SUPERVISED GRAPH REPRESENTATION LEARNING FOR MODELING THE RELATIONSHIP BETWEEN STRUCTURAL AND FUNCTIONAL BRAIN CONNECTIVITY

Yang Li[†], Rasoul Shafipour[†], Gonzalo Mateos[†] and Zhengwu Zhang^{*}

[†]Dept. of Electrical and Computer Engineering, University of Rochester, Rochester, NY, USA *Dept. of Biostatistics and Computational Biology, University of Rochester, Rochester, NY, USA

ABSTRACT

In this paper, we propose a supervised graph representation learning method to model the relationship between brain functional connectivity (FC) and structural connectivity (SC) through a graph encoderdecoder system. The graph convolutional network (GCN) model is leveraged in the encoder to learn lower-dimensional node representations (i.e. node embeddings) integrating information from both node attributes and network topology. In doing so, the encoder manages to capture both direct and indirect interactions between brain regions in the node embeddings which later help reconstruct empirical FC networks. From node embeddings, graph representations are learnt to embed the entire graphs into a vector space. Our end-toend model utilizes a multi-objective loss function to simultaneously learn node representations for FC network reconstruction and graph representations for subject classification. The experiment on a large population of non-drinkers and heavy drinkers shows that our model can provide a characterization of the population pattern in the SC-FC relationship, while also learning features that capture individual uniqueness for subject classification. The identified key brain subnetworks show significant between-group difference and support the promising prospect of GCN-based graph representation learning on brain networks to model human brain activity and function.

Index Terms— Brain networks, graph representation learning, graph convolutional network, supervised encoder-decoder system, graph embedding, SC-FC relationship.

1. INTRODUCTION

Our brain is an efficient network with various regions conducting individual tasks while sharing information with each other [1]. Accordingly, graph-centric tools and methods of network science become indispensable for studying brain structure (of neural connections often referred to as the connectome) [2], as well as function [3].

Human brain can be characterized by networks of brain regions connected by functional connectivity (FC) or anatomical tracts (structural connectivity, or SC) [4]. Previous studies have shown that FC correlates with SC at an aggregate level [5], while strong functional connections do not typically imply strong direct SC between the same regions [6]. These findings revealed that functional interactions between brain regions are shaped by higher-level indirect anatomical connections [7] and motivated relative research into predicting FC from SC [5, 6] and inspecting generation process of FC from underlying SC [8,9]. In this context, we model the SC-FC relationship by pursuing SC-to-FC mapping represented by lowerdimensional representations which can be tackled by the state-of-art machine learning frameworks for network data. Recent deep learning developments [10] have significantly improved the performance of various machine learning tasks such as computer vision and natural language processing, due to invariance and stability properties of the well-defined convolution operator for signals supported on regular domains [11]. Driven by the large amount of non-Euclidean data residing on a graph, geometric deep learning [12] was proposed to generalize (convolutional) neural network models for network data [13]. Such generalization motivated the definition of graph convolutional networks (GCNs) with convolutional features computed in the graph spectral and/or spatial domain; see [13, 14] for recent surveys and the references therein.

One important branch of geometric deep learning is graph representation learning [15] which manages to learn lower-dimensional representations of nodes [16–18] for node classification or link prediction, or of graphs [19, 20] for graph classification or similarity ranking. To date, most recent graph representation learning methods are evaluated on social networks [21] and their potential for neuroimaging data analysis is yet to be explored and understood. In particular, generalizing the inductive graph representation learning models is not apparent for brain networks with learnable parameters which are shared among large cohort of graphs.

In this paper, we propose a supervised graph representation learning method to: (i) learn node embeddings generated from SC networks to reconstruct empirical FC networks; and (ii) learn graph embeddings to represent the whole graph for subject classification by building a graph encoder-decoder system. For each node in the graph, the encoder outputs a low-dimensional node embedding that integrates both nodal attributes (when available) and the local graph topology information. To aggregate information from multiple hops within the graph, graph convolution operations are used to compute nodal features in order to capture indirect interactions across the network. This makes the GCN-based encoder a proper model for connectome embeddings by considering indirect (functional) connections within brain networks. The graph embeddings are obtained by averaging all node embeddings [22], and a logistic regression classifier is built for subject classification. We propose this multitask learning model in Section 2.2 with the the goal of reconstructing brain FC using input SC data while achieving subject classification with learnt graph embeddings. The first objective learns a universal SC-FC relationship in the population, while the second objective uses subject labels and takes individual uniqueness into consideration. Consequently, the learnt representations capture the population pattern in SC-FC relationship and also reflect individual properties of each subject regarding the labels used for classification. We train and test our model on a population of 466 subjects with 245 non-drinkers and 221 heavy drinkers from the Human Connectome Project¹, and obtain satisfactory FC reconstruction performance and subject classification accuracy in Section 3. Additionally, via statistical tests on the reconstructed FC networks, key brain sub-networks are identified to show significant difference between groups. Con-

