the same graph predictor of GNNs for the task prediction. As is shown in Figure 2(b), GNNs with our learnable graph gains more improvements compared with other graphs, indicating that our learnable graph is more informative and compatible with GNNs.

3.4. Comparison with different model ablations (RQ3)

In this section, we further verify the influence of two major regularizaters in the graph generator: the Group Loss (GL) (including both Group Inner Loss and Group Intra Loss) and the Sparsity Loss (SL). Starting from the CE loss alone, we add one regularizer each time and see how it influences the performance compared with the model. The ablation results are shown in Figure 2(c). We observe that the full model improves more stably than its three downgraded versions and finally achieves the highest performance. Specifically, CE+SL achieves close-to-optimal performance, demonstrating the effectiveness of sparsity constraints in generating informative brain networks. The smaller AUROC standard deviation of CE+CL compared with CE indicates that incorporating group losses can obtain more stable performance.
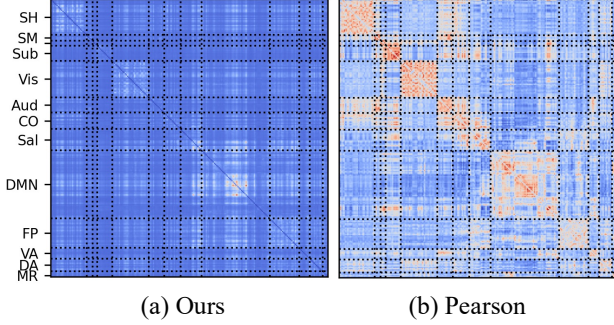
(a) Ours　　　　　(b) Pearson

Figure 3: Visualizations of different brain networks. Abbreviations of neural systems: SH (Somatomotor Hand), SM (Somatomotor Mouth), Sub (Subcortical), Vis (Visual), Aud (Auditory), CO (Cingulo-opercular), Sal (Salience), DMN (Default mode), FP (Fronto-parietal), VA (Ventral attention), DA (Dorsal attention), MR (Memory retrieval). (a) is heatmaps from the mean of learnable brain graphs, while (b) is heatmaps from the mean of Pearson brain graphs.

3.5. Visualization

In this section, we visualize and compare our learnable graphs with the most commonly used existing functional brain networks, i.e., the Pearson Graphs (Smith et al., 2011). Results show that our learnable graph is more task-oriented and advantageous in capturing differences among classes.

Since we can obtain a Pearson graph or learnable graph from each sample, we calculate all samples' learnable graphs and average them to gain the mean of learnable brain graphs to show more meaningful result. From the same steps, we can derive the mean of Pearson brain graphs. Since 32 nodes cannot be grouped into any functional modules. After removing them, 232 of the 264 nodes that were associated with the resting state networks were used for our connectivity analysis. The visualization of the adjacency matrixes with heatmaps are shown in Figure 3. It is observed that our graph distinctively highlights the default mode network (DMN), while in the Pearson graph, the most significant positive components are the connections within functional modules. These within-module connections mainly reflect intrinsic brain functional organizations but are not necessarily informative for the gender prediction task. The highlight of DMN network in our learnable graph is consistent with the neurobiological findings on the PNC data (Satterthwaite et al., 2015) that regions with significant differences between genders are located in the DMN. This comparison indicates that our learnable graphs can learn to highlight the prediction task-specific brain regions and connections, which are potentially useful for the analysis of other clinical prediction tasks where disease-region relevance is unclear.

Besides, to demonstrate that our learnable graphs possess the discrimination ability among classes, in our experiments,

we divide these learnable graphs based on genders. T-test is applied to identify edges $\mathcal{E}^d$ with significant difference ($p < 0.05$) between genders. Next, a difference score T is calculated for each predefined functional module u,

$$T_u = \sum_{(p,q) \in \mathcal{E}^d} \frac{\mathbb{1}(p \in \mathcal{M}_u) + \mathbb{1}(q \in \mathcal{M}_u)}{|\mathcal{M}_u|}$$

where $\mathcal{M}_u$ is a set containing all indexes of node belonging to the module u. Higher score indicates larger difference between genders. The top 3 modules are ventral attention network, default mode network and memory retrieval network. Our findings align well with results from previous work (Satterthwaite et al., 2015) which show that most of the brain regions that demonstrate significant gender differences are located within these three modules. This further validates that our learnable graphs can effectively capture group differences in brain networks, which facilitates imaging-based classification.

4. Conclusion

In this paper, we present a framework which can generate the brain connectivity matrix and predict clinical outcomes simultaneously from fMRI BOLD signal series. The significant improvement in our experiment result presents our framework is more effective than other methods. Besides, the analysis of learnable brain networks shows the generated networks is task-oriented and possess the discrimination ability among classes, which provides a better interpretability. In the near future, it would be interesting to explore the utility of our framework in more brain imaging datasets and downstream prediction tasks.

References

Anirudh, R. and Thiagarajan, J. J. Bootstrapping graph convolutional neural networks for autism spectrum disorder classification. *ICASSP*, pp. 3197–3201, 2019.

Bullmore, E. and Sporns, O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10:186–198, 2009.

Deco, G., Jirsa, V. K., and McIntosh, A. R. Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience*, 12: 43–56, 2011.

Franceschi, L., Niepert, M., Pontil, M., and He, X. Learning discrete structures for graph neural networks. In *ICML*, 2019.

