

Learning to Model the Relationship Between Brain Structural and Functional Connectomes

Yang Li , Member, IEEE, Gonzalo Mateos , Senior Member, IEEE, and Zhengwu Zhang

Abstract—Recent neuroimaging advances along with algorithmic innovations in statistical learning from network data offer a unique pathway to integrate brain structure and function, and thus facilitate revealing some of the brain’s organizing principles at the system level. In this direction, we develop a supervised graph representation learning framework to model the relationship between brain structural connectivity (SC) and functional connectivity (FC) via a graph encoder-decoder system, where the SC is used as input to predict empirical FC. A trainable graph convolutional encoder captures direct and indirect interactions between brain regions-of-interest that mimic actual neural communications, as well as to integrate information from both the structural network topology and nodal (i.e., region-specific) attributes. The encoder learns node-level SC embeddings which are combined to generate (whole brain) graph-level representations for reconstructing empirical FC networks. The proposed end-to-end model utilizes a multi-objective loss function to jointly reconstruct FC networks and learn discriminative graph representations of the SC-to-FC mapping for downstream subject (i.e., graph-level) classification. Comprehensive experiments demonstrate that the learnt representations of said relationship capture valuable information from the intrinsic properties of the subject’s brain networks and lead to improved accuracy in classifying a large population of heavy drinkers and non-drinkers from the Human Connectome Project. Our work offers new insights on the relationship between brain networks that support the promising prospect of using graph representation learning to discover more about brain function.

Index Terms—Brain connectomes, encoder-decoder system, graph classification, graph convolutional network, graph representation learning, graph signal processing.

I. INTRODUCTION

THE human brain is a complex yet efficient information processing network, whose distributed organization entails different regions conducting individual tasks while actively interacting with each other [2]. This *integrative* nature of brain function along with recent advances in neuroimaging, motivate well the adoption of graph-centric signal and information

Manuscript received 18 December 2021; revised 24 June 2022 and 20 September 2022; accepted 20 September 2022. Date of publication 23 September 2022; date of current version 3 October 2022. This work was supported in part by NSF Awards under Grants CCF-1750428, CCF-1934962, and ECCS-1809356. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Siheng Chen. (Corresponding author: Gonzalo Mateos.)

Yang Li and Gonzalo Mateos are with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627 USA (e-mail: yangli@roc.edu; gmateos@ece.rochester.edu).

Zhengwu Zhang is with the Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599 USA (e-mail: zz10c@email.unc.edu).

Digital Object Identifier 10.1109/TSIPN.2022.3209097

processing tools to study the interplay between brain structure (of neural connections often referred to as the structural connectome) [3] and functional activity [4]. Indeed, understanding the underpinnings of brain function is one of the fundamental scientific challenges of this century.

Brain connectivity broadly consists of two types of networks. *Structural connectivity* (SC) has to do with anatomical tracts of axonal bundles [5], which can be extracted using tractography algorithms applied to diffusion magnetic resonance imaging (dMRI). On the other hand, *functional connectivity* (FC) represents pairwise statistical correlation between activation signals in various brain regions of interest (RoIs) [4], [5]. These so-termed blood oxygen level-dependent (BOLD) signals are measured by functional MRI (fMRI). The interplay between SC and FC is of great importance and a timely area of research in network neuroscience. There is consensus that further understanding of such coupling could offer new and important insights on the inner working mechanics of the brain [6], [7], [8]. Previous studies have revealed that FC correlates with SC at an aggregate level [9]. Although FC is shaped by the actual anatomical connections within the brain, strong functional connections exist between brain regions that are not directly linked by structural connections, or have rather limited anatomical pathways [10]. These findings provide strong evidence that functional interactions between brain regions depend on both direct and indirect anatomical connections [11], and motivate multiple previous works in predicting FC from SC [9], [10], [12], [13], and in estimating SC from FC [8]. Recently, [14] developed a deep learning model to predict FC from SC, and reported that the structure-function coupling in human brain networks is substantially tighter than previously suggested.

In this paper, we model the structure-function relationship in the brain by learning low-dimensional representations of the SC-to-FC mapping using graph convolutional networks (GCNs) [15], [16], [17]. This way, we aim at better exploiting the relational inductive biases present in brain connectome data. While the benefits of graph representation learning (GRL) methods have been well documented in applications ranging from recommender systems [18] to social network analysis [19], their impact to network neuroscience is yet to fully materialize [20], [21]. In the context of brain network analysis, it is not apparent how to generalize inductive GRL models with learnable parameters shared among a large cohort of subjects; we explore this avenue here. Specifically, we propose and test a supervised GRL framework to jointly learn: (i) low-dimensional *node* embeddings generated from SC networks to reconstruct

empirical FC; and (ii) *graph* embeddings to represent the whole graphs for subject-level classification.

Our overall architecture can be viewed as a graph encoder-decoder system inspired by the graph autoencoder model [22]; see also [23]. For each ROI in the input SC graph, the encoder outputs low-dimensional node embeddings that integrate both nodal attributes (when available) and the local graph topology information. To aggregate information from multiple hops within the network, layered graph convolution operations interleaved with point-wise nonlinearities are used to compute nodal features. This makes the GCN a proper encoder model to generate connectome embeddings that capture indirect connections inherent to brain networks, and the SC-FC relationship in particular. Whole graph embeddings are then obtained by applying certain global pooling mechanisms on the node embeddings [24]. All in all, we propose a multi-task learning model with the dual goal of reconstructing brain FC from input SC data, and performing subject classification using the learnt graph embeddings as features. The graph estimation/regression task learns a parsimonious representation of the population-level SC-FC relationship, and the classification task uses subject labels as additional inputs for supervised learning. Accordingly, the model strikes the right balance between deciphering population patterns that shape the SC-FC coupling, and distilling subject-level variability to facilitate graph classification with regard to drinking habits.

We train and test our graph encoder-decoder model on a neuroimaging dataset of 412 subjects from the Human Connectome Project (HCP) [25], and obtain satisfactory FC reconstruction performance as well as subject classification accuracy (classes are 191 non-drinkers and 221 heavy drinkers). The proposed GRL pipeline outperforms baseline auto-encoder methods that only rely on SC or FC, supporting the thesis that SC-to-FC mapping provides valuable information to better discriminate among the classes beyond static graph topology, specifically in the study of alcohol's effect on the human brain. Via statistical tests on the reconstructed FC networks, brain sub-graphs exhibiting significant difference between groups are also unveiled. All in all, the novelty of our work is in learning low-dimensional, latent representations that capture the intrinsic attributes of the brain SC-FC coupling across a large population, while being discriminative for subject-level classification. This paper is not about fundamental innovations in the GRL space, but instead it offers a novel application of this framework to neuroimaging data analysis with the distinct perspective of modeling the SC-FC relationship in the brain.

A. Related Works

GNNs are versatile signal and information processing architectures [15], [16], [17], which comprise stacked layers of graph (convolutional) filters followed by point-wise nonlinearities; see [17], [23], [26], [27], [28], [29] for recent surveys and the references therein. From early spectral convolutions [15], [30] to distributed implementations of (equivalent) shift-invariant polynomial graph filters [19], [26], [31], GNNs integrate information from both the graph topology and nodal attributes to learn representations of network data. Indeed, the GRL paradigm

is to learn low-dimensional embeddings of individual vertices, edges, or the graph itself [23], [32], [33], [34], which can then be used in e.g., (semi-supervised) node classification [15], link prediction [35], graph clustering [36], [37], and graph classification [38]. Recently, GRL ideas have permeated to neuroimaging data analysis for behavioral state classification [39], to study the relationship between SC and FC [13], [40], and to extract representations for subject classification [41], [42], [43]. Recent works in [44] and [45] utilize spatio-temporal-based GNNs to capture the temporal information within the brain signals, and demonstrate superior performance in fMRI classification. Although these prior works have probed the area of brain connectomics in several novel directions and achieved solid performance in multiple regression or classification tasks, they mostly rely on a single type of brain network and use feed-forward models to predict subject labels from input graphs. Our different approach is to distill actionable information from the relationship between FC and SC. Instead of learning representations of a certain type of brain network, we propose a multi-task GRL pipeline to model the SC-FC coupling, thus seamlessly integrating brain structure and function in a principled way.

