

MAPPING BRAIN STRUCTURAL CONNECTIVITIES TO FUNCTIONAL NETWORKS VIA GRAPH ENCODER-DECODER WITH INTERPRETABLE LATENT EMBEDDINGS

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, the relationship between functional and structural brain networks is investigated by training a graph encoder-decoder system to learn the mapping from brain structural connectivity (SC) to functional connectivity (FC). Our work leverages a graph convolutional network (GCN) model in the encoder which integrates both nodal attributes and the network topology information to generate new graph representations in lower dimensions. Using brain SC graphs as inputs, the novel GCN-based encoder-decoder system manages to account for both direct and indirect interactions between brain regions to reconstruct the empirical FC networks. In doing so, the latent variables within the system (i.e., the learnt low-dimensional embeddings) capture important information regarding the relation between functional and structural networks. By decomposing the reconstructed functional networks in terms of the output of each graph convolution filter, we can extract those brain regions which contribute most to the generation of FC networks from their SC counterparts. Experiments on a large population of healthy subjects from the Human Connectome Project show our model can learn a generalizable and interpretable SC-FC relationship. Overall, results here support the promising prospect of using GCNs to discover more about the complex nature of human brain activity and function.

Index Terms— Brain networks, graph convolutional network, encoder-decoder system, graph embedding, structural connectivity-functional connectivity relationship.

1. INTRODUCTION

Understanding brain function represents one of the most fundamental and pressing scientific challenges of our time. Driven by advances in neuroimaging technology, brain data have increased in volume and complexity, and accordingly graph-centric tools and methods of network science have become indispensable for mapping and modeling brain structure (of neural connections often referred to as the connectome) [1, 2], as well as function [3].

Brain connectivity broadly consists of networks of brain regions connected by functional associations (functional connectivity, or FC) or anatomical tracts (structural connectivity, or SC) [4]. FC captures statistical dependencies shaped by an underlying structural backbone, including higher-level interactions between brain regions with no or limited structural connections [5, 6]. As a result, exploring the mapping between FC and SC is a popular topic for research as such relation could offer important additional insights on the inner workings of the brain. Previous studies revealed that functional connectivity correlates with structural connectivity at an aggregate level [7], while strong functional connections exist among brain regions that

are not directly structurally linked or have rather limited anatomical connections between themselves [8]. These findings provide strong evidence that functional interactions between brain regions are shaped by indirect anatomical connections [9]. This also motivated the problem of predicting FC from SC [5, 7, 8] and to inspect how brain FC is generated in the context of the underlying physical brain fiber connections [10]. In this context, our novel perspective here is to view the pursuit of the SC-to-FC mapping as a regression problem, which can be tackled with state-of-the-art machine learning algorithms for (non-Euclidean) network data.

Recent machine learning developments in computer vision, speech processing, and natural language processing have brought significant performance improvements in various tasks [11], due to invariance and stability properties of the well-defined convolution operator for signals supported on regular domains [12]. To accommodate nowadays ubiquitous non-Euclidean data residing on a graph, there have been attempts to generalize (convolutional) neural networks models for network data [13]. This body of work has been sometimes referred to as geometric deep learning [14]. A graph convolution operator was first introduced in the spectral domain using eigenvectors of the graph Laplacian matrix [15], which play a similar role to the Fourier basis for information processing of time-varying signals. This generalization motivated the definition of graph convolutional networks (GCNs) with convolutional features computed in the graph spectral and/or spatial domain; see [13, 16] for recent surveys and the references therein. While the benefits of GCNs have been well documented in applications ranging from recommendation systems to resource allocation for wireless communications, their potential for neuroimaging data analyses is yet to be explored and fully realized.

In this paper, we put forth an unsupervised learning method to reconstruct brain FC patterns from SC networks by building a graph encoder-decoder system. The proposed architecture is inspired by the graph autoencoder model [17], where the goal is to recover input graphs using a GCN-based encoder. For each node of the graph, the encoder outputs a low-dimensional feature vector (also known as node embedding) [18]. Such compact node representations integrate both nodal attributes (when available) and the local graph topology information. By computing nodal features via graph convolution operations, information is aggregated from multiple hops within the graph, thus capturing indirect interactions across the network [19]. Such property makes the GCN-based encoder a suitable model to capture indirect (functional) connections within brain networks. Inspired by [20], which shows that connectome embeddings using the node2vec algorithm [21] can map higher-order relations between SC and FC, we propose a graph encoder-decoder system in Section 2.2 to reconstruct brain FC using input SC data. We train and test our model on a population of 1058 healthy subjects from the Human Connectome Project¹, and obtain satisfactory reconstruction performance in Section 3. Additionally, the latent variables (i.e., the

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.