Work in this paper was supported in part by the NSF award CCF-1750428. Emails: yli131@ur.rochester.edu, {rshafipo, gmateosb}@ece.rochester.edu, zhengwu_zhang@urmc.rochester.edu.

¹https://www.humanconnectome.org/

cluding remarks are given in Section 4, along with an outline of future directions related to this promising, but preliminary, line of research. The novelty of our work lies in the learning of lowerdimensional representations that capture the intrinsic quality of the relationship between brain SC and FC networks in a large population for subject classification.

2. PRELIMINARIES AND PROPOSED MODEL

In this section, a brief introduction of graph convolutional models are provided; see also [13,14] for further details. Then we formally state the problem of supervised graph representation learning for modeling SC-FC relationship and introduce the proposed graph encoder-decoder model.

2.1. Graph convolutional networks

Consider a weighted, undirected graph denoted by $\mathcal{G} := (\mathcal{V}, \mathbf{A})$, where \mathcal{V} is a set of N nodes corresponding to brain regions, and $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the symmetric adjacency matrix with \mathbf{A}_{ij} representing the functional (or structural) connection strengths between brain region iand j. The graph Laplacian matrix is defined as $\mathbf{L} := \mathbf{D} - \mathbf{A}$, where \mathbf{D} is the diagonal degree matrix. The Laplacian \mathbf{L} is a symmetric matrix and can be further decomposed as $\mathbf{L} = \mathbf{U}\mathbf{A}\mathbf{U}^T$, where $\mathbf{U} \in \mathbb{R}^{N \times N}$ denotes the set of orthonormal eigenvectors and $\mathbf{\Lambda}$ contains all eigenvalues on its diagonal. As in classical signal processing, the eigenvectors in \mathbf{U} serve as the Fourier basis [23]. Consider a vertex-valued process $\mathbf{x} \in \mathbb{R}^N$ where x_i denotes the signal value at node i, for example, the nodal attributes on the brain FC (or SC) network. Then the graph Fourier transform (GFT) is defined as $\hat{\mathbf{x}} =$ $\mathbf{U}^T \mathbf{x}$ [23,24]. The graph convolution can be defined as

$$\mathbf{H}\mathbf{x} = (\sum_{i=0}^{K} h_i \mathbf{L}^i) \mathbf{x} = \mathbf{U}(\sum_{i=0}^{K} h_i \mathbf{\Lambda}^i) \mathbf{U}^T \mathbf{x} = \mathbf{U}\hat{\mathbf{H}}\hat{\mathbf{x}}$$
(1)

where $\mathbf{H} = \sum_{i=0}^{K} h_i \mathbf{L}^i$ is the graph filter with coefficients $\mathbf{h} := [h_0, \dots, h_K]^T$ and frequency response $\hat{\mathbf{H}} = \sum_{i=0}^{K} h_i \mathbf{\Lambda}^i$. The first spectral convolutional neural network was proposed in

The first spectral convolutional neural network was proposed in [25] to learn the filter response in the spectral domain. However, this method needs to compute the eigenvectors of a (fixed) graph, which may be computationally infeasible for very large networks. Also, the filters depend on the eigenbasis of the Laplacian thus the parameters cannot be shared across different graphs. To overcome these limitations, the ChebNet was proposed in [18] by defining a filter in terms of Chebyshev polynomials of the diagonal matrix of eigenvalues thus the frequency response of the filter is approximated as

$$\hat{\mathbf{H}} = \sum_{i=0}^{K} h_i \mathbf{\Lambda}^i \approx \sum_{i=0}^{K} \theta_i T_i(\tilde{\mathbf{\Lambda}}), \qquad (2)$$