Ganis, G. and Kosslyn, S. M. Neuroimaging. In *Encyclopedia of the Human Brain*, pp. 493–505. 2002.

Insel, T. R. and Cuthbert, B. N. Brain disorders? precisely. *Science*, 348:499–500, 2015.

Kawahara, J., Brown, C. J., Miller, S. P., Booth, B. G., Chau, V., Grunau, R. E., Zwicker, J. G., and Hamarneh, G. BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage*, 146:1038–1049, 2017.

Kipf, T. N. and Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*, 2017.

Li, X., Dvornek, N., Zhou, Y., Zhuang, J., Ventola, P., and Duncan, J. Graph neural network for interpreting task-fmri biomarkers. *MICCAI*, 11768:485–493, 2019.

Li, X., Zhou, Y., Gao, S., Dvornek, N., Zhang, M., Zhuang, J., Gu, S., Scheinost, D., Staib, L., Ventola, P., and Duncan, J. BrainGNN: Interpretable Brain Graph Neural Network for fMRI Analysis. *bioRxiv*, pp. 2020.05.16.100057, 2020.

Lindquist, M. A. The statistical analysis of fmri data. *Statistical science*, 23:439–464, 2008.

Power, J., Cohen, A., Nelson, S., Wig, G., Barnes, K., Church, J., Vogel, A., Laumann, T., Miezin, F., Schlaggar, B., and Petersen, S. Functional Network Organization of the Human Brain. *Neuron*, 72:665–678, 2011.

Satterthwaite, T. D., Elliott, M. A., Ruparel, K., Loughead, J., Prabhakaran, K., Calkins, M. E., Hopson, R., Jackson, C., Keefe, J., Riley, M., Mentch, F. D., Sleiman, P., Verma, R., Davatzikos, C., Hakonarson, H., Gur, R. C., and Gur, R. E. Neuroimaging of the Philadelphia Neurodevelopmental Cohort. *NeuroImage*, 86:544–553, 2014.

Satterthwaite, T. D., Wolf, D. H., Roalf, D. R., Ruparel, K., Erus, G., Vandekar, S., Gennatas, E. D., Elliott, M. A., Smith, A., Hakonarson, H., Verma, R., Davatzikos, C., Gur, R. E., and Gur, R. C. Linked Sex Differences in Cognition and Functional Connectivity in Youth. *Cerebral Cortex*, 25:2383–2394, 2015.

Shang, C., Chen, J., and Bi, J. Discrete graph structure learning for forecasting multiple time series. In *ICLR*, 2021.

Simpson, S. L., Bowman, F. D., and Laurienti, P. J. Analyzing complex functional brain networks: fusing statistics and network science to understand the brain. *Statistics surveys*, 7:1, 2013.

Smith, S. M. The future of fmri connectivity. *Neuroimage*, 62:1257–1266, 2012.

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., and Beckmann, C. F. Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106:13040–13045, 2009.

Smith, S. M., Miller, K. L., Salimi-Khorshidi, G., Webster, M., Beckmann, C. F., Nichols, T. E., Ramsey, J. D., and Woolrich, M. W. Network modelling methods for FMRI. *NeuroImage*, 54, 2011.

Veličković, P., Fedus, W., Hamilton, W. L., Liò, P., Bengio, Y., and Hjelm, R. D. Deep Graph Infomax. In *ICLR*, 2019.

Wang, Y. and Guo, Y. A hierarchical independent component analysis model for longitudinal neuroimaging studies. *NeuroImage*, 189:380–400, 2019.

Wang, Y., Kang, J., Kemmer, P. B., and Guo, Y. An Efficient and Reliable Statistical Method for Estimating Functional Connectivity in Large Scale Brain Networks Using Partial Correlation. *Frontiers in Neuroscience*, 10:123, 2016.

Williams, L. M. Precision psychiatry: a neural circuit taxonomy for depression and anxiety. *The Lancet Psychiatry*, 3:472–480, 2016.

Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How Powerful are Graph Neural Networks? In *ICLR*, 2019.

Yahata, N., Morimoto, J., Hashimoto, R., Lisi, G., Shibata, K., Kawakubo, Y., Kuwabara, H., Kuroda, M., Yamada, T., Megumi, F., Imamizu, H., Sr, J. E. N., Takahashi, H., Okamoto, Y., Kasai, K., Kato, N., Sasaki, Y., Watanabe, T., and Kawato, M. A small number of abnormal brain connections predicts adult autism spectrum disorder. *Nature Communications*, 7:11254, 2016.

Yan, Y., Zhu, J., Duda, M., Solarz, E., Sripada, C., and Koutra, D. Groupinn: Grouping-based interpretable neural network-based classification of limited, noisy brain data. In *KDD*, 2019.

Yang, C., Zhuang, P., Shi, W., Luu, A., and Li, P. Conditional structure generation through graph variational generative adversarial nets. In *NeurIPS*, 2019.

Yang, C., Xiao, Y., Zhang, Y., Sun, Y., and Han, J. Heterogeneous network representation learning: A unified framework with survey and benchmark. *TKDE*, 2020.

Zhang, M. and Chen, Y. Link prediction based on graph neural networks. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *NeurIPS*, 2018.

Zhang, Y. and Huang, H. Information Processing in Medical Imaging, 26th International Conference, IPMI 2019, Hong Kong, China, June 2–7, 2019, Proceedings. *Lecture Notes in Computer Science*, pp. 669–681, 2019.