¹<https://www.humanconnectome.org/>

graph convolutional features) learned by the encoder-decoder system can summarize the relation between brain SC and FC as well as provide valuable insights regarding the generation of FC patterns from SC. By further decomposing the reconstructed FC network into small building blocks, each of which is the output of one graph convolution filter [22], we are able to extract certain brain regions and sub-networks that contribute to the SC-FC relationship. We contend these regions may play a vital role in shaping brain functional activity and may be fruitfully used later for a downstream patient-control classification task. Concluding remarks are given in Section 4, along with an outline of future directions related to this promising, but admittedly preliminary, line of research. To the best of our knowledge, this is the first time that a GCN model is applied to study the relation between brain structural and functional connectivities.

2. PRELIMINARIES AND PROPOSED MODEL

In this section, we first review some established graph convolutional models used to build GCN architectures; see also [13, 16] for further details. Then we formally state the problem of learning the SC-to-FC mapping 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 $\mathbf{A}$ is the symmetric adjacency matrix with $\mathbf{A}_{ij}$ representing the functional (or structural) connection strengths between brain region i and 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 $\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, where $\mathbf{U}$ 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 $\tilde{\mathbf{x}} = \mathbf{U}^T \mathbf{x}$ [23, 24]. Given a filter $\mathbf{h} \in \mathbb{R}^N$, the graph convolution operator can be defined as

$$\mathbf{x} * \mathbf{h} = \mathbf{U}((\mathbf{U}^T \mathbf{h}) \odot (\mathbf{U}^T \mathbf{x})) \quad (1)$$

where $\odot$ denotes the element-wise multiplication. The definition in (1) comes from the analogy of convolution among temporal signals, where convolution in the time domain becomes element-wise multiplication of the signal and filter's Fourier transform.

The first spectral convolutional neural network was proposed in [15] 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 [25] by defining a filter in terms of Chebyshev polynomials of the diagonal matrix of eigenvalues expressed as

$$h(\tilde{\mathbf{\Lambda}}) = \sum_{i=0}^K \theta_i T_i(\tilde{\mathbf{\Lambda}}), \quad (2)$$

where θ_i are the polynomial filter coefficients to be learned, $\tilde{\mathbf{\Lambda}} = 2\mathbf{\Lambda}/\lambda_{\max} - \mathbf{I}_N$, with $\lambda_{\max}$ 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 $T_k(x) = 2xT_{k-1}(x) - T_{k-2}(x)$ with $T_0(x) = 1$, $T_1(x) = x$. In this case, the graph convolution can be defined as

$$\mathbf{x} * \mathbf{h} = \sum_{i=0}^K \theta_i T_i(\tilde{\mathbf{L}}) \mathbf{x} \quad (3)$$




Fig. 1: The model scheme. The rows in $\mathbf{Y}$ are the low-dimensional nodal embeddings learned by the encoder from the input structural networks.

with $\tilde{\mathbf{L}} = 2\mathbf{L}/\lambda_{\max} - \mathbf{I}_N$. 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 $\mathcal{G}$.

A first-order approximation of the ChebNet was introduced in [26] with $K = 1$, $\lambda_{\max} = 2$, $\theta = \theta_0 = -\theta_1$ to simplify (3) as

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

The compact version of (4) motivates the simple per-layer filtering update implemented in [26] of the form

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

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 filter 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 [26]. Next, for our proposed architecture, we build upon the graph convolution in (5) due to its simplicity and remarkable performance [26].

2.2. Problem statement and model architecture

Given a brain SC network, the goal is to build and train a model to reconstruct the brain FC network and learn the SC-to-FC mapping. The latent variables within the model can provide novel insights regarding the generation of FC patterns from SC. To this end, we propose a graph encoder-decoder model as shown in Fig. 1.

