



HHS Public Access

Author manuscript

Med Image Comput Assist Interv. Author manuscript; available in PMC 2020 November 29.

Published in final edited form as:

Med Image Comput Assist Interv. 2020 October ; 12267: 528–538.

doi:10.1007/978-3-030-59728-3_52.

Spatio-Temporal Graph Convolution for Resting-State fMRI Analysis

Soham Gadgil^{#1}, Qingyu Zhao^{#2}, Adolf Pfefferbaum^{2,3}, Edith V. Sullivan², Ehsan Adeli^{1,2}, Kilian M. Pohl^{2,3}

¹Computer Science Department, Stanford University, Stanford, USA

²School of Medicine, Stanford University, Stanford, USA

³Center of Health Sciences, SRI International, Menlo Park, USA

These authors contributed equally to this work.

Abstract

The Blood-Oxygen-Level-Dependent (BOLD) signal of resting-state fMRI (rs-fMRI) records the temporal dynamics of intrinsic functional networks in the brain. However, existing deep learning methods applied to rs-fMRI either neglect the functional dependency between different brain regions in a network or discard the information in the temporal dynamics of brain activity. To overcome those shortcomings, we propose to formulate functional connectivity networks within the context of spatio-temporal graphs. We train a spatio-temporal graph convolutional network (ST-GCN) on short sub-sequences of the BOLD time series to model the non-stationary nature of functional connectivity. Simultaneously, the model learns the importance of graph edges within ST-GCN to gain insight into the functional connectivities contributing to the prediction. In analyzing the rs-fMRI of the Human Connectome Project (HCP, $N = 1,091$) and the National Consortium on Alcohol and Neurodevelopment in Adolescence (NCANDA, $N = 773$), ST-GCN is significantly more accurate than common approaches in predicting gender and age based on BOLD signals. Furthermore, the brain regions and functional connections significantly contributing to the predictions of our model are important markers according to the neuroscience literature.

1 Introduction

The BOLD signal of rs-fMRI characterizes the intrinsic functional organization of the human brain by measuring its spontaneous activity at rest [3]. One commonly used approach for identifying impact of factors, such as age and sex, on intrinsic functional networks is to predict their values by applying deep neural networks to the rs-fMRI of individual subjects [7,13,14,16]. Accurate predictors of those factors can then enhance the understanding of functional neurodevelopment across the life span [15], characterize developmental disruption caused by neurological disorders, and explain sex-specific differences in cognitive performance [22].

It is not trivial to select an appropriate network architecture for analyzing rs-fMRI as their (average) BOLD signals of each brain region are structured time series. A natural representation of such data are spatio-temporal graphs [24,25], where the temporal graph characterizes the dynamics of brain activity at each region and the spatial graph characterizes the functional interaction between different brain regions. However, most existing deep learning works applied to rs-fMRI analysis fail to consider both aspects simultaneously. Methods that only incorporate spatial graph convolution [13,16,26] often transform the time series data into hand-crafted features, such as partial correlation [17], thereby potentially losing the fine temporal information of the BOLD signal. Methods based on recurrent neural networks (RNN) [6,7,14] can learn temporal features from the BOLD signal but neglect the functional dependency between regions-of-interests (ROIs). On the other hand, methods based on spatio-temporal networks often perform spatial convolution according to the topological arrangement among ROIs in the physical space [13,30], which cannot model interactions among distal ROIs with similar functional properties. To address these issues, we suggest deep neural networks to incorporate spatio-temporal convolution on functional connectivity graphs, i.e., spatio-temporal Graph Convolution (ST-GC) [24,25].