B. Summary of Contributions

In summary, we develop a GNN-based supervised encoder-decoder system to learn parsimonious latent-space representations of the SC-FC relationship from HCP neuroimaging data. Different from the surveyed deep learning approaches for brain network analyses, to the best of our knowledge this is the first work to explore learnt representations of the mapping between two brain connectivity modalities – among which a tight coupling is well documented [9], [12], [40]. Our model can extract valuable and interpretable information from the intrinsic properties of the subjects' brain networks, leading to an improved performance trade-off between FC reconstruction error and graph classification accuracy relative to baseline methods. We also identified brain sub-networks exhibiting significant functional connectivity differences among heavy drinkers and non-drinkers. Such data-driven findings may serve as explanations behind the improved classification performance. They could also offer additional insights to guide follow-up studies about alcohol's effects on the brain.

Relative to the conference precursor [1], here we consider GRL-based modeling of the SC-FC mapping through a unified presentation along with full-blown technical details (including required background material, model refinements, extended discussions and interpretation of results, as well as unpublished experimental tests). Noteworthy novel pieces include:

- **Exploration of GNN architectures.** We examine different GNN-based encoder architectures with multiple layers and various numbers of filters per layer. Accordingly, the representation space of the resulting model is markedly richer than the *shallow GNN encoder* in [1]. Inspired by [31], we also find that concatenating the outputs from each GNN layer as a readout function improves model performance.

As an important component in our architectural design not present in [1], we explore various global (mean/max/sum) pooling mechanisms to form graph-level representations from the set of node embeddings.

- **Comparison with baseline methods.** Preliminary tests in [1] did not include comparisons with single-modality FC or SC auto-encoder baselines, while the results presented here allow us to justify our claim on the merit of learning SC-FC representations (that inherently integrate brain structure and function). New comparisons with a graph isomorphism network (GIN)-based classifier [31] are included in the Appendix.
- **Reproducibility.** The code we developed to train and optimize the model architecture in publicly accessible from <https://github.com/yli131/brainGRL/>. The HCP data we used is open access but cannot be shared in our repository, please refer to the HCP's data usage terms.

II. GRAPH NEURAL NETWORKS BACKGROUND

In this section, the required background of graph neural networks (GNNs) is briefly outlined. As the GCN [15] is adopted for the encoder in the GRL pipeline, we focus on graph convolutional models; see also [17], [23], [27], [28], [29] for comprehensive surveys with further details.

A. Network Models, Graph Signals and Transforms

Consider a weighted, undirected graph denoted by $G := (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \dots, N\}$ is a set of N vertices corresponding to brain ROIs and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ are edges in which $(i, j) \in \mathcal{E}$ is a structural connection joining i and j . We define the neighborhood of vertex i as $\mathcal{N}(i) = \{j \in \mathcal{V} : (i, j) \in \mathcal{E}\}$. The symmetric adjacency matrix $\mathbf{A} \in \mathbb{R}_+^{N \times N}$ has entries $A_{ij} \geq 0$ representing the structural connection strengths between ROIs i and j ; $A_{ij} = 0$ indicates $(i, j) \notin \mathcal{E}$. Henceforth we will use Σ to denote the $N \times N$ adjacency matrix of the FC network, making it explicit that this is a correlation matrix (and different from \mathbf{A}). Network data models often augment G with a vertex-valued signal $\mathbf{x} \in \mathbb{R}^N$, where x_i denotes the signal value at node i , for example, the nodal attributes (or features) on the brain FC (or SC) network. Extensions to feature vectors (of length d) $\mathbf{x}_i \in \mathbb{R}^d$ per node are straightforward, leading to a signal matrix $\mathbf{X} \in \mathbb{R}^{N \times d}$ that serves as input to the GNN.

The main building block of spectral graph theory is the graph Laplacian matrix defined as $\mathbf{L} := \mathbf{D} - \mathbf{A}$, where $\mathbf{D} := \text{diag}(\mathbf{A}\mathbf{1})$ is the diagonal matrix of nodal degrees. The Laplacian \mathbf{L} is symmetric and positive semidefinite. Accordingly, it can be decomposed as $\mathbf{L} = \mathbf{U}\Lambda\mathbf{U}^\top$, where $\mathbf{U} \in \mathbb{R}^{N \times N}$ denotes the set of orthonormal eigenvectors and the diagonal matrix Λ contains all the non-negative eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N := \lambda_{\max}$. The eigenvectors in \mathbf{U} serve as Fourier modes to decompose graph signals [46]. Specifically, the graph Fourier transform (GFT) of \mathbf{x} is defined as $\tilde{\mathbf{x}} := \mathbf{U}^\top \mathbf{x}$, where $\tilde{\mathbf{x}} = [\tilde{x}_1, \dots, \tilde{x}_N]^\top$ collects the graph spectral coefficients of \mathbf{x} at frequencies given by the Laplacian eigenvalues [46], [47].

B. Graph Convolutional Networks

To process signals \mathbf{x} efficiently while incorporating G 's topology information, one defines the graph convolution as

$$\mathbf{H}\mathbf{x} = \sum_{k=0}^K h_k \mathbf{L}^k \mathbf{x}, \quad (1)$$

where $\mathbf{H} := \sum_{k=0}^K h_k \mathbf{L}^k$ is a graph filter with coefficients $\mathbf{h} := [h_0, \dots, h_K]^\top$. The graph convolution operation aggregates signal values from K -hop neighborhoods in G . Indeed, computation of $\mathbf{L}^K \mathbf{x} = \mathbf{L}(\mathbf{L}^{K-1} \mathbf{x})$ entails a sequence of K one-hop aggregations (shifts or diffusions) through multiplication with \mathbf{L} . Graph filters can be more generally defined as polynomials of GSOs beyond the combinatorial Laplacian \mathbf{L} , for instance the adjacency matrix \mathbf{A} or suitable normalized variants of the aforementioned algebraic graph topology descriptors [26], [46]. With this interpretation, the analogy of (1) with temporal convolutions implemented via (shift and sum) finite impulse response (FIR) filters should be apparent. Graph convolutions are thus linear shift-invariant operators because $\mathbf{H}\mathbf{S} = \mathbf{S}\mathbf{H}$, where $\mathbf{S} = \mathbf{L}$ in (1). The intuitiveness of the definition can be further appreciated in the graph frequency domain, since the GFT of the filtered signal (1) becomes

$$\mathbf{U}^\top \mathbf{H}\mathbf{x} = \mathbf{U}^\top \sum_{k=0}^K h_k (\mathbf{U}\Lambda\mathbf{U}^\top)^k \mathbf{x} = \left(\sum_{k=0}^K h_k \Lambda^k \right) \tilde{\mathbf{x}}. \quad (2)$$

Notice how the GFT diagonalizes the filter, so that the convolution becomes the element-wise (per frequency λ_i) multiplication between the filter's frequency response $\tilde{h}_i := \sum_{k=0}^K h_k \lambda_i^k$ and the signal's GFT coefficient \tilde{x}_i .

GCNs broadly consist of stacked layers of learnable graph convolutional filters and point-wise nonlinearities. Convolutional NNs for network data can be traced to [30], where the filter's frequency response $\tilde{\mathbf{h}}^{(\ell)} = [\tilde{h}_1^{(\ell)}, \dots, \tilde{h}_N^{(\ell)}]^\top$ at each layer ℓ is learnt using stochastic gradient descent. While intuitive this architecture has several drawbacks. First, computing the eigen-decomposition of \mathbf{L} may become computationally infeasible for large graphs, and the number of trainable parameters grows with N . Also, the filters depend on the eigenbasis of the Laplacian and thus the parameters cannot be shared across different graphs, which limits its usage in an inductive setting. Finally, without a smoothness constraint in the frequency response one obtains filters that are not localized in the vertex domain [cf. the smooth polynomial frequency response in (2), resulting in a K -hop localized graph filter (1)]. To overcome these limitations, the ChebNet was proposed in [48] by defining a filter in terms of Chebyshev polynomials of the diagonal matrix of eigenvalues Λ . In terms of expressive power, this choice is essentially equivalent to the polynomial graph filters (1); see also [17], [26].

Adoption of first-order ($K = 1$) graph convolutional filters was advocated for the GCN model in [15], working with the degree-normalized Laplacian $\mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$ and letting $\theta = h_0/2 = -h_1$ so that (1) simplifies to

$$\mathbf{H}\mathbf{x} = \theta(\mathbf{I}_N + \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}) \mathbf{x}. \quad (3)$$

This motivates a simple per-layer filtering update implemented to refine the nodal embeddings, namely

$$\mathbf{X}^{(\ell)} = \text{ReLU} \left(\tilde{\mathbf{A}} \mathbf{X}^{(\ell-1)} \Theta^{(\ell)} \right), \quad (4)$$

where $\tilde{\mathbf{A}} := \mathbf{I}_N + \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ is a graph shift operator (GSO), $\mathbf{X}^{(\ell)} \in \mathbb{R}^{N \times d_\ell}$ are the nodal representations at layer ℓ , $\Theta^{(\ell)} \in \mathbb{R}^{d_{\ell-1} \times d_\ell}$ stores the learnable parameters of d_ℓ filters acting on $d_{\ell-1}$ input features, and $\text{ReLU}(x) = \max(0, x)$ is a point-wise nonlinear activation function. Notice that actually [15] advocates a slightly different GSO than $\tilde{\mathbf{A}}$, which is obtained from \mathbf{A} using a renormalization trick that we describe in Section III.