With the input brain SC network where each node represents one brain region, the encoder generates a lower dimension representation for each node. Being a suitable method to capture indirect connections within brain networks as in (3), GCN is used here in the encoder to generate latent variables that consolidate 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. In our architecture, a single-layer GCN modified from (5) is used for the encoder, given by

$$\mathbf{Y} = \text{Relu}(\tilde{\mathbf{A}} \mathbf{X} \Theta) \quad (6)$$

where $\text{Relu}(x) = \max(0, x)$. 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 $\tilde{\mathbf{A}} := \tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2}$, where $\tilde{\mathbf{A}} = \mathbf{I} + \mathbf{A}$ and $\tilde{\mathbf{D}}$ is the degree matrix of $\tilde{\mathbf{A}}$ (4). Matrix $\mathbf{A}$, in our case, stands for the SC brain network. Note that the identity matrix is added to the input graph $\mathbf{A}$, introducing self loops so that during the graph convolution, the attribute on the node itself could also contribute to the new node embedding. The weight matrix $\Theta \in \mathbb{R}^{N \times F}$ collects the filter coefficients for the

GCN to be learned during training. The hyperparameter F stands for the length of the node embeddings, which is also the number of filters in the graph convolutional layer. It is a common practice to have $F < T$ to perform dimensionality reduction while the neighboring information is efficiently summarized. The *Relu* activation function [27] is used to speed up the training and avoid the problem of vanishing gradients.

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

$$\mathbf{Z} = \tanh(\text{Relu}(\mathbf{Y}\mathbf{Y}^T)) \quad (7)$$

to approximate the empirical FC networks via the adjacency matrix estimate $\mathbf{Z}$. The choice of the activation functions for the decoder highly depends on the nature of the data. In the present setup described in Section 3.1, the functional networks contain limited number of negative edges with much smaller magnitude compared with the vast amount of positive edges. Using only *tanh* activation while keeping the negative FC connections generates graphs with positive edge weights. This situation resembles the classical data imbalance problem in machine learning where classifiers are more biased towards the majority classes. As a result, we restrict the FC weights to the range $[0, 1]$ by removing all the negative FC connections as in previous works dealing with functional brain connectivity [28, 29]. To ensure that the output of the decoder is in the same range, we select the combination of *tanh* and *Relu* over *sigmoid* function as the latter results in slower training and larger test errors. Mean squared error (MSE) between the reconstructed graph $\mathbf{Z}$ and the desired output, i.e. the empirical FC network is used as the loss function for training the model. MSE has been widely used and proven effective for weighted graph reconstruction [19].

Note that (7) can also be written as $\mathbf{Z} = \tanh(\text{Relu}(\mathbf{Y}\mathbf{Y}^T)) = \tanh(\text{Relu}(\sum_{i=1}^F \mathbf{y}_i \mathbf{y}_i^T))$, where $\mathbf{y}_i$ corresponds to the i th column in $\mathbf{Y}$ (respectively in Θ). Before the nonlinear activation, we can view $\mathbf{Y}\mathbf{Y}^T$ as a rank- F approximation to the FC graph. This prompts us to extract and analyze each of the rank-one components $\mathbf{y}_i \mathbf{y}_i^T$, and consider the component graphs

$$\mathbf{Z}_i = \tanh(\text{Relu}(\mathbf{y}_i \mathbf{y}_i^T)), \quad i = 1, \dots, F \quad (8)$$

which correspond to the outputs of the individual filters in the graph convolutional layer. These F graphs can be regarded as building blocks of the FC network, that we believe could reveal more details about the process in the generation of FC patterns from SC networks.

3. NUMERICAL TEST CASES

Here we evaluate the effectiveness of the proposed architecture on a real-world neuroimaging dataset.

3.1. Neuroimaging data

A dataset from the Human Connectome Project (HCP) containing 1058 healthy subjects is used in this paper. For each subject, the brain SC network is extracted from the raw diffusion MRI (dMRI) and structural MRI (sMRI) data using the pipeline in [30]. The brain functional activities are measured by the blood oxygen-level dependent (BOLD) signals residing on each brain region using functional MRI (fMRI) scanning. The brain FC networks are then obtained by computing the Pearson correlation between the BOLD signals².

The Desikan-Killiany atlas [31] is used to define brain regions corresponding to the nodes in both the FC and SC networks thus both networks have $N = 68$ cortical surface regions with 34 nodes in




Fig. 2: Correlation between (left) SC and empirical FC; (right) reconstructed FC and empirical FC. Higher correlation between the reconstructed FC originated from SC and empirical FC shows that the model captures the relation between SC and FC.

each hemisphere. Since the FC networks are obtained from the brain signals, in our tests we do not use the averaged BOLD timecourses as node attributes and use one-hot encoding instead to define graph signals. Accordingly, we set $\mathbf{X} = \mathbf{I}_{68}$ in (6).

3.2. Implementation

The GCN-based encoder-decoder model is implemented using Tensorflow. The number of filters F in the graph convolutional layer is set to 32 and the weight coefficients Θ are initialized following [32]. 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 5 consecutive training epochs. The Adam [33] optimizer is used to minimize the MSE with learning rate 0.001.