where θ_i are the Chebyshev coefficients to be learned, $\tilde{\mathbf{\Lambda}} = 2\mathbf{\Lambda}/\lambda_{\max} - \mathbf{I}_N$, with λ_{\max} denoting the largest eigenvalue of the Laplacian, and \mathbf{I}_N is the $N \times N$ identity matrix. The Chebyshev polynomials are recursively given as $T_k(\tilde{\mathbf{\Lambda}}) = 2\tilde{\mathbf{\Lambda}}T_{k-1}(\tilde{\mathbf{\Lambda}}) - T_{k-2}(\tilde{\mathbf{\Lambda}})$ with $T_0(\tilde{\mathbf{\Lambda}}) = 1, T_1(\tilde{\mathbf{\Lambda}}) = \tilde{\mathbf{\Lambda}}$. In this case, the graph convolution can be defined as

$$\mathbf{H}\mathbf{x} = \mathbf{U}\hat{\mathbf{H}}\hat{\mathbf{x}} = \mathbf{U}(\sum_{i=0}^{K} \theta_i T_i(\tilde{\mathbf{\Lambda}}))\mathbf{U}^T \mathbf{x} = \sum_{i=0}^{K} \theta_i T_i(\tilde{\mathbf{L}})\mathbf{x}$$
(3)

with $\mathbf{L} = 2\mathbf{L}/\lambda_{\text{max}} - \mathbf{I}_N$. Note that (3) now is K-hop localized with K being the order of the Chebyshev polynomial. This means the convolution operation on a node depends only on other nodes that are at most K hops away in \mathcal{G} .



Fig. 1: The model scheme. Rows in **Y** are low-dimensional node embeddings learned by the encoder from input SC networks. Logistic regression classifier outputs predicted label \hat{l} and compares with actual label l.

A first-order approximation of the ChebNet was introduced in [17] with K = 1, $\lambda_{\max} = 2$, $\theta = \theta_0 = -\theta_1$ to simplify (3) as

$$\mathbf{H}\mathbf{x} = \theta(\mathbf{I}_N + \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2})\mathbf{x}.$$
 (4)

The compact version of (4) motivates the simple per-layer filtering update implemented in [17] of the form

$$\mathbf{X} \leftarrow \mathbf{A}\mathbf{X}\mathbf{\Theta},$$
 (5)

where $\tilde{\mathbf{A}} = \mathbf{I}_N + \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$, \mathbf{X} is a set of multiple observations of the graph signal \mathbf{x} , and Θ stores the learnable weight parameters. The output $\tilde{\mathbf{X}}$ integrates both the nodal attributes in \mathbf{X} and the graph topology information in $\tilde{\mathbf{A}}$. A neural network model based on graph convolution can be built by stacking multiple convolutional layers as in (5), each one followed by a point-wise non-linearity [17]. Next, for our proposed architecture, we build upon the graph convolution in (5) due to its simplicity and remarkable performance [17].

2.2. Problem statement and model architecture

Given brain SC networks, the goal is to build and train a model to reconstruct FC networks through node embeddings and also to learn graph embeddings for subject classification. The learnt representations then contain information of both population patterns and individual uniqueness. To this end, we propose a supervised graph encoder-decoder model as shown in Fig. 1.

Upon the input brain SC network with each node representing one brain region, the encoder generates a lower-dimensional representation for each node. Among various node embedding models, neighbor aggregation methods such as GCN are permutationinvariant and inductive [20]. As a result, GCN is used in the encoder to generate latent variables that consolidates both nodal attributes such as the known intrinsic properties of each brain region, and the network topology information like the connection strengths among regions in SC networks. Here, we use a single-layer GCN for the encoder which is modified from (5) and is given by

$$\mathbf{Y} = \operatorname{Relu}(\mathbf{\hat{A}}\mathbf{X}\mathbf{\Theta}),\tag{6}$$

where $\operatorname{Relu}(x) = \max(0, x)$. In (6), $\mathbf{X} \in \mathbb{R}^{N \times T}$ is the input signal matrix where each row represents a nodal attribute of length T and N is the number of graph nodes. The normalized adjacency matrix is $\hat{\mathbf{A}} := \tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2}$, where $\tilde{\mathbf{A}} = \mathbf{I} + \mathbf{A}$ and $\tilde{\mathbf{D}}$ is the degree matrix of $\tilde{\mathbf{A}}$ (4). Matrix \mathbf{A} here stands for the SC brain network. Identity matrix is added to the input graph \mathbf{A} to introduce self loops to assure that the attribute on the node itself also contributes to the new node embedding during graph convolution. Weight matrix $\mathbf{\Theta} \in \mathbb{R}^{N \times F}$ contains learnable graph filter coefficients for the GCN-based encoder. F represents the length of the node embeddings, which is also the number of filters in the graph convolutional layer. *Relu*

activation function is used to speed up training and avoid the problem of vanishing gradients in the case of deeper architectures.