In the sequel, for the proposed encoder architecture we leverage GCN layers in (4) due to their simplicity and satisfactory performance on weighted graphs [15], [31], [49], such as brain SC. While using only first-order filters, GCNs have the capability to integrate multi-hop information within G using $L > 1$ stacked layers. Adoption of more expressive filters where $K > 1$ is left as future work, and may be prudent for larger neuroimaging datasets than the one described next.

III. PROBLEM SETUP AND PROPOSED METHOD

Given brain SC graphs, we build and train a model with the twofold goal of: (i) reconstructing the corresponding FC networks from nodal embeddings; and (ii) classifying a cohort of subjects using graph-level representations learnt from binary labels in a supervised fashion. This way, learnt representations should capture both population patterns and subject-level variability. We start by describing the dataset in Section III-A. The proposed supervised graph encoder-decoder model is then presented in Section III-B.

A. Data Description and Network Construction

We adopt a neuroimaging dataset with $P = 412$ subjects from the Human Connectome Project (HCP) [50], [51]. The cohort is partitioned into two classes: 191 non-drinkers and 221 heavy drinkers according to information available on the lifetime maximum number of drinks had in a single day. Subjects with no more than 3 drinks in a single day are considered non-drinkers, and heavy drinkers are those who had over 16 drinks. Refer to Appendix A for more details on the subject label delineation. The Desikan-Killiany atlas is used to specify brain RoIs [52]. Hence, \mathcal{V} in both FC and SC networks correspond to $N = 68$ cortical surface regions, with 34 nodes in each hemisphere. Based on the data-processing pipeline in [53], [54], the SC network \mathbf{A} of each subject is extracted from the dMRI and structural MRI data. Brain functional activities on each RoI are given by the resting-state BOLD time courses measured using fMRI [55]. Brain FC networks Σ are then constructed so that edge weights are the Pearson correlation coefficient between the BOLD signals at the incident RoIs. We find that the resulting FC graphs contain few negative edge weights with much smaller magnitude compared to a vast majority of positive edges. To mitigate such data imbalance problem, we discard all functional edges with negative weights and restrict




Fig. 1. Visualization of SC and FC networks of a sample subject, where RoIs have been ordered according to the hemisphere they belong to. Note the increased edge density in the main diagonal blocks of the SC adjacency matrix, corresponding to a modular structure with more connections within hemispheres than between. The FC graph exhibits a more disassortative pattern of connections among brain RoIs. Increased functional correlations between regions that belong to different hemispheres are observed.

ourselves to entries $\Sigma_{ij} \in [0, 1]$, as it is customarily done in prior FC studies [56], [57]. For additional details about the data, preprocessing, and network construction steps, refer to [53], [54] and <http://www.humanconnectome.org/>.

In Fig. 1, we depict the resulting adjacency matrices of the SC and FC networks of a sample subject in the dataset. Nodes are rearranged by hemisphere, with the first 34 vertices corresponding to RoIs that belong to the left brain. Note how edges densely concentrate in the main diagonal blocks of the SC network, thus indicating a markedly larger amount of anatomical connections within hemispheres rather than across the left and right brains. On the other hand, the FC network exhibits a more disassortative (or less modular) connectivity pattern. This supports the well-documented finding that functional links exist between RoIs that are loosely-connected in a structural sense, thus also depending on indirect (multi-hop) anatomical pathways. These general observations are consistent across subjects.

B. Graph Encoder-Decoder Model and Architecture

Here we introduce the proposed graph encoder-decoder system used to model the relationship between structural and functional brain connectomes; see also the schematic description of the architecture in Fig. 2 and the recent survey in [23]. The top branch in Fig. 2 implements a regression model to reconstruct FC from SC. The bottom branch is a supervised graph classification algorithm to predict heavy-drinking subjects. Notice that the encoder module is shared and trained using HCP data to accomplish both tasks. This way, the learnt representations capture population patterns in the SC-FC relationship while also reflecting subject-specific variability to discriminate among classes.

Next, we introduce the main modules of the architecture and explain various design choices made.

Input. The required input for the GRL model is an SC network represented by its (symmetric) adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, where $N = 68$ is the number of RoIs from the Desikan atlas. Optionally, nodal attributes or graph signals $\mathbf{X} \in \mathbb{R}^{N \times d_0}$ can be fed to the model to provide additional information via RoI-specific feature vectors of length d_0 . Since we did not extract




Fig. 2. The GRL encoder-decoder architecture. The inputs are SC networks \mathbf{A} and optional nodal attributes $\mathbf{X}^{(0)}$. Graph convolution/information propagation occurs within the deep GCN encoder. Rows in $\mathbf{X}^{(L)}$ are low-dimensional node embeddings at the final layer of the GCN encoder. An outer-product decoder reconstructs the FC network $\hat{\Sigma}$ and thus implicitly models the SC-to-FC mapping. Given graph (i.e., subject)-level representations \mathbf{x}_G obtained via pooling, a logistic regression classifier outputs the predicted binary label \hat{y} . Training is performed over a large cohort of HCP subjects, using empirical FC Σ and labels y to define a judicious loss function.

RoI attributes from the dataset outlined in Section III-A, we use one-hot encoding to define graph signals and accordingly set the initial signal input as $\mathbf{X}^{(0)} = \mathbf{I}_{68}$ ($d_0 = 68$). Future work will be devoted to collect and investigate the merits of incorporating additional meaningful subject-related nodal attributes, in addition to the graph structure itself. Examples of potential graph signals include RoI volume, node degree in the SC network, or BOLD timecourses. The latter two alternatives have been considered in network-based machine learning (ML) studies of neuroimaging data [10], [43], [47].

For training, empirical functional networks Σ and subject (i.e., graph-level) labels y related to drinking are used to define a loss function we describe later in this section. Accordingly, both these quantities could be viewed as inputs to the model during the training phase.

Encoder. The deep encoder module is a parametric function

$$\mathbf{X}_C = \text{ENC}(\tilde{\mathbf{A}}, \mathbf{X}^{(0)}; \Theta_E) \quad (5)$$

that takes a normalized version of the SC adjacency matrix $\tilde{\mathbf{A}}$ as well as graph signals $\mathbf{X}^{(0)}$ as inputs, and generates low-dimensional representations for each node that we stack as rows of \mathbf{X}_C . Θ_E represents the trainable parameters in the encoder. Following the renormalization idea in [15], the normalized adjacency matrix is $\tilde{\mathbf{A}} := \hat{\mathbf{D}}^{-1/2} \hat{\mathbf{A}} \hat{\mathbf{D}}^{-1/2}$, where $\hat{\mathbf{A}} := \mathbf{I}_N + \mathbf{A}$ and $\hat{\mathbf{D}} = \text{diag}(\hat{\mathbf{A}}\mathbf{1})$ is the degree matrix of $\hat{\mathbf{A}}$. Among various possible node embedding approaches, neighborhood-aggregation methods such as the GCN model discussed in Section II-B are simple, permutation-invariant, and inductive [58]. As a result, we opted for a GCN-based encoder to generate latent variables capturing network topology information such as the connection strengths among RoIs in SC networks.

Unlike the single-layer GCN model in our previous work [1], here we consider a deep GCN encoder with $L \geq 1$ stacked layers and multiple filters d_ℓ per layer ℓ as listed in Table I. The methodology used to choose the values of the hyperparameters L and $\{d_\ell\}$ is outlined in Section IV-A. Another alternative to gain in expressivity is to utilize graph convolutional filters (1) with $K > 1$. But given the limited number of training samples and the satisfactory performance obtained with first-order filters, we decided to stick with a GCN architecture. Recall from (4) that the weight matrix $\Theta^{(\ell)} \in \mathbb{R}^{d_{\ell-1} \times d_\ell}$ contains learnable graph

TABLE I
GCN ARCHITECTURES FOR HYPERPARAMETER SEARCH

$L = 1$ layer	$d_0 \times \{128; 64; 32; 16; 8\}$
$L = 2$ layers	$d_0 \times \{128 \times 64; 64 \times 32; 32 \times 16; 16 \times 8\}$
$L = 3$ layers	$d_0 \times \{128 \times 64 \times 32; 64 \times 32 \times 16; 32 \times 16 \times 8\}$

filter coefficients at layer ℓ . So overall, the trainable parameters in the deep GCN-based encoder (5) are $\Theta_E := \{\Theta^{(\ell)}\}_{\ell=1}^L$. The ReLU activation function is used to speed up training and avoid the problem of vanishing gradients in multi-layer settings [59].