In computer vision, ST-GC Networks (ST-GCN) are popular for solving problems that base prediction on graph-structured time series [5,24,25]. In the context of rs-fMRI analysis, these networks have the potential to automatically extract features that jointly characterize functional connectivity patterns of the brain and their temporal dynamics within the BOLD series. To the best of our knowledge, we are the first to use ST-GCN for building predictive models based on rs-fMRI data. We train our network on short sub-sequences of BOLD time series to model the non-stationary nature of functional connectivity [20,28]. Further, learning of an edge importance matrix associated with the ST-GC operation improves the interpretability of the model as it allows us to identify selective functional connections significantly contributing to the prediction. We apply ST-GCN to predict the age and gender of healthy individuals of two large publicly available rs-fMRI datasets: Human Connectome Project (HCP, $N=1,096$) [21] and National Consortium on Alcohol and Neurodevelopment in Adolescence (NCANDA, $N=773$) [2]. The optimal window size for the sub-sequences highly coincide between the two datasets despite their distinct imaging protocols and data processing pipelines. Furthermore, the resulting prediction accuracy is significantly higher than traditional RNN-based methods. Finally, the learned edge importance localizes meaningful brain regions and functional connections associated with aging effects and sexual differences.¹

2 ST-GCN for rs-fMRI Analysis

We first relate functional networks to spatio-temporal graphs and then define the ST-GC convolution on them. Next, we build the classifier ST-GCN, which uses ST-GC in each layer and models the importance of graph edges in the decision process across layers. Lastly, we train ST-GCN on short BOLD sequences.

¹The code is available at https://github.com/sgadgil6/cnslab_fmri.

Representing Functional Networks as Spatio-Temporal Graphs:

To encode the functional networks captured by rs-fMRI, let $\mathcal{G} := (\mathcal{V}, \mathcal{E})$ (Fig. 1) be an undirected spatio-temporal graph consisting of a set of edges \mathcal{E} capturing temporal and spatial connections between a set of nodes $\mathcal{V} = \{v_{t,i} \mid t = 1, \dots, T; i = 1, \dots, N\}$ defined across N ROIs and T time points. For each ROI and time point, an edge in the temporal graph connects the corresponding node to the node of the same ROI at the preceding time point. The edges of the spatial graph connect all nodes of the same time point, where the weight of an edge is defined by the functional affinity between the corresponding regions. To define the functional affinity, we now assume that spontaneous activation for each of the N regions can be quantified by the average BOLD time series measured within that region. First, those series are concatenated across all subjects within each ROI. Then, the affinity between two regions $d(v_{tj}, v_{ti})$ is defined by the magnitude of correlation between their concatenated series. Note, this affinity is impartial to the time point.

Spatio-Temporal Graph Convolution (ST-GC):

To define a convolution on such a graph structure, we denote $f_{in}(v_{ti})$ as the input feature at node v_{ti} (e.g., the average BOLD signal of ROI i at time t) and $B(v_{ti})$ as the spatio-temporal neighbourhood of v_{ti} , i.e.,

$$B(v_{ti}) := \{v_{qj} \mid d(v_{tj}, v_{ti}) \leq K, \mid q - t \mid \leq \lfloor \Gamma / 2 \rfloor\}, \quad (1)$$

where K defines the size of the spatial neighborhood (i.e., spatial kernel size) and Γ the temporal neighborhood (i.e., temporal kernel size). An ST-GC operation on node v_{ti} with respect to a convolutional kernel $\mathbf{w}(\cdot)$ and a normalization factor Z_{ti} can then be defined as [24]

$$f_{out}(v_{ti}) := \frac{1}{Z_{ti}} \sum_{v_{qj} \in B(v_{ti})} f_{in}(v_{qj}) \cdot \mathbf{w}(v_{qj}). \quad (2)$$