The learnt node embeddings $\mathbf{Y} \in \mathbb{R}^{N \times F}$ then go through the outer product decoder

$$\mathbf{Z} = \tanh(\operatorname{Relu}(\mathbf{Y}\mathbf{Y}^T)) \tag{7}$$

to generate an estimate of the empirical FC networks via the adjacency matrix **Z**. In the current setup described in Section 3.1, FC networks contain few negative edges with much smaller magnitude compared with the vast amount of positive edges. To avoid such data imbalance problem, we remove all negative FC connections and restrict FC weights to the range [0, 1] as in other works on FC [26,27]. To have the model output in the same range, we choose the combination of *tanh* and *Relu* over *sigmoid* function as the latter suffers from slow training and large errors [9]. Mean squared error (MSE) between the reconstructed graph **Z** and the empirical FC network, denoted as $\mathcal{L}_{MSE}(\mathbf{Z}, \mathbf{FC})$, is used as reconstruction loss for training.

Given the node embeddings \mathbf{Y} , we further obtain the graph embedding by taking the row-wise average of all the node embeddings which results in a vector summarizing SC-FC relationship. Nodewise average pooling is a simple yet effective procedure which has been used in many studies; see e.g. [22]. The graph embeddings are then fed to the logistic regression classifier to predict subject labels. The overall loss function is given as

$$\mathcal{L} = \mathcal{L}_{\text{MSE}}(\mathbf{Z}, \mathbf{FC}) + \lambda \times \mathcal{L}_{\text{CLA}}(\hat{l}, l), \tag{8}$$

where \hat{l} is the predicted label and l is the ground truth label of the subject. Sigmoid cross-entropy loss denoted by \mathcal{L}_{CLA} is used to evaluate classification performance. Hyperparameter λ tunes the trade-off between FC reconstruction and classification performance and its value is determined via grid search in Section 3.2. By training the model end-to-end with the loss function as (8), we aim to find a balance between FC reconstruction and subject classification and achieve acceptable performance on both objectives. As a result, the learnt representations would summarize both the population patterns and the subject uniqueness in the SC-FC relationship. These are further elaborated upon in the next section.

3. NUMERICAL TEST CASES

In this section, the performance of the proposed architecture is evaluated on a real-world neuroimaging dataset.

3.1. Neuroimaging data

A dataset from the Human Connectome Project (HCP) with 466 subjects (245 non-drinkers, 221 heavy-drinkers) is used. Based on the the pipeline in [28], SC network of each subject is extracted from the raw diffusion MRI (dMRI) and structural MRI (sMRI) data. Brain functional activities are measured by the blood oxygen-level dependent (BOLD) signals on each brain region collected by functional MRI (fMRI). Brain FC networks are then constructed with edges weighing the Pearson correlation between the BOLD signals².

The Desikan-Killiany atlas [29] is used to define brain regions thus both FC and SC networks contain N = 68 cortical surface regions with 34 nodes in each hemisphere. In our work we use onehot encoding to define graph signals to avoid data redundancy and accordingly, set $\mathbf{X} = \mathbf{I}_{68}$ in (6). Future work will be devoted to collect and investigate meaningful subject-related information which <u>could be used as additional</u> node attributes on top of the network.

²For more information about the data and preprocessing steps, refer to [28] and https://www.humanconnectome.org/



Fig. 2: Reconstruction (red)/ classification (blue) loss with different λ . Each point on the curve is the average of 10 realizations with corresponding λ .

3.2. Implementation

The supervised graph encoder-decoder system is implemented in TensorFlow. The number of filters F in the graph convolutional layer is set to 32 determined by grid search. Xavier initialization [30] was used to initialize the weight coefficients Θ . To determine the choice of λ in (8), an extensive grid search was carried out with results shown in Fig. 2. As λ increases, the reconstruction performance deteriorates as the training puts more weight on the individual uniqueness for subject classification. When $\lambda = 0$, the model is merely learning population patterns by achieving very good reconstruction performance thus the classification loss is at the peak. As λ increases, the classification loss first decreases, then fluctuates and saturates at a certain level. This is intuitive as the graph embeddings for classification depend on the learnt node embeddings through reconstruction. When the reconstruction is getting worse, the classification performance is confined as expected. In the end, λ was set to be 0.1 based on the grid search outcomes.