It is natural to consider the output $\mathbf{X}^{(L)} \in \mathbb{R}^{N \times d_L}$ of layer L as the sought nodal embeddings. Inspired by [31], [58], we also evaluated the contribution from node embeddings learnt in intermediate GCN layers $\ell = 1, \dots, L-1$ via concatenation, thus forming $\mathbf{X}_C \in \mathbb{R}^{N \times \sum_\ell d_\ell}$.

Decoder and loss function. Starting with the FC reconstruction branch of the GRL architecture in Fig. 2, node embeddings \mathbf{X}_C go through an outer product decoder

$$\hat{\Sigma} = \text{ReLU}(\mathbf{X}_C \mathbf{X}_C^T) \quad (6)$$

to generate a predicted FC adjacency matrix $\hat{\Sigma}$. Notice that we utilize a rectifier to enforce the constraint $\hat{\Sigma}_{ij} \geq 0$; recall we have dropped those few functional connections exhibiting negative correlations (Section III-A). To also enforce $\hat{\Sigma}_{ij} \leq 1$ we have tried other non-linearities such as the sigmoid function or the hyperbolic tangent (useful even when retaining the negative correlations), and found training to be slower while the reconstruction performance is comparable to the chosen ReLU [1]. The mean squared error (MSE) between the reconstructed graph $\hat{\Sigma}$ and the empirical FC Σ , averaged over a training set and denoted as $\mathcal{L}_{\text{MSE}}(\hat{\Sigma}, \Sigma)$, is used as reconstruction loss for training the regression branch of the model. All in all, while we did not normalize the decoder output to impose $\hat{\Sigma}_{ij} \leq 1$ we instead let the MSE reconstruction loss adjust the scale of the predicted Pearson correlations.

Moving on to the supervised classification branch of the model, we compute a *graph-level embedding* by taking the row-wise average of all nodal representations in \mathbf{X}_C . This results

in a vector \mathbf{x}_G that summarizes information about the SC-FC relationship. Admittedly simple, average pooling has been shown effective in multiple studies [16], [60]. For a thorough comparison, we also investigated max pooling and sum pooling which are common alternatives [19]. Attention mechanisms could also be applied to determine the relevance of each ROI to the subject-level task [61]. The graph embeddings \mathbf{x}_G are then fed to a logistic regression classifier to predict the binary subject labels. As our main objective is to assess how well the graph convolution operation helps integrate information within brain networks, a simple logistic regression classifier is preferred rather than advanced models such as multi-layer perceptrons, where training may be biased towards the fully connected layers post GCN.

Finally, the GRL model is trained on HCP data in an end-to-end manner by considering the regression and classification tasks simultaneously. The overall loss function is given as

$$\mathcal{L} = \mathcal{L}_{\text{MSE}}(\hat{\Sigma}, \Sigma) + \lambda \times \mathcal{L}_{\text{CLA}}(\hat{y}, y), \quad (7)$$

where \hat{y} is the predicted label and y is the ground truth label of the subject. The sigmoid cross-entropy loss denoted by \mathcal{L}_{CLA} is used to evaluate classification performance. The hyperparameter λ controls the trade-off between FC reconstruction and classification performance. Its value is determined via grid search and predefined criteria described in the next section. By training the model end-to-end with the loss function as in (7), we aim to strike the right balance between FC reconstruction and subject classification and achieve satisfactory performance on both objectives. The results are presented in Section V.

IV. TRAINING AND EVALUATION FRAMEWORK

Training details and hyperparameter search criteria are outlined in Section IV-A; the code to implement the model is publicly accessible from <https://github.com/yli131/brainGRL/>. We also describe our comprehensive evaluation protocol in Section IV-B, listing several baselines that will guide the presentation of numerical results in Section V.

A. Training Details

The supervised graph encoder-decoder system is implemented in TensorFlow [62]. The details about the model architectures we tested are given in Table I. We carry out a two-stage training procedure for model selection.

The first stage is a low-resolution model selection module. We set $\lambda = 0.1$ following our preliminary results in [1], and search for the optimal architecture among the options in Table I. With fixed λ , during training of each candidate model architecture we use 10-fold cross validation whereby the whole dataset of $P = 412$ subjects is randomly partitioned into 80% training, 10% validation, and 10% test set. The random seed for each fold is fixed for a fair comparison between model architectures. We use Xavier initialization for the weight coefficients $\Theta^{(\ell)}$ of each GCN layer ℓ [63]. The Adam algorithm with learning rate 0.001 was adopted to optimize the tunable parameters in Θ_E as well as those of the logistic regression classifier [64]. To avoid overfitting, early stopping is applied to monitor the validation

loss and to stop the training once the loss increases during 10 consecutive training epochs. Altogether, we trained 57 models because we also considered three global pooling mechanisms (mean/max/sum pooling) and node embedding concatenation for the multi-layer GCN encoders (with $L = 2$ or $L = 3$).

For each of the 57 model architectures, metrics including the graph reconstruction error and classification accuracy are monitored across all 10 folds to evaluate the performance on both regression and classification. We also compute the F score in its traditional form –defined as the harmonic mean of the classifier’s precision and recall– in order to better evaluate the mis-classification between groups. Training time is also taken into consideration to offer insights into the performance cost ratio. A selection criterion \mathcal{C}_{low} is developed by combining the aforementioned metrics, which is given by

$$\mathcal{C}_{\text{low}} = \frac{\text{classification accuracy} \times \text{F score}}{\text{MSE} \times \text{training time}}. \quad (8)$$

The three models with highest \mathcal{C}_{low} are singled out for the second stage. Therein, an extensive grid search over λ values from 0 to 1.5 with step size 0.1 is carried out to determine the optimal λ for each selected GCN architecture.

Recall the loss function (7). For large values of λ , the reconstruction performance deteriorates because training is biased towards the subject classification objective. When $\lambda = 0$, the model achieves the best possible FC reconstruction while the classification loss is at its peak. As λ increases from zero, the classification loss first decreases, then fluctuates and saturates at a certain level. This is intuitive as the graph embeddings that serve as input features for classification depend on the learnt node embeddings that are also driven by the reconstruction objective. When the reconstruction is getting worse, as expected the classification performance will also eventually degrade. The λ search results in Section V-A will corroborate such intuition.

In the end, among the 3 picked encoder architectures and all 16 possible λ values, we choose the final model for downstream analysis and visualization as the one that maximizes

$$\mathcal{C}_{\text{high}} = \frac{\text{classification accuracy} \times \text{F score}}{\text{MSE}}. \quad (9)$$

B. Evaluation Protocol and Baselines

To corroborate the effectiveness of the proposed model, in Section V-B we perform a comprehensive comparative study involving four baseline methods.

First, instead of training the graph classification framework in an end-to-end manner, we train the graph reconstruction branch with $\lambda = 0$ [13]. Subsequently, the graph embeddings obtained by row-wise average pooling of node embeddings are fed to a logistic regression classifier that is separately trained for subject-level classification. Two-step algorithms have been observed to yield sub-optimal performance for supervised learning application [65]. Accordingly, we expect that training our model in an end-to-end fashion while jointly accounting for both objectives in (7) will lead to better performance relative to this *two-step* baseline.

Second, as one of the objectives of our framework is to reconstruct FC and distill information from the structure-function relationship as a byproduct, a natural alternative would be to consider using FC as an *additional input* besides SC. To this end, we carry out an experiment where both the SC and FC networks of the same subject are inputs to a feed-forward GCN model, which is trained to predict the subject-level binary labels associated to drinking habits. Such comparison would help reveal whether the SC-FC mapping offers additional information for classification on top of the raw graph topologies for both modalities.

Third, we explore whether the information within SC or FC *individually* suffices for subject classification. With this purpose, the scheme in Fig. 2 is modified so that the desired output graph is the same as the one fed at the input. As a result, with the input being SC or FC, the model boils down to a graph autoencoder. The training details remain the same and consequently, the learnt representations capture information with respect to either SC or FC only. Their discriminative power for subject classification is evaluated and compared against the learnt representations of the SC-FC mapping. In Appendix C, we also investigate the potential of another GNN model, the graph isomorphism network (GIN) [41] in a feed-forward classifier trained on a single brain network modality.