Adopting a similar implementation as in [24], we approximate the spatio-temporal convolutional kernel $\mathbf{w}(\cdot)$ by decomposing it to a spatial graph convolutional kernel $\mathbf{W}_{SG} \in \mathbb{R}^{C \times M}$ represented in the Fourier domain and a temporal convolutional kernel $\mathbf{W}_{TG} \in \mathbb{R}^{M \times \Gamma}$. Specifically, we denote $\mathbf{f}_t \in \mathbb{R}^{N \times C}$ as the C -channel input features of the N ROIs at the t^{th} frame, $\mathbf{f}'_t \in \mathbb{R}^{N \times M}$ as the M -channel output features, and \mathbf{A} as the aforementioned affinity matrix. The spatial graph convolution at a time point t is then defined with respect to the diagonal matrix \mathbf{A} (where $\Lambda^{ii} = \sum_j \Lambda^{ij} + 1$) as [11]

$$\mathbf{f}'_t := \mathbf{\Lambda}^{-\frac{1}{2}} (\mathbf{A} + \mathbf{I}) \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{f}_t \mathbf{W}_{SG}. \quad (3)$$

Next, the temporal convolution is performed on the resulting features. Let $\mathbf{f}'_i \in \mathbb{R}^{M \times T}$ be the features of node v_i defined on the temporal graph of length T with regular grid spacing. We

then perform a standard 1D convolution $\mathbf{f}_i \otimes \mathbf{W}_{TG} \in \mathbb{R}^{T \times T}$ as the final output of ST-GC for v_i .

Classifying BOLD Time Series by ST-GCN:

Our proposed ST-GCN is composed of 3 layers of ST-GC units. The input to ST-GCN are 1-channel spatio-temporal features $\mathbf{f} \in \mathbb{R}^{N \times T}$ representing the average BOLD signals of the N ROIs. Consistent with the setup in [24], each ST-GC layer produces 64-channel outputs with the temporal kernel size $T=11$, a stride of 1, and a dropout rate of 0.5. The output of the last ST-GC layer is fed to a global average pooling and its output vector of length 64 is transformed to *class* probabilities by a fully connected layer with a sigmoid activation.

To determine the importance of spatial graph edges in defining class probabilities, we integrate a positive and symmetric “edge importance” matrix $\mathbf{M} \in \mathbb{R}^{N \times N}$ into our model. This matrix is shared across all ST-GC layers by replacing $\mathbf{A} + \mathbf{I}$ in Eq. (3) with $(\mathbf{A} + \mathbf{I}) * \mathbf{M}$ (element-wise multiplication). As such, while performing spatial graph convolution on node v_{ti} , the contribution from its neighbouring nodes $B(v_{ti})$ will be re-scaled according to the importance weights learned in the i^{th} row of \mathbf{M} . Thus, the diagonal entries of \mathbf{M} (self-connection) quantify importance for each ROI, while off-diagonal entries do so for each functional connection. Note, the original proposal by Yan et al. [24] learns a separate \mathbf{M} for each individual ST-GC layer. Their strategy generally results in negative, asymmetric importance matrices that are difficult to interpret and vary across layers. We ease the interpretability by enforcing \mathbf{M} to be consistent across layers and be both positive and symmetric.

Training ST-GCN:

Recent rs-fMRI studies have revealed that patterns of intrinsic functional connectivity are not stationary across the full rs-fMRI scan but exhibit considerable fluctuations. These dynamics are often analyzed by dividing the entire rs-fMRI sequence into sub-sequences according to a fixed window size (often chosen empirically) and then assessing the connectivity within these segments. Accordingly, we also consider training ST-GCN on sub-sequences of window size T' sampled from the full data. Specifically, at each training iteration, we sample a sub-sequence of length T' starting at a random time frame from the full sequence of each training subject in the mini-batch. The models are then trained by stochastic gradient descent with a learning and weight decay rate of 0.001. At the testing stage, we sample S sub-sequences, each starting at a random time frame, of length T' for each testing subject, derive the ST-GCN prediction for each sub-sequence, and perform a simple voting to produce the final subject-level prediction, e.g., the average of sigmoid values in the case of binary classification (Fig. 1).