With the fixed λ , we use 10-fold cross validation where each time the whole dataset is partitioned randomly into 80% training, 10% validation and 10% testing set. To avoid overfitting, early stopping is applied to monitor the validation loss and stop the training once the validation loss increases in 10 consecutive training epochs. The Adam [31] optimizer is used with learning rate 0.001.

3.3. Results

Reconstruction performance. The model is trained for 10 folds and the average reconstruction MSE is 0.034174 with a standard deviation of 0.00208. The reconstruction performance is secondary to the case with $\lambda = 0$ where the target is solely FC reconstruction. We deem that the slightly higher reconstruction error reflects that the reconstructed FC networks also preserve some key connections that can be used for investigation of between-group difference. We show such significant group-wise difference in the following discussion.

Classification performance. One of the objectives in our work is to learn lower-dimensional graph embeddings, i.e. a single vector to represent the whole graph. Averaging the learnt node embeddings from FC reconstruction, we obtain the graph embeddings as the input to the logistic regression classifier for categorizing subjects into either non-drinker or drinker group. Training the model in Fig. 1 via the 10-fold cross validation gives an average classification accuracy of 67.4% on the test set with 2% standard deviation. After cross validation, the model with the highest test accuracy among the 10 fold is reused to implement on the whole dataset. All the SC networks of the 466 subjects go through the encoder-decoder system and for each subject, we obtain the reconstructed FC network and a 32-by-1 vector modeling the SC-FC relationship. The graph embeddings learnt are visualized in 2D space by t-SNE [32] shown in Fig. 3a. We can distinguish a separation, although not perfectly, between the two groups, indicating that the learnt graph representations contain



Fig. 3: Visualization of graph representations via t-SNE for $\lambda = 0$ and $\lambda = 0.1$. Group-wise separation can be observed in (a) but not in (b).

essential information for subject classification. In Fig. 3b, we also plot the learnt graph embeddings obtained when $\lambda = 0$, i.e., optimizing only the FC reconstruction loss. In that case, all the subjects are mixed together with no clear boundary to separate the two groups.

Using graph embeddings, the model predicts subject labels through logistic regression. We compare the classification results with two baseline methods. The first one is our model with $\lambda = 0$ which boils down to our previous work in [9] and solely captures the population patterns. The second uses node2vec [16] to learn node embeddings, and then takes their average as graph embeddings for logistic regression classification. Table 1 presents statistical measures of binary classification performance of our proposed methods and the compared baseline approaches. ACC represents the accuracy as the proportion of correctly classified labels; TPR (TNR) reflects the true positive (negative) rate describing the percentage of actual positives (negatives) which are correctly identified.

	ACC	TPR	TNR
Model with $\lambda = 0.1$	0.749	0.79	0.7
Model with $\lambda = 0$	0.575	0.6	0.55
Node2vec + logistic regression	0.489	0.55	0.42

 Table 1: Subject classification performance. ACC: accuracy; TPR: true positive rate; TNR: true negative rate.

Table 1 shows that our model outperforms the other two as it is trained in a supervised manner using information from both SC and FC networks. The model with $\lambda = 0$ only focuses on the reconstruction task thus the embeddings only consist of population patterns. Node2vec methods, similar to other embedding techniques [15], can only be applied to a single network thus cannot capture the rich information in the relationship between networks.

Analysis on reconstructed FC networks. Through subject classification, features which help distinguish subject labels are embedded in the learnt representations, thus shall be reflected from the output reconstructed FC networks. To investigate the significance of the between-group difference in the output FC networks, statistical tests are carried out on each functional connection within the FC networks. For each connection, i.e. each element in the $N \times N$ adjacency matrix representing the FC networks, t-test with FDR correction is implemented to check if it shows statistically significant between-group difference (p < 0.05). One sub-network with connections that are stronger in non-drinker group (i.e. weaker in drinker group) is shown in Fig. 4a. Edge weights correspond to the average functional connection strengths in the non-drinker group, and regions with top-10 nodal degree are labeled on the axes. Connections related with regions such as Entorhinal and Parahippocampus show that the limbic system (Parahippocampal gyrus) are impaired in alcohol-dependent subjects [33], which may lead to problems in memory encoding and retrieval. Additionally, a universal decrease in connection strengths regarding regions in frontal, cingulate, parietal, temporal and occipital cortex as visualized in Fig. 4c