3.3. Results

Reconstruction performance. The model is trained for 10 folds and the average test error is 0.0304 with a standard deviation of 0.0011. This indicates that the proposed graph encoder-decoder system can reconstruct FC networks from SC networks effectively with a stable performance. Fig. 2 shows the correlation coefficients between the vectorization of the input SC networks and the empirical FC networks, and between vectorization of the reconstructed FC networks and the empirical FC networks for all 1058 subjects. It clearly shows that the graph encoder-decoder system is able to generate estimated FC networks highly associated to the empirical FC networks even when the SC networks and FC networks do not correspond well.

In Fig. 3, we further illustrate the ground-truth and reconstructed FC networks of two randomly drawn test samples (test error = 0.021 and 0.0306, respectively) which depicts similarities between the reconstructed FC network and the empirical FC network for each subject. Note that since the model is trained to capture the SC-FC relation in a very large cohort of subjects, the latent variables used to reconstruct the FC networks capture population patterns. Thus the model should offer satisfactory performance on most subjects, that is we expect it will generalize well as the test MSE corroborates. As the goal is to study the overall relation between SC and FC networks and to find out general patterns within SC-to-FC mapping, we believe that our model is well-suited to this end.

The reconstructed FC networks are learned by considering the topology of SC networks and the interactions between brain regions beyond direct anatomical pathways. The apparent similarity between the reconstructed and the empirical FC networks as shown in Fig. 3 suggests that the embeddings $\mathbf{Y}$ have learned and captured valuable information related to the unknown SC-FC relationship, as will be further elaborated in the next section.

Network decomposition. As mentioned in Section 2.2, the reconstructed FC networks can be further decomposed into multiple graphs. Each of the graphs is built from one column in the output of

²For more information about the data and preprocessing steps, refer to [30] and <https://www.humanconnectome.org/>




Fig. 3: (a)(d) Structural network (b)(e) Recovered FC network (c)(f) Empirical FC network for two subjects (top and bottom row, respectively). While the SC shows the pattern of highly within-hemisphere connections and sparser between-hemisphere connections, both reconstructed FC and empirical FC show less structured (i.e., modular) pattern of connections.




Fig. 4: The sub-networks identified through decomposition of reconstructed FC network. Each subplot corresponds to a brain sub-network that contributes to the SC-to-FC mapping.

the encoder, i.e. the output of each filter in the graph convolutional layer. As the superposition of these rank-one outer products (before the nonlinear activation) gives rise to the reconstructed graph structure at the decoder output, we can think of these graphs as building blocks that contain vital information on a finer scale regarding the mapping from SC to FC; see (8).

To evaluate what the filters learn during the training of the graph encoder-decoder system, we choose one model with the smallest test error among 10 folds. The SC networks of all 1058 subjects are fed to the system and the reconstructed FC network is generated for each subject that best approximates the individual empirical FC network. Since the number of filters is set to $F = 32$, each reconstructed FC network is decomposed into 32 graphs. For each filter, a representative graph is obtained by taking the average of the outputs across all 1058 subjects. Among the 32 representative graphs, some stand out to correspond well to known brain regions and brain sub-networks. We believe that these regions and sub-networks contribute to the mapping from SC to FC and may provide vital information regarding the formation of brain functional activity. Fig. 4 presents the connectivity matrices of the identified sub-networks which are visualized in brain domain in Fig. 5.




Fig. 5: The sub-networks in Fig. 4 visualized via BrainNet Viewer [34].

The identified brain sub-networks present nice correspondence with brain functions. For example, Fig. 4d and Fig. 5d present a brain sub-network centering three regions with high intra-connections: parsopercularis (POP), parsorbitalis (POB) and pars-triangularis (PT). These three regions consist the Inferior Frontal Gyrus which is critically involved in a wide range of complex operations [35]. Fig. 4b and Fig. 5b also present another sub-network starting with precentral (PRC), paracentral (PARA), postcentral (POC) and superior parietal (SP) regions. The first two deal with motor and sensory functions and are located in the frontal lobe. The last two are in the parietal lobe and are involved with spatial orientation and somatosensory functions. The results indicate that multiple identified building blocks highly correspond to functional brain sub-networks and the decomposition process serves as functional segregation to dissect the brain FC network into individual components of functional activities.