3 Experiments

Understanding the dramatic neurodevelopment and the emerging sexual differences during adolescence is an important topic in neuroscience [19,22]. We investigate the age and sex difference using ST-GCN and a variety of baseline approaches on the NCANDA dataset. We further investigate sex differences in young adults of the Human Connectome Project (HCP)

S1200 [21]. Note, identifying significant aging effects on HCP is unlikely due to functional organization reaching maturity after young adulthood [23].

NCANDA:

The publicly released baseline data consisted of 773 rs-fMRI scans (269 frames, TR = 2 s). Among the 773 adolescents (ages 12–21 years, 388 younger than the mean age of 16 years vs. 385 older adolescents, 376 male vs. 397 females). 638 met the no-to-low alcohol drinking criteria of NIAAA (373 young vs. 265 old adolescents, 315 boys vs. 323 girls) [2]. Each rs-fMRI scan was preprocessed by the NCANDA pipeline, which registered the mean BOLD image to subject-specific T1 MRI and then non-rigidly to the standard MNI space [29]. The cortical surface was parcellated to $N=34$ ROIs according to [12]. The average BOLD signal in each ROI was normalized to z -scores.

HCP:

The data set consisted of rs-fMRI of 1096 young adults (ages 22–35 years). Excluding 5 rs-fMRIs with less than 1200 frames, we used the first session (15 min, $T=1200$ frames, TR = 0.72 s) for each of the 498 females and 593 males. Each rs-fMRI went through the minimal processing pipeline of HCP, *fMRISurface* [9], which mapped each volume time series to the standard CIFTI grayordinates space. The cortical surface was parcellated to $N=22$ major ROIs [8], and the average BOLD signal in each ROI was normalized to z -scores.

Experimental Setup:

On NCANDA, ST-GCN first distinguished younger from older participants by performing 5-fold cross-validation on the 773 rs-fMRIs. The training and testing were repeated on sub-sequences of different window sizes, from short segments of $T'=16$ to the full sequence. The number of sub-sequences used for voting in the testing stage was fixed at $S=64$ ($S=1$ for the full sequence). To ensure the prediction results were not confounded by alcohol drinking, the test accuracy was also recorded for the no-to-low drinking cohort, denoted as “ST-GCN-no-ex”. Based on the optimal window size determined through cross-validation, we trained the model on the entire data set to produce an edge importance matrix summarizing the aging effects within the entire cohort. To reduce uncertainty in the estimation caused by stochastic gradient descent, we repeated the training 20 times and derived the “average” edge importance matrix as our final outcome. Next, this experiment was repeated with respect to sex classification on NCANDA and HCP.

Baselines:

The simplest approach for our comparison was a Multi-Layer Perceptron (MLP) applied to the upper triangular correlation matrix, which quantified the static functional connectivity between ROIs during the full scan. The MLP had 2 hidden layers, each with 64 neurons and ReLu activation. The next two baseline methods relied on end-to-end training on the BOLD time series based on Long Short-Term Memory (LSTM) model [10], a form of RNN frequently used to analyze rs-fMRI data [7,14]. Adopting a similar implementation as in [14], the first implementation of LSTM consisted of a recurrent cell with hidden states of dimension 256. The output of the last hidden state was fed into a fully connected layer (with

a dropout of 0.5) to produce a final label. The second implementation (**GC-LSTM**) first extracted features from the BOLD signal at each time frame using spatial graph convolution. The features were then fed into an LSTM model for temporal analysis. The number of voting sequences for these two LSTM-based methods were again fixed to $S = 64$.