The last baseline method is an ML pipeline that relies on handcrafted features constructed from graph summary statistics extracted from SC, FC, or both. Network measures [5], [57], [66] capture certain properties of the graph, e.g., community structure, connectivity, and graph components or subgraph structure. Here, we extracted scalar graph measures including average path length, global efficiency, clustering coefficient, graph radius/diameter, transitivity as well as graph density. We combined all aforementioned graph-level measures into a feature vector summarizing network structure. For more information regarding the said graph measures, the interested reader is referred to [57], [66]. After constructing such a feature vector for each subject, we implement workhorse ML models such as SVM, logistic regression and tree-based classifiers; see e.g., [67]. Such a comparison with baseline ML-based methods is expected to reveal the power of the proposed GRL model, by capitalizing on task-driven learnt representations (cf. handcrafted features) that exploit the geometric structure of network data.

V. RESULTS

Here we present and discuss the results of our experimental validation with HCP data.

A. Training and Model Selection

As discussed in Section IV-A, the first stage of training is a low-resolution model selection process, where we tested the 12 GCN-based encoder architectures given in Table I. Comparisons were also conducted among global pooling methods, and between the inclusion and exclusion of node embedding concatenation in multi-layer ($L > 1$) settings. Fig. 3 depicts the results of the performance comparison for 3-layer GCN encoders. We can tell that concatenating node embeddings from previous layers

TABLE II
END-TO-END MODEL VERSUS TWO-STEP BASELINE

Model	MSE	Accuracy	F score
End-to-end	0.0398 ± 0.003	0.6610 ± 0.043	0.6962 ± 0.030
Two-step	0.0377 ± 0.003	0.5072 ± 0.033	0.5071 ± 0.033

to form \mathbf{X}_C helps improve the performance in both subject classification (second row) and FC graph reconstruction (first row). This finding is consistent with prior results reported in the literature [19], [58], suggesting that the embeddings in early layers also contain valuable information about localized signal propagation over the graph G . Based on the criteria \mathcal{C}_{low} defined in (8), we singled-out the three best models. These are the models with $L = 3$ GCN layers (see the third row in Table I), together with node embedding concatenation and global average pooling.

Starting with these three models, the second stage was an extensive grid search of the λ parameter in (7) to determine the optimal trade-off between graph reconstruction and subject classification. The classification accuracy and FC reconstruction MSE on the three selected models were evaluated and presented in Fig. 4. It is apparent that $\lambda = 0$ gives the lowest regression MSE as the model training merely focuses on graph reconstruction. As λ gets larger, the regression error keeps increasing while the classification accuracy slightly fluctuated around a saturation value. Finally, the best model for the downstream tasks and further analysis as well as comparison with baselines is chosen based on the criteria $\mathcal{C}_{\text{high}}$ defined in (9). This corresponds to an encoder architecture with $L = 3$ layers and $32 \times 16 \times 8$ filters, where $\lambda = 0.2$. Recall we concatenate nodal representations from all layers, and we use global average pooling to obtain \mathbf{x}_G .

B. Comparison With Baselines

In this section, we discuss the results of the comparisons with the baselines outlined in Section IV-B.

Two-step GRL followed by classification procedure. Unlike the model in Fig. 2 which is trained in an end-to-end manner, GRL-based classification can also be conducted in a step-wise fashion. The first step is to train the FC reconstruction branch alone ($\lambda = 0$). In the second step, graph embeddings \mathbf{x}_G are constructed from the output node representations \mathbf{X}_C in the first step, and then used in a logistic regression classifier for subject classification. When it comes to training in an end-to-end or stage-wise fashion, 10-fold cross validation is adopted for both schemes with the same random seeds for fair comparison. The performance of these two training procedures is summarized in Table II. As expected, the reconstruction MSE is marginally lower for the two-step procedure, as the encoder is optimized for FC reconstruction performance. When it comes to subject (i.e., graph) classification, the proposed model outperforms the two-step GRL baseline both in terms of accuracy and F score, matching the findings in [65]. All in all, when trained in an end-to-end fashion the proposed model attains a markedly better classification performance, while incurring a marginal penalty in FC reconstruction MSE relative to a baseline that was optimized for this latter objective.




Fig. 3. Graph reconstruction and subject classification performance of the three selected model architectures with $\lambda = 0.1$ and $L = 3$ (number of filters per layer shown in the legend where ‘con’ represents concatenation of node embeddings from each layer and ‘no_con’ indicates that only the node embedding from the last layer is considered). Apparently, the combination of node embedding concatenation and global average pooling attains better performance in graph reconstruction with lower MSE and subject classification with higher accuracy. For performance comparison for model architectures with $L = 1$ and $L = 2$, refer to Appendix B.

For visualization purposes, the model with the highest test accuracy among the 10 folds is retained and re-run on the whole dataset. In more details, the SC networks of all $P = 412$ subjects are fed through the trained encoder-decoder system in Fig. 2 to yield per-subject reconstructed FC networks and the low-dimensional graph embeddings that model the SC-FC relationship. We use the t-SNE algorithm [68] to reduce the dimensionality of the graph embeddings \mathbf{x}_G (56-dimensional vectors) so that they can be visualized in the 2D plane. The results are depicted in Fig. 5(a). We can distinguish a separation, albeit not perfect, between the two groups, indicating that the learnt graph representations contain valuable information for subject classification. In Fig. 5(b), we also plot the learnt graph embeddings obtained when $\lambda = 0$, i.e., optimizing only the FC reconstruction loss. In that case, there is no apparent separation among the classes. These findings suggest that end-to-end training captures discriminative properties of the subjects’ networks.

Combined FC-SC input. The proposed model learns the representation of the SC-FC relationship to distinguish heavy drinkers from non-drinkers. One valid question is whether using both SC and FC as the combined graph input to a feed-forward GCN-based classifier could also achieve competitive or similar performance. If so, there is no gain in using advanced GRL methods and considering the FC graph reconstruction pipeline as a regression component in the model. Moreover, this would imply that the SC to FC mapping we model is not as informative to distinguish among classes. In this comparative study we examine

the aforementioned question. To this end, we combined both SC and FC networks of each subject into a $2N \times 2N$ adjacency matrix, where the upper left block corresponds to SC and lower right block contains FC. The other entries are set to zero. The combined SC plus FC adjacency matrix is then used in the same GCN architectures as in Table I for a fair comparison. Fig. 6 depicts the classification performance comparison between all 12 different GCN architectures ($L = 1, 2, 3$ and different amounts of filters per layer), with node embedding concatenation and global average pooling. One finds that the proposed method outperforms this baseline in almost all of the architectural settings, and importantly this always holds for those models that yield the highest classification rates. The general trend remains the same without embedding concatenation (i.e., $\mathbf{X}_C = \mathbf{X}^{(L)}$), or, when using other global pooling schemes (max/sum).

FC and SC autoencoders. We also compared the proposed encoder-decoder system against two graph autoencoder baselines that individually process SC and FC, respectively. Fig. 6 shows that the proposed approach exhibits superior performance relative to the autoencoder baselines, both in terms of classification accuracy and F score. This offers further evidence supporting our hypothesis which states that, beyond the information encoded in the graph topology of each individual connectome (SC or FC), the relationship between them appears to provide additional insights and more discriminative learnt representations to improve subject classification (while capturing intrinsic properties of the brain structure-function coupling as a byproduct).




Fig. 4. FC reconstruction MSE and classification accuracy with respect to different λ values, for the three selected model architectures with $32 \times 16 \times 8$, $64 \times 32 \times 16$, $128 \times 64 \times 32$ filter (left to right). As λ increases, the training is biased towards classification; thus, the MSE increases. The classification performance saturates for large λ values as the reconstruction is compromised at these values. The chosen model that maximizes $\mathcal{C}_{\text{high}}$ has $L = 3$ layers with $32 \times 16 \times 8$ filters (left), node embedding concatenation and global average pooling, where $\lambda = 0.2$.




Fig. 5. 2D visualization of the graph embeddings x_G learnt by the proposed model (left) and the two-step baseline method (right). Group-wise separation can be observed on the left but not on the right, supporting the difference in classification performance reported in Table II.