Fig. 4: Sub-networks with connections that are significantly (a)(c) weaker, and (b)(d) stronger in heavy drinkers. For visualization purpose, nodes are assigned with different colors and edges use the color of the source nodes.

indicates that alcohol may cause not only changes in regional brain activity but also the patterns of brain functional organization [34].

Another sub-network, shown in Fig. 4b and Fig. 4d, contains connections that are statistically stronger in the drinker group. These connections involve regions in various cortices and may lead to the concept of neural compensation which has been widely studied in cognitive analysis [35]. Neural compensation relates with both structural and functional networks thus is more likely to be revealed in models that analyze both SC and FC networks at the same time. Such a finding indicates that brain may recruit regions and raise additional connections to compensate for damage caused by alcohol [36]. We leave this as a future research direction as it may reveal more about the robustness and resilience of human brain activity.

4. CONCLUSIONS AND FUTURE WORK

The relationship between brain functional and structural networks has long been an intriguing research topic with the valuable potential to reveal the intrinsic patterns of brain activity. In this paper, a supervised graph representation learning procedure was developed by building a GCN-based encoder-decoder system to simultaneously reconstruct FC from SC networks and carry out subject classification. By tuning the balance between reconstruction and classification, the proposed model outperforms baseline methods with higher subject classification accuracy, while also preserving individual uniqueness in the reconstructed FC networks showing betweengroup difference that matches existing literatures. The learnt representations show capability to model SC-FC relationship and capture significant between-group difference. Instead of working on a single type of networks, our model is able to extract hidden features in the relationship between SC and FC networks and information from both SC and FC contributes to the learning process. Future research will be devoted to build advanced models to integrate more information from the node attributes and consider information retrieval from multi-hops within network topology.

5. REFERENCES

- [1] M. Van Den Heuvel and H. Pol, "Exploring the brain network: a review on resting-state fMRI functional connectivity," *Eur. Neuropsychopharmacol.*, vol. 20, no. 8, pp. 519–534, 2010.
- [2] A. Fornito, A. Zalesky, and E. Bullmore, Fundamentals of Brain Network Analysis, Academic Press, 2016.
- [3] J. D. Power, A. L. Cohen, S. M. Nelson, et al., "Functional network organization of the human brain," *Neuron*, vol. 72, no. 4, pp. 665–678, 2011.
- [4] E. Bullmore and O. Sporns, "Complex brain networks: Graph theoretical analysis of structural and functional systems," *Nat. Rev. Neurosci.*, vol. 10, no. 3, pp. 186, 2009.
- [5] C. Honey, O. Sporns, L. Cammoun, X. Gigandet, J.-P. Thiran, R. Meuli, and P. Hagmann, "Predicting human resting-state functional connectivity from structural connectivity," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 106, no. 6, pp. 2035–2040, 2009.
- [6] J. Goñi, M. P. van den Heuvel, A. Avena-Koenigsberger, N. V. de Mendizabal, R. F. Betzel, A. Griffa, P. Hagmann, B. Corominas-Murtra, J.-P. Thiran, and O. Sporns, "Resting-brain functional connectivity predicted by analytic measures of network communication," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 111, no. 2, pp. 833–838, 2014.
- [7] C. Stam, E. Van Straaten, E. Van Dellen, P. Tewarie, et al., "The relation between structural and functional connectivity patterns in complex brain networks," *Intl. J. Psychophysiol.*, vol. 103, pp. 149–160, 2016.
- [8] Y. Li and G. Mateos, "Identifying structural brain networks from functional connectivity: A network deconvolution approach," *IEEE Intl. Conf. Acoust., Speech and Signal Process.* (*ICASSP*), pp. 1135–1139, 2019.
- [9] Y. Li, R. Shafipour, G. Mateos, and Z. Zhang, "Mapping brain structural connectivities to functional networks via graph encoder-decoder with interpretable latent embeddings," *IEEE Global Conf. Signal and Info. Process. (GlobalSIP)*, pp. 1–5, 2019.
- [10] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436, 2015.
- [11] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Adv. Neural. Inf. Process. Syst.*, pp. 1097–1105, 2012.
- [12] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: going beyond Euclidean data," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 18–42, 2017.
- [13] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *arXiv* preprint arXiv:1901.00596, 2019.
- [14] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, "Convolutional neural network architectures for signals supported on graphs," *IEEE Trans. Signal Process.*, vol. 67, no. 4, pp. 1034– 1049, 2018.
- [15] W. L. Hamilton, R. Ying, and J. Leskovec, "Representation learning on graphs: Methods and applications," *IEEE Data Engineering Bulletin*, 2017.
- [16] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining (KDD), pp. 855–864, 2016.
- [17] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *Proc. ICLR*, 2017.