4 Results and Analysis

NCANDA:

Figure 2 shows the prediction accuracy derived from the 5-fold cross-validation for different methods. The accuracy scores measured on all NCANDA subjects were not significantly different from results on the no-to-low drinking cohort (denoted as “no-ex”, paired t -test $p > 0.5$), indicating that the predictions were not confounded by alcohol consumption in individuals. On both data sets, all three deep-learning-based methods were generally more accurate than MLP when applied to shorter sub-sequences. ST-GCN achieved significantly higher accuracy than the two LSTM-based methods across different window sizes ($p < .0001$, two-sample t -test). Performing further architecture search for the two LSTM-based methods, such as varying the hidden-state dimension and adding more LSTM layers, did not improve the results in Fig. 2. Moreover, adding more convolution layers or feature channels to GC-LSTM reduced the accuracy substantially, which indicated that the strategy of training spatial graph convolution on each individual frame without incorporating temporal convolution was flawed. With respect to the no-to-low drinking participants, the highest prediction accuracy of ST-GCN with respect to age was 77.7% ($T' = 24$ or 52.8 s) and sex was 79.8% ($T' = 32$ or 70.4 s). The accuracy of sex prediction was relatively stable across $T' = [24, 40]$ (52.8–88 s). These predictions were not confounded by subject motion based on the insignificant correlation between the prediction scores and the number of outlier frames in each rs-fMRI.

We visualized the edge importance matrices associated with the optimal models in Fig. 3 with the images on the left (Fig. 3a+c) showing the importance of each ROI (diagonal entries) for prediction and the images on the right (Fig. 3b+d) showing functional connections (off-diagonal entries) with importance value higher than 0.3. For sex prediction, the most important ROI identified by ST-GCN was the inferior temporal lobe, which echoed findings from other resting-state studies [4] and from a structural MRI analysis on the NCANDA cohort [27]. We also identified a significant effect in the frontal-posterior-cingulate (PCC) connection (red in Fig. 3b), which defines the default mode network, a signature intrinsic network frequently linked to sexual differences [18]. For age prediction, the most critical ROIs were the supramarginal and par opercularis regions (Fig. 3c). Their functional connection (red in Fig. 3d) defined the inferior temporo-parieto-frontal network, which was shown to decrease in older adolescents within the NCANDA cohort by a longitudinal rs-fMRI study [29]. All the results above demonstrate that our strategy of edge importance learning can accurately localize functional properties of the brain related to significant aging and sex effects.

HCP:

Figure 4a plots the accuracy in predicting sex on the HCP dataset in relation to those of prior studies by Weis et al. [22] and Smith et al. [19]. These works applied linear classifiers and used correlation coefficients as input features respectively, which was similar to MLP. Note that their results can not be strictly compared to ours as different cohorts and data processing pipelines were used for analysis². However, their accuracies were similar to the ones of MLP, which were suboptimal compared to those produced by the deep-learning-based methods applied to sub-sequences of $T' \in [32, 128]$ (23.04–92.16 s). This finding indicates that the dynamical properties of functional interactions among brain regions can not be fully captured by the static correlation coefficients and require more comprehensive spatio-temporal modeling.

An example of more comprehensive spatio-temporal modeling was ST-GCN, which achieved the highest accuracy of 83.7% at $T' = 128$ (92.16 s) and produced higher accuracies than the two LSTM-based methods across all window sizes. Increasing the number of voters did not further increase testing accuracy (Fig. 4b). Moreover, ST-GCN produced similarly optimal accuracy scores for $T' \in [64, 128]$ (46.08–92.16 s). This range was highly consistent with the one revealed in the NCANDA experiment despite the difference in imaging protocols, length of BOLD signal, brain parcellation, and data processing pipelines between the two studies. These results also converge with recent understanding in neuroscience literatures stating that dynamical functional connectivity generally have dwell times of tens of seconds [1,20]. By visualizing the edge importance matrix derived from the HCP subjects (Fig. 4c), we found that regions with significant sexual differences in young adults were spatially more concentrated compared to adolescents (the NCANDA experiment) and mainly located in the visual cortex. This phenomenon could be potentially linked to the evidence that sexual dimorphism in intrinsic functional organization diminishes with age (sex-age interaction) during adolescence [18].