ML with handcrafted graph-theoretic features. Starting from both FC and SC connectomes, we compute network-centric measures that summarize different global graph properties such as modular structure, paths lengths, connectivity, and/or graph components as well as subgraph structures [5], [57], [66]. Specifically, we evaluate average path length, global efficiency, clustering coefficient, graph radius/diameter, transitivity, and graph density for the SC and FC networks of each subject [57]. These summary statistics are collected in a feature vector representing the connectivity patterns of the subject's FC, SC, or both. Using these handcrafted feature vectors and the given labels, we train several classifiers for non-graph inputs. These include logistic

regression, linear and kernel SVMs, decision trees, random forests and gradient-boosted decision trees (GBDTs). Results are presented in Fig. 7. The conclusion is that using handcrafted features relative to learnt representations accounting for symmetries and invariances in graph data, leads to inferior classification performance. Needless to say, this principle has been verified in a multitude of other domains and is at the heart of the deep learning revolution we have witnessed over the last dozen years.

C. Analysis on Reconstructed FC Networks

By virtue of the combined loss function (7) in the model, learnt nodal representations serve as discriminative features to distinguish among the two drinking-related classes. These representations are also used to reconstruct the FC networks $\hat{\Sigma}$, so we expect to see some of this embedded discriminability in the reconstructed graphs themselves. To examine and interpret these patterns, we test for significant between-group (i.e., heavy drinkers and non-drinkers) differences in the topology of the graphs $\hat{\Sigma}$. In more detail, for each potential functional connection $(i, j) \in \mathcal{E}$, i.e., each element $\hat{\Sigma}_{ij}$ in the, say, upper-triangular half of $\hat{\Sigma}$, we perform a t -test with False Discovery Rate (FDR) correction for multiple testing. We report statistically-significant differences between groups at the 5% level ($p < 0.05$), and collections of significant links forming connected subgraphs are examined next. Results are shown in Fig. 8, where the top row depicts the adjacency matrix of the subgraphs and the lower row offers a visualization of the networks themselves.




Fig. 6. Classification performance comparison between proposed model and baselines. The proposed encoder-decoder model (shown in the rightmost red box) outperforms autoencoders that learn representations to reconstruct either FC or SC. It is also consistently superior to a baseline method where both SC and FC are fed as a combined input to a GCN-based feed-forward classifier.




Fig. 7. Classification performance of various ML models using handcrafted graph-theoretic feature vectors. The performance is not competitive with that of the proposed GRL encoder-decoder system marked by the red horizontal line in each subplot.

Fig. 8(a) and 8(c) depict a subgraph of functional connections that are significantly weaker for the group of heavy drinkers. For reference, edge weights in Fig. 8(a) and 8(b) correspond to the average functional connection strengths in the non-drinker group and the drinker group, respectively. Fig. 8(c) shows an important connectivity decrease for brain RoIs that can be traced to the parietal, cingulate, temporal, occipital, and frontal cortical regions. Recent literature has identified these regions as being susceptible to adverse effects due to binge drinking or heavy alcohol consumption. For instance, [69] reports that past and recent patterns of intermittent heavy alcohol consumption are associated with reduced frontal cortical thickness. A high dose of alcohol detrimentally affects a number of functions associated with the prefrontal and temporal lobes [70]; and marked alterations in the inferior and superior fronto-occipital fasciculus white matter bundles may underlie the visual processing deficits found in alcohol-dependent subjects [71]. The spatial breadth of the observed FC sparsification suggests that alcohol may not only affect localized regional brain linkages, but also affect global patterns of brain functional organization; see [72], [73] for related findings.

Interestingly, the observed effect of drinking does not only reflect FC weakening and sparsification. A different subgraph

shown in Fig. 8(b) and 8(d), illustrates functional edges that were found to be significantly stronger in the group of heavy-drinkers. Nodes in said subgraph span various cortical RoIs and can be attributed to a neural compensation or equalization effect, which has been investigated extensively in the field of cognitive analysis [74]. Prior findings on neural compensation suggest it correlates with a combination of brain structure and function-related effects, hence it is more likely to be revealed in models that jointly account for SC and FC as we do here. All in all, our results are aligned with the literature [75] that reports the brain tends to establish additional connections to compensate for neuro-functional damage caused by alcohol. A deeper neuroscientific interpretation in the broader context of robustness and resilience of human brain activity is well beyond the scope of this methodological paper. Nonetheless, we believe these preliminary analyses can be valuable from the perspective of explainability, further motivating the application of GRL advances for neuroimaging data research.




Fig. 8. Subgraphs with edges that were found to be significantly (a), (c) weaker, and (b), (d) stronger among the group of heavy drinkers ($p < 0.05$). For visualization purposes, in the bottom plots edges are color-coded to match the color assigned to the source node. The difference in sparsity (edge densities) between the subgraphs in the (a), (c) and (b), (d) appears to suggest a marked weakening of functional links in heavy-drinkers, with some limited level of compensation via strengthened connections.

VI. CONCLUDING SUMMARY AND OUTLOOK

We developed a GCN-based encoder-decoder system to learn low-dimensional embeddings that carry information about the relationship between brain structural and functional connectomes. Although there have been ample SC- and FC-based studies to investigate the difference between brain connectomes of patients and healthy controls, using GRL to model the structure-function coupling for subject (i.e., graph) classification has been unexplored so far. By putting forth a flexible multi-objective framework the learnt representations can be used in various types of downstream tasks, thus markedly broadening the range of network neuroscience applications. In the proposed model, two learning objectives are considered. Given SC networks as inputs, the first objective is to reconstruct the FC of the same subject as a regression problem. The second objective is to generate graph-level embeddings that capture individual subject's variability and are discriminative among the classes of heavy drinkers and non-drinkers. This enables supervised subject-level classification.

We conducted a rather broad exploration of the architectural design space. Specifically, we investigated a wide range of GCN models and also tested different global pooling mechanisms. For deep GCNs, nodal embeddings that concatenate the output of earlier layers were shown to improve performance in both graph reconstruction and subject classification. Through experimental comparisons with a broad lineup of baseline methods,

we showed that: (i) end-to-end multi-objective training attains better classification performance than a two-step method where graph reconstruction and classification are trained separately; and (ii) the proposed GRL model could extract succinct, yet informative latent representations of the SC-FC mapping. Relative to baseline models that learn from SC, FC (using autoencoders), or both networks fed as inputs to a feed-forward GCN, the obtained SC-FC representations offer additional information, leading to higher classification accuracy and F score. Moreover, (iii) the proposed method outperformed ML models trained on hand-crafted features collecting various graph-theoretic summary statistics computed from the FC and SC networks. This highlights the superior discriminative power of representations learnt using geometric deep learning, which favorably exploit symmetries and invariances in graph data. Finally, from the reconstructed FC for all subjects, statistically significant differences in edges between drinkers and non-drinkers were revealed. Our findings suggest a marked weakening of functional links in heavy-drinkers, with some limited level of neural compensation in the form of sparse strengthened connections.

This being the first work to build a GRL model of SC-FC brain networks, several extensions are possible. For example, we used one-hot encoding as the initial nodal attributes for each node. An open challenge concerns situations where nodal attributes can be collected from meaningful subject-related information, such as the physical volume of ROIs. In the convolutional layers of the encoder, one could also adopt higher-order graph filters to further gain in expressive power. Node- or edge-varying graph filters are other valuable alternatives in this direction, possibly departing from convolutional models. Rich temporal information within the BOLD signals can also be leveraged via spatio-temporal graph convolutions and dynamic FC networks. Also, attention mechanisms can also be applied for the global pooling. Our focus here was on a simple, yet powerful architecture to demonstrate conceptual advantages of the model – not to push e.g., subject classification accuracy to the limit. In a broader context, the proposed multi-task GRL framework could be applied (modulo application-specific modifications) to other domains beyond network neuroscience. We believe it is especially well suited to combine dynamic or multi-view (multi-aspect) network data, while being capable to extract interpretable and discriminative patterns for supervised learning applications. Examples include snapshots of traffic networks, dynamic brain connectomes at consecutive time stamps, or sensor-collected human co-location graphs along with Facebook ties for friendship recommendation via link prediction.

APPENDIX A SUBJECT LABEL DELINEATION

The HCP dataset contains more subjects than those $P = 412$ used in our experiments. Fig. 9 depicts the distribution of subjects with different levels of drinking habits, as measured by the lifetime maximum number of drinks consumed in a single day. In our study, we consider the 191 subjects with no more than 3 drinks as the group of non-drinkers, and label the 221 subjects with 16 or more drinks as the group of heavy drinkers. This




Fig. 9. Distribution of subjects with different levels of drinking habits, as measured by the lifetime maximum number of drinks consumed in single day. In our experiments, we consider the 191 subjects with no more than 3 drinkers as the non-drinker group (most left-wise bin), and label the 221 subjects with 16 or more drinks as heavy drinkers (most right-wise bin).




Fig. 10. Graph reconstruction and subject classification performance of selected model architectures with $\lambda = 0.1$ and $L = 2$.

way, the two classes are balanced, well separated, and contain sufficient data points to effectively train and test the proposed GRL-based classifier.