4. CONCLUSIONS AND FUTURE WORK

The brain SC network and FC network are closely related and how the brain functional activities are generated from and beyond the underlying anatomical pathways has been an intriguing research topic. In this paper, a graph encoder-decoder system was proposed to reconstruct FC graphs using the information from brain SC networks. Through training with a large population of subjects, the model is able to learn population patterns regarding the SC-FC relationship in human brains. The latent variables within the encoder-decoder model can be extracted and interpreted as building blocks that approximate the brain FC networks. Multiple building blocks correspond to key brain sub-networks contributing to the generation of functional activities from anatomical connections. Such findings provide a novel perspective to exploring the SC-to-FC mapping and locating key regions and sub-networks that generate complex brain functional activities. Currently the model is trained on healthy subjects using simple one-hot encoding as node attributes. Future work shall be done to include alternative node features, and train the model to identify differences in the SC-FC relationship between patient and control groups and to locate potential regions as bio-markers for disease diagnosis.

5. REFERENCES

- [1] A. Fornito, A. Zalesky, and E. Bullmore, *Fundamentals of Brain Network Analysis*, Academic Press, 2016.
- [2] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, V. J. Wedeen, and O. Sporns, "Mapping the structural core of human cerebral cortex," *PLoS Biol.*, vol. 6, no. 7, pp. e159, 2008.
- [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. J. Honey, J.-P. Thivierge, and O. Sporns, "Can structure predict function in the human brain?," *Neuroimage*, vol. 52, no. 3, pp. 766–776, 2010.
- [6] M. D. Greicius, K. Supekar, V. Menon, and R. F. Dougherty, "Resting-state functional connectivity reflects structural connectivity in the default mode network," *Cerebral Cortex*, vol. 19, no. 1, pp. 72–78, 2009.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436, 2015.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Adv. Neural. Inf. Process. Syst.*, pp. 1097–1105, 2012.
- [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] 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.
- [15] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, "Spectral networks and locally connected networks on graphs," *Proc. ICLR*, 2014.
- [16] 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.
- [17] T. N. Kipf and M. Welling, "Variational graph auto-encoders," *NIPS Workshop on Bayesian Deep Learning*, 2016.
- [18] H. Cai, V. W. Zheng, and K. C.-C. Chang, "A comprehensive survey of graph embedding: Problems, techniques, and applications," *IEEE Trans Knowl Data Eng.*, vol. 30, no. 9, pp. 1616–1637, 2018.
- [19] W. L. Hamilton, R. Ying, and J. Leskovec, "Representation learning on graphs: Methods and applications," *IEEE Data Engineering Bulletin*, 2017.
- [20] G. Rosenthal, F. Váša, A. Griffa, P. Hagmann, E. Amico, J. Goñi, G. Avidan, and O. Sporns, "Mapping higher-order relations between brain structure and function with embedded vector representations of connectomes," *Nat. Commun.*, vol. 9, no. 1, pp. 2178, 2018.
- [21] 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.
- [22] P. Lee, M. Choi, D. Kim, S. Lee, H.-G. Jeong, and C. E. Han, "Deep learning based decomposition of brain networks," *Intl. Conf. on Artificial Intelligence in Information and Communication (ICAIIIC)*, pp. 349–354, 2019.
- [23] A. Ortega, P. Frossard, J. Kovačević, J. M. Moura, and P. Vandergheynst, "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, D. S. Bassett, A. Ribeiro, and D. Van De Ville, "A graph signal processing perspective on functional brain imaging," *Proc. IEEE*, vol. 106, no. 5, 2018.
- [25] 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.
- [26] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *ICLR*, 2017.
- [27] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," *Intl. Conf. on Machine Learning (ICML)*, pp. 807–814, 2010.
- [28] 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.
- [29] M. Rubinov and O. Sporns, "Complex network measures of brain connectivity: uses and interpretations," *Neuroimage*, vol. 52, no. 3, pp. 1059–1069, 2010.
- [30] 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.
- [31] R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson, 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.
- [32] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," *Intl. Conf. on Artificial Intelligence and Statistics*, pp. 249–256, 2010.
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 2015.
- [34] M. Xia, J. Wang, and Y. He, "Brainnet viewer: a network visualization tool for human brain connectomics," *PloS one*, vol. 8, no. 7, pp. e68910, 2013.
- [35] J. D. Greenlee, H. Oya, H. Kawasaki, I. O. Volkov, M. A. Severson III, M. A. Howard III, and J. F. Brugge, "Functional connections within the human inferior frontal gyrus," *J. Comp. Neurol.*, vol. 503, no. 4, pp. 550–559, 2007.