5 Conclusion

We introduced a framework for analyzing rs-fMRI data based on spatio-temporal graph convolution networks. By analyzing the rs-fMRI data of two large-scale neuroimaging studies, we showed that ST-GCN could accurately predict age and gender of the study participants based on short sequences of BOLD time series. The similar optimal window sizes (of the sub-sequences) between the two datasets highlighted the usage of short BOLD sequences for modeling dynamic functional connectivity. Future work will focus on defining ST-GC with respect to non-static graph structures accommodating dynamic functional states, exploring automatic determination of sliding window size, applying ST-GCN to a fine-grained brain parcellation, and identifying functional biomarkers linked to neuropsychiatric disorders. Accomplishing these goals will then show if our strategy for learning graph edge importance within the context of model interpretation is valuable for advancing knowledge in neuroscience.

²The 87% accuracy in [19] on 131 HCP subjects was based on multi-modal data.

Acknowledgment.

This research was supported in part by NIH grants AA021697, AA005965, AA013521, and AA010723.

References

1. Allen E, Damaraju E, Plis S, Erhardt E, Eichele T, Calhoun V: Tracking whole-brain connectivity dynamics in the resting state. *Cereb. Cortex* 24(3), 663–676 (2014). 10.1093/cercor/bhs352 [PubMed: 23146964]
2. Brown S, Brumback T, Tomlinson K, et al.: The national consortium on alcohol and neurodevelopment in adolescence (NCANDA): a multisite study of adolescent development and substance use. *J. Stud. Alcohol Drugs* 76(6), 895–908 (2015). 10.15288/jsad.2015.76.895 [PubMed: 26562597]
3. Buckner R, Krienen F, Yeo B: Opportunities and limitations of intrinsic functional connectivity MRI. *Nat. Neurosci* 16(7), 832–837 (2013) [PubMed: 23799476]
4. Conrin SD, et al.: From default mode network to the basal configuration: sex differences in the resting-state brain connectivity as a function of age and their clinical correlates. *Front. Psychiatry* 9, 365 (2018) [PubMed: 30150944]
5. Covert I, et al.: Temporal graph convolutional networks for automatic seizure detection. In: *Machine Learning for HealthCare* (2019)
6. Cui Y, et al.: Identifying brain networks of multiple time scales via deep recurrent neural network. In: *21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part III*, pp. 284–292 (September 2018)
7. Dvornek NC, Ventola P, Pelphrey KA, Duncan JS: Identifying autism from resting-state fMRI using long short-term memory networks In: Wang Q, Shi Y, Suk H-I, Suzuki K (eds.) *MLMI 2017. LNCS*, vol. 10541, pp. 362–370. Springer, Cham (2017). 10.1007/978-3-319-67389-9_42
8. Glasser M, et al.: A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178 (2016) [PubMed: 27437579]
9. Glasser M, et al.: The minimal preprocessing pipelines for the human connectome project. *NeuroImage* 80, 105 (2013). 10.1016/j.neuroimage.2013.04.127 [PubMed: 23668970]
10. Hochreiter S, Schmidhuber J: Long short-term memory. *Neural Comput.* 9(8), 1735–1780 (1997) [PubMed: 9377276]
11. Kipf TN, Welling M: Semi-supervised classification with graph convolutional networks. In: *ICLR* (2017)
12. Klein A, Tourville J: 101 labeled brain images and a consistent human cortical labeling protocol. *Front. Neurosci* 6, 171 (2012). 10.3389/fnins.2012.00171 [PubMed: 23227001]
13. Ktena SI, et al.: Distance metric learning using graph convolutional networks: application to functional brain networks In: Descoteaux M, Maier-Hein L, Franz A, Jannin P, Collins DL, Duchesne S (eds.) *MICCAI 2017. LNCS*, vol. 10433, pp. 469–477. Springer, Cham (2017). 10.1007/978-3-319-66182-7_54
14. Li H, Fan Y: Brain decoding from functional MRI using long short-term memory recurrent neural networks In: Frangi AF, Schnabel JA, Davatzikos C, Alberola-López C, Fichtinger G (eds.) *MICCAI 2018. LNCS*, vol. 11072, pp. 320–328. Springer, Cham (2018). 10.1007/978-3-030-00931-1_37
15. Li H, Satterthwaite TD, Fan Y: Brain age prediction based on resting-state functional connectivity patterns using convolutional neural networks. In: *ISBI* (2018)
16. Li X, Dvornek NC, Zhou Y, Zhuang J, Ventola P, Duncan JS: Graph neural network for interpreting Task-fMRI biomarkers In: Shen D, et al. (eds.) *MICCAI 2019. LNCS*, vol. 11768, pp. 485–493. Springer, Cham (2019). 10.1007/978-3-030-32254-0_54
17. Marrelec G, et al.: Partial correlation for functional brain interactivity investigation in functional MRI. *Neuroimage* 32(1), 228–237 (2006) [PubMed: 16777436]
18. Müller-Oehring E, et al.: Influences of age, sex, and moderate alcohol drinking on the intrinsic functional architecture of adolescent brains. *Cereb. Cortex* 28, 1–15 (2017). 10.1093/cercor/bhx014