APPENDIX B PERFORMANCE OF SHALLOWER GCN ENCODERS

For completeness, we present the MSE and accuracy values achieved by selected shallower GCN encoder architectures with only one or two layers. Due to space constraints, we only select and present 2 architectures from each layer setting. As shown in Figs. 10 and 11, node embedding concatenation together with global average pooling offers better performance with respect to FC reconstruction and subject classification in most architecture settings. While MSE performance is comparable with the chosen model ($L = 3$), the classification accuracy of these shallower encoders tends to degrade.




Fig. 11. Graph reconstruction and subject classification performance of selected model architectures with $\lambda = 0.1$ and $L = 1$. Node embedding concatenation is not available in the single-layer setting.

TABLE III
GIN-BASED CLASSIFIER ON THRESHOLDED SC

Threshold	Accuracy	F score
20%	0.524 ± 0.070	0.528 ± 0.141
40%	0.529 ± 0.056	0.567 ± 0.073
60%	0.522 ± 0.029	0.559 ± 0.279
80%	0.515 ± 0.034	0.489 ± 0.320

TABLE IV
GIN-BASED CLASSIFIER ON THRESHOLDED FC

Threshold	Accuracy	F score
20%	0.554 ± 0.085	0.564 ± 0.097
40%	0.549 ± 0.055	0.569 ± 0.055
60%	0.520 ± 0.079	0.533 ± 0.099
80%	0.512 ± 0.085	0.523 ± 0.107

APPENDIX C GRAPH ISOMORPHISM NETWORK

For the sake of comparison, we add an additional baseline that implements the graph isomorphism network (GIN)-based encoder [31] in a feed-forward classifier to test the predictions of subject drinking labels using a single brain network modality (SC or FC alone). We opt to use the same architecture settings (3 layers with dimension $32 \times 16 \times 8$). As GIN works on binary graphs, we tested different threshold levels to binarize the weighted brain networks by keeping a certain portion of the edges. For each level, a 10-fold cross validation is conducted where each time the whole dataset is partitioned into 80% training, 10% validation and 10% test set. We compute the classification accuracy and F score and report the results in Tables III and IV of input SC and input FC, respectively.

From the tables we can see that the threshold that retains a moderate amount of edges (40%) gives the best classification

performance. This is as expected since keeping too few edges will cause significant information loss from the original graphs, while retaining too many edges leads dense binary graphs with less discriminative power. The results shown in the tables above are inferior to the results presented in Table I, which again suggests that the representations learnt from the SC-FC relationship provide more discriminative power than the information embedded within a single brain network modality, regardless of the GNN models selected for the encoder. Although GIN has been shown effective in subject classification tasks from neuroimaging data [41], the information loss due to binary quantization of e.g., FC links may lead to inferior performance in some cases. As one of the unique features in GIN is the concatenation of node embeddings, which we also implemented here to form \mathbf{X}_C , we believe that GCN with node embedding concatenation suits better for our purpose and helps capture the high-resolution differences of the structural connection strengths between drinkers and non-drinkers.

ACKNOWLEDGMENT

Part of the results in this paper were presented at the 2020 ICASSP Conference [1].

REFERENCES

- [1] Y. Li, R. Shafipour, G. Mateos, and Z. Zhang, “Supervised graph representation learning for modeling the relationship between structural and functional brain connectivity,” in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 9065–9069.
- [2] 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, Aug. 2010.
- [3] A. Fornito, A. Zalesky, and E. Bullmore, *Fundamentals of Brain Network Analysis*. New York, NY, USA: Academic, 2016.
- [4] J. D. Power et al., “Functional network organization of the human brain,” *Neuron*, vol. 72, no. 4, pp. 665–678, 2011.
- [5] E. Bullmore and O. Sporns, “Complex brain networks: Graph theoretical analysis of structural and functional systems,” *Nature Rev. Neurosci.*, vol. 10, no. 3, pp. 186–198, 2009.
- [6] O. Sporns, *Networks of the Brain*. Cambridge, MA, USA: MIT Press, 2010.
- [7] F. Abdelnour, H. U. Voss, and A. Raj, “Network diffusion accurately models the relationship between structural and functional brain connectivity networks,” *Neuroimage*, vol. 90, pp. 335–347, 2014.
- [8] Y. Li and G. Mateos, “Identifying structural brain networks from functional connectivity: A network deconvolution approach,” in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2019, pp. 1135–1139.
- [9] C. J. Honey et al., “Predicting human resting-state functional connectivity from structural connectivity,” *Proc. Nat. Acad. Sci.*, vol. 106, no. 6, pp. 2035–2040, 2009.
- [10] J. Goñi et al., “Resting-brain functional connectivity predicted by analytic measures of network communication,” *Proc. Nat. Acad. Sci. USA*, vol. 111, no. 2, pp. 833–838, 2014.
- [11] C. J. Stam et al., “The relation between structural and functional connectivity patterns in complex brain networks,” *Int. J. Psychophysiol.*, vol. 103, pp. 149–160, 2016.
- [12] 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.
- [13] Y. Li, R. Shafipour, G. Mateos, and Z. Zhang, “Mapping brain structural connectivities to functional networks via graph encoder-decoder with interpretable latent embeddings,” in *Proc. IEEE Glob. Conf. Signal Inf. Process.*, 2019, pp. 1–5.
- [14] T. Sarwar, Y. Tian, B. T. Yeo, K. Ramamohanarao, and A. Zalesky, “Structure-function coupling in the human connectome: A machine learning approach,” *NeuroImage*, vol. 226, 2021, Art. no. 117609.
- [15] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *Proc. Int. Conf. Learn. Representations*, 2017, pp. 1–14.
- [16] D. Duvenaud et al., “Convolutional networks on graphs for learning molecular fingerprints,” in *Proc. Int. Conf. Neural. Inf. Process. Syst.*, 2015, pp. 2224–2232.
- [17] 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, Feb. 2019.
- [18] R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec, “Graph convolutional neural networks for web-scale recommender systems,” in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 974–983.
- [19] W. L. Hamilton, R. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” in *Proc. Int. Conf. Neural. Inf. Process. Syst.*, 2017, pp. 1025–1035.
- [20] X. Yue et al., “Graph embedding on biomedical networks: Methods, applications and evaluations,” *Bioinformatics*, vol. 36, no. 4, pp. 1241–1251, 2020.
- [21] S. I. Ktena et al., “Metric learning with spectral graph convolutions on brain connectivity networks,” *NeuroImage*, vol. 169, pp. 431–442, 2018.
- [22] T. N. Kipf and M. Welling, “Variational graph auto-encoders,” in *Proc. NeurIPS Workshop Bayesian Deep Learn.*, 2016, pp. 1–3.
- [23] I. Chami, S. Abu-El-Haija, B. Perozzi, C. Ré, and K. Murphy, “Machine learning on graphs: A model and comprehensive taxonomy,” *J. Mach. Learn. Res.*, vol. 23, no. 89, pp. 1–64, 2022.
- [24] M. Cheung, J. Shi, L. Jiang, O. Wright, and J. M. Moura, “Pooling in graph convolutional neural networks,” in *Proc. Asilomar Conf. Signals Syst. Comput.*, 2019, pp. 462–466.
- [25] “Connectome coordination facility,” Accessed: Oct. 12, 2021. [Online]. Available: <https://www.humanconnectome.org/>
- [26] F. Gama, E. Isufi, G. Leus, and A. Ribeiro, “Graphs, convolutions, and neural networks: From graph filters to graph neural networks,” *IEEE Signal Process. Mag.*, vol. 37, no. 6, pp. 128–138, Nov. 2020.
- [27] 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, Jul. 2017.
- [28] F. Errica, M. Poddar, D. Bacciu, and A. Micheli, “A fair comparison of graph neural networks for graph classification,” in *Proc. Int. Conf. Learn. Representations*, 2020, pp. 1–14.
- [29] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, “A comprehensive survey on graph neural networks,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 1, pp. 4–24, Jan. 2021.
- [30] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, “Spectral networks and locally connected networks on graphs,” in *Proc. Int. Conf. Learn. Representations*, 2014, pp. 1–14.
- [31] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, “How powerful are graph neural networks?,” in *Proc. Int. Conf. Learn. Representations*, 2019, pp. 1–17.
- [32] W. L. Hamilton, *Graph Representation Learn.* Berlin, Germany: Springer, 2020, pp. 1–159.
- [33] W. L. Hamilton, R. Ying, and J. Leskovec, “Representation learning on graphs: Methods and applications,” *IEEE Data Eng. Bull.*, 2017, pp. 52–74.
- [34] L. Gutiérrez-Gómez and J.-C. Delvenne, “Unsupervised network embedding for graph visualization, clustering and classification,” *Appl. Netw. Sci.*, vol. 4, no. 82, pp. 1–21, 2019.
- [35] M. Zhang and Y. Chen, “Link prediction based on graph neural networks,” in *Proc. Int. Conf. Neural. Inf. Process. Syst.*, 2018, pp. 5171–5181.
- [36] A. Tsitsulin, J. Palowitch, B. Perozzi, and E. Müller, “Graph clustering with graph neural networks,” in *Proc. Int. Workshop Mining Learn. Graphs*, 2020, pp. 1–21.
- [37] F. Tian, B. Gao, Q. Cui, E. Chen, and T.-Y. Liu, “Learning deep representations for graph clustering,” in *Proc. AAAI Conf. Artif. Intell.*, 2014, pp. 1293–1299.
- [38] A. Narayanan, M. Chandramohan, R. Venkatesan, L. Chen, Y. Liu, and S. Jaiswal, “Graph2vec: Learning distributed representations of graphs,” in *Proc. Int. Workshop Mining Learn. Graphs*, 2017, pp. 1–8.
- [39] P. Y. Wang, S. Sapra, V. K. George, and G. A. Silva, “Generalizable machine learning in neuroscience using graph neural networks,” *Front. Artif. Intell.*, vol. 4, 2021, Art. no. 618372.
- [40] G. Rosenthal et al., “Mapping higher-order relations between brain structure and function with embedded vector representations of connectomes,” *Nature Commun.*, vol. 9, no. 1, 2018, Art. no. 2178.
- [41] B.-H. Kim and J. C. Ye, “Understanding graph isomorphism network for rs-fMRI functional connectivity analysis,” *Front. Neurosci.*, vol. 14, pp. 1–15, 2020.