- [18] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," *Adv. Neural. Inf. Process. Syst.*, pp. 3844–3852, 2016.
- [19] A. Narayanan, M. Chandramohan, R. Venkatesan, L. Chen, Y. Liu, and S. Jaiswal, "graph2vec: Learning distributed representations of graphs," *arXiv preprint arXiv:1707.05005*, 2017.
- [20] Y. Bai, H. Ding, Y. Qiao, A. Marinovic, K. Gu, T. Chen, et al., "Unsupervised inductive graph-level representation learning via graph-graph proximity," *Intl. Joint Conf. Artificial Intelligence*, pp. 1988–1994, 2019.
- [21] X. Yue, Z. Wang, et al., "Graph embedding on biomedical networks: Methods, applications, and evaluations," *arXiv preprint arXiv:1906.05017*, 2019.
- [22] D. Duvenaud, D. Maclaurin, et al., "Convolutional networks on graphs for learning molecular fingerprints," *Adv. Neural. Inf. Process. Syst.*, pp. 2224–2232, 2015.
- [23] A. Ortega, P. Frossard, et al., "Graph signal processing: Overview, challenges, and applications," *Proc. IEEE*, vol. 106, no. 5, pp. 808–828, 2018.
- [24] W. Huang, T. A. Bolton, J. D. Medaglia, et al., "A graph signal processing perspective on functional brain imaging," *Proc. IEEE*, vol. 106, no. 5, 2018.
- [25] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, "Spectral networks and locally connected networks on graphs," *Proc. ICLR*, 2014.
- [26] J. D. Power, D. A. Fair, B. L. Schlaggar, and S. E. Petersen, "The development of human functional brain networks," *Neuron*, vol. 67, no. 5, pp. 735–748, 2010.
- [27] M. Rubinov and O. Sporns, "Complex network measures of brain connectivity: uses and interpretations," *Neuroimage*, vol. 52, no. 3, pp. 1059–1069, 2010.
- [28] Z. Zhang, G. I. Allen, H. Zhu, and D. Dunson, "Tensor network factorizations: Relationships between brain structural connectomes and traits," *Neuroimage*, vol. 197, pp. 330–343, 2019.
- [29] R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, et al., "An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest," *Neuroimage*, vol. 31, no. 3, pp. 968–980, 2006.
- [30] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feed forward neural networks," *Intl. Conf. on Artificial Intelligence and Statistics*, pp. 249–256, 2010.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 2015.
- [32] L. Maaten and G. Hinton, "Visualizing data using t-SNE," J. Mach. Learn. Res., vol. 9, no. Nov, pp. 2579–2605, 2008.
- [33] M. Oscar-Berman and K. Marinkovic, "Alcoholism and the brain: an overview.," *Alcohol Research & Health*, vol. 27, no. 2, pp. 125–133, 2003.
- [34] N. D. Volkow, Y. Ma, W. Zhu, J. S. Fowler, J. Li, et al., "Moderate doses of alcohol disrupt the functional organization of the human brain," *Psychiatry Res. Neuroimaging*, vol. 162, no. 3, pp. 205–213, 2008.
- [35] D. Barulli and Y. Stern, "Efficiency, capacity, compensation, maintenance, plasticity: emerging concepts in cognitive reserve," *Trends Cogn. Sci.*, vol. 17, no. 10, pp. 502–509, 2013.
- [36] M. H. Parks, V. L. Morgan, Pickens, et al., "Brain fMRI activation associated with self-paced finger tapping in chronic alcohol-dependent patients," *Alcohol.: Clin. Exp. Res.*, vol. 27, no. 4, pp. 704–711, 2003.