19. Smith SM, Vidaurre D, Beckmann CF, et al.: Functional connectomics from resting-state fMRI. *Trends Cogn. Sci* 17(12), 666–682 (2013) [PubMed: 24238796]
20. Taghia J, Ryali S, et al.: Bayesian switching factor analysis for estimating time-varying functional connectivity in fMRI. *NeuroImage* 155, 271–290 (2017). 10.1016/j.neuroimage.2017.02.083 [PubMed: 28267626]
21. Van Essen DC, et al.: The WU-Minn human connectome project: an overview. *Neuroimage* 80, 62–79 (2013) [PubMed: 23684880]
22. Weis S, Patil KR, Hoffstaedter F, Nostro A, Yeo BTT, Eickhoff SB: Sex classification by resting state brain connectivity. *Cereb. Cortex* 30, 824 (2019). 10.1093/cercor/bhz129
23. Westlye L, et al.: Life-span changes of the human brain white matter: diffusion tensor imaging (DTI) and volumetry. *Cereb. Cortex* 20, 2055–68 (2010). 10.1093/cercor/bhp280 [PubMed: 20032062]
24. Yan S, Xiong Y, Lin D: Spatial temporal graph convolutional networks for skeleton-based action recognition. In: AAAI (2018)
25. Yu B, Yin H, Zhu Z: Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. In: IJCAI (2018)
26. Zhang Y, Bellec P: Functional annotation of human cognitive states using graph convolution networks. In: NeurIPS 2019 Workshop Neuro AI (2019)
27. Zhao Q, Adeli E, Pfefferbaum A, Sullivan EV, Pohl KM: Confounder-aware visualization of ConvNets In: Suk H-I, Liu M, Yan P, Lian C (eds.) MLMI 2019. LNCS, vol. 11861, pp. 328–336. Springer, Cham (2019). 10.1007/978-3-030-32692-0_38
28. Zhao Q, Honnorat N, Adeli E, Pfefferbaum A, Sullivan EV, Pohl KM: Variational autoencoder with truncated mixture of Gaussians for functional connectivity analysis In: Chung ACS, Gee JC, Yushkevich PA, Bao S (eds.) IPMI 2019. LNCS, vol. 11492, pp. 867–879. Springer, Cham (2019). 10.1007/978-3-030-20351-1_68
29. Zhao Q, et al.: Longitudinally consistent estimates of intrinsic functional networks. *Hum. Brain Mapp* 40, 2511–2528 (2019). 10.1002/hbm.24541 [PubMed: 30806009]
30. Zhao Y, et al.: Modeling 4D fMRI data via spatio-temporal convolutional neural networks (ST-CNN) In: Frangi AF, Schnabel JA, Davatzikos C, Alberola-López C, Fichtinger G (eds.) MICCAI 2018. LNCS, vol. 11072, pp. 181–189. Springer, Cham (2018). 10.1007/978-3-030-00931-1_21