[42] X. Zhang, L. He, K. Chen, Y. Luo, J. Zhou, and F. Wang, "Multi-view graph convolutional network and its applications on neuroimage analysis for Parkinson's disease," in *Proc. AMIA Annu. Symp.*, 2018, Art. no. 1147, pp. 1147–1156.

[43] S. Parisot et al., "Disease prediction using graph convolutional networks: Application to autism spectrum disorder and Alzheimer's disease," *Med. Image Anal.*, vol. 48, pp. 117–130, 2018.

[44] S. Gadgil, Q. Zhao, A. Pfefferbaum, E. V. Sullivan, E. Adeli, and K. M. Pohl, "Spatio-temporal graph convolution for resting-state fMRI analysis," in *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Interv.*, 2020, pp. 528–538.

[45] B.-H. Kim, J. C. Ye, and J.-J. Kim, "Learning dynamic graph representation of brain connectome with spatio-temporal attention," in *Proc. Int. Conf. Neural. Inf. Process. Syst.*, 2021, pp. 4314–4327.

[46] A. Ortega, P. Frossard, J. Kovačević, J. M. F. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proc. IEEE*, vol. 106, no. 5, pp. 808–828, May 2018.

[47] 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," in *Proc. IEEE*, vol. 106, no. 5, May 2018, pp. 868–885.

[48] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Proc. Int. Conf. Neural. Inf. Process. Syst.*, 2016, pp. 3844–3852.

[49] J. You, Z. Ying, and J. Leskovec, "Design space for graph neural networks," in *Proc. Int. Conf. Neural. Inf. Process. Syst.*, 2020, Art. no. 1427.

[50] D. C. van Essen et al., "The WU-Minn human connectome project: An overview," *Neuroimage*, vol. 80, pp. 62–79, 2013.

[51] M. F. Glasser et al., "The human connectome project's neuroimaging approach," *Nature Neurosci.*, vol. 19, no. 9, pp. 1175–1187, 2016.

[52] R. S. Desikan 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.

[53] 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.

[54] Z. Zhang et al., "Mapping population-based structural connectomes," *NeuroImage*, vol. 172, pp. 130–145, 2018.

[55] J. Richiardi, S. Achard, H. Bunke, and D. Van De Ville, "Machine learning with brain graphs: Predictive modeling approaches for functional imaging in systems neuroscience," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 58–70, May 2013.

[56] 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.

[57] M. Rubinov and O. Sporns, "Complex network measures of brain connectivity: Uses and interpretations," *Neuroimage*, vol. 52, no. 3, pp. 1059–1069, 2010.

[58] Y. Bai et al., "Unsupervised inductive graph-level representation learning via graph-graph proximity," in *Proc. Int. Joint Conf. Artif. Intell.*, 2019, pp. 1988–1994.

[59] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. Int. Conf. Mach. Learn.*, 2013, pp. 1–6.

[60] H. Dai, B. Dai, and L. Song, "Discriminative embeddings of latent variable models for structured data," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 2702–2711.

[61] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, "Gated graph sequence neural networks," in *Proc. Int. Conf. Learn. Representations*, 2016, pp. 1–20.

[62] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in *Proc. USENIX Symp. Operating Syst. Des. Implementation*, 2016, pp. 265–283.

[63] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feed forward neural networks," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2010, pp. 249–256.

[64] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Representations*, 2015, pp. 1–15.

[65] J. Weston, F. Ratle, H. Mobahi, and R. Collobert, "Deep learning via semi-supervised embedding," in *Neural Networks: Tricks of the Trade*. Berlin, Germany: Springer, 2012, pp. 639–655.

[66] E. D. Kolaczyk, *Statistical Analysis of Network Data: Methods and Models*. New York, NY, USA: Springer, 2009.

[67] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. Berlin, Germany: Springer, 2009.

[68] L. Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, 2008.

[69] Y. Mashhoon, C. Czerkawski, D. J. Crowley, J. E. Cohen-Gilbert, J. T. Sneider, and M. M. Silveri, "Binge alcohol consumption in emerging adults: Anterior cingulate cortical 'thinness' is associated with alcohol use patterns," *Alcoholism Clin. Exp. Res.*, vol. 38, no. 7, pp. 1955–1964, 2014.

[70] J. B. Peterson, J. Rothfleisch, P. D. Zelazo, and R. O. Pihl, "Acute alcohol intoxication and cognitive functioning," *J. Stud. Alcohol*, vol. 51, no. 2, pp. 114–122, 1990.

[71] D. Bagga et al., "Decreased white matter integrity in fronto-occipital fasciculus bundles: Relation to visual information processing in alcohol-dependent subjects," *Alcohol*, vol. 48, no. 1, pp. 43–53, 2014.

[72] N. D. Volkow 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.

[73] M. Oscar-Berman and K. Marinovic, "Alcoholism and the brain: An overview," *Alcohol Res. Health*, vol. 27, no. 2, pp. 125–133, 2003.

[74] 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.

[75] M. H. Parks et al., "Brain fMRI activation associated with self-paced finger tapping in chronic alcohol-dependent patients," *Alcoholism Clin. Exp. Res.*, vol. 27, no. 4, pp. 704–711, 2003.



Yang Li (Member, IEEE) received the B.E. degree in automation from Tianjin University, Tianjin, China in 2013, and the M.Sc. and Ph.D. degrees in electrical engineering from the University of Rochester, Rochester, NY, USA, in 2016 and 2022, respectively. His research interests include machine learning on graphs, with a focus on graph signal processing, and graph neural networks for neuroimaging data analysis.



Gonzalo Mateos (Senior Member, IEEE) received the B.Sc. degree in electrical engineering from the Universidad de la República, Montevideo, Uruguay, in 2005, and the M.Sc. and Ph.D. degrees in electrical engineering from the University of Minnesota, Minneapolis, MN, USA, in 2009 and 2011, respectively. He is currently an Associate Professor with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY, USA, and an Asaro Biggar Family Fellow of data science. During 2013, he was a Visiting Scholar with the CS Department, Carnegie Mellon University, USA. From 2004 to 2006, he was a Systems Engineer at Asea Brown Boveri, Uruguay. His research interests include statistical learning from complex data, network science, decentralized optimization, and graph signal processing.



Zhengwu Zhang received the M.S. and Ph.D. degrees in statistics from Florida State University, Alachua, FL, USA, in 2013 and 2015, respectively. From 2015 to 2017, he was a Postdoctoral Fellow with the Statistical and Applied Mathematical Sciences Institute in Research Triangle Park, North Carolina. He is currently an Assistant Professor of Statistics with the University of North Carolina at Chapel Hill, Chapel Hill, NC, USA. His research interests include brain connectomics, imaging genetics, shape and functional data analysis, computational neuroscience, bayesian statistics, and machine learning.