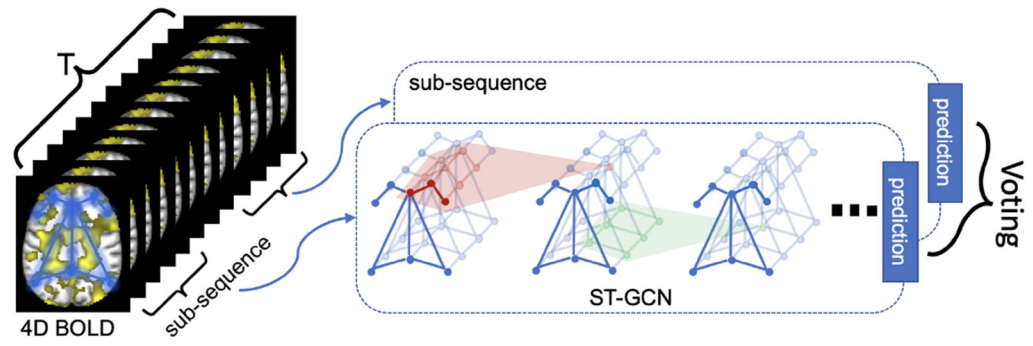


Fig. 1. Framework of classifying BOLD time series by applying a spatio-temporal graph convolutional network to sub-sequences.




Fig. 2. Accuracy of age and sex prediction on the NCANDA dataset. “no-ex” denotes the accuracy score confined to the no-to-low drinking cohort.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Fig. 3. Importance maps for age (a, b) and sex (c, d) prediction of the NCANDA study. (a, c): importance of functional dynamics within each ROI; (b, d): functional connections between ROIs with importance higher than 0.3 are shown by the ROIs having the same color. Highlighted by the red arrows are the default mode network for sex prediction and the inferior temporo-parieto-frontal network for age prediction.

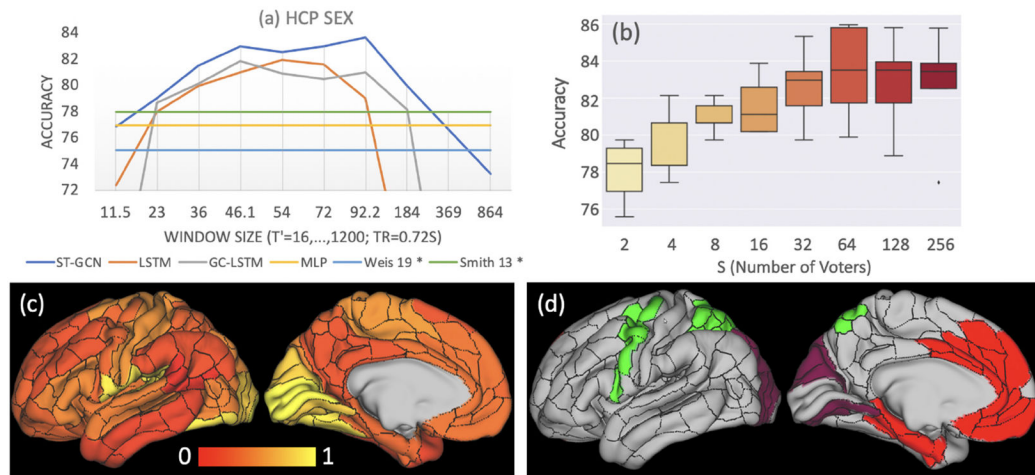


Fig. 4.

Experiments on HCP: (a) sex prediction accuracy w.r.t window size. * denotes results from prior studies. (b) Distribution of prediction accuracy scores from cross-validation w.r.t. the number of sub-sequences used for voting. (c) Importance of functional dynamics within each ROI; (d) functional connections between ROIs with importance higher than 0.3 are shown by the ROIs having the same color.