

HHS Public Access

Author manuscript

Proc IEEE Int Symp Biomed Imaging. Author manuscript; available in PMC 2019 October 01.

Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2019 April; 2019: 414-417. doi:10.1109/ISBI.2019.8759531.

GRAPH CONVOLUTIONAL NEURAL NETWORKS FOR ALZHEIMER'S DISEASE CLASSIFICATION

Tzu-An Song^{1,*}, Samadrita Roy Chowdhury^{1,*}, Fan Yang¹, Heidi Jacobs², Georges El Fakhri², Quanzheng Li², Keith Johnson², and Joyita Dutta^{1,2}

¹Department of Electrical and Computer Engineering, University of Massachusetts Lowell, Lowell, MA

²Gordon Center for Medical Imaging, Massachusetts General Hospital & Harvard Medical School, Boston, MA

Abstract

Graph convolutional neural networks (GCNNs) aim to extend the data representation and classification capabilities of convolutional neural networks, which are highly effective for signals defined on regular Euclidean domains, e.g. image and audio signals, to irregular, graph-structured data defined on non-Euclidean domains. Graph-theoretic tools that enable us to study the brain as a complex system are of great significance in brain connectivity studies. Particularly, in the context of Alzheimer's disease (AD), a neurodegenerative disorder associated with network dysfunction, graph-based tools are vital for disease classification and staging. Here, we implement and test a multi-class GCNN classifier for network-based classification of subjects on the AD spectrum into four categories: cognitively normal, early mild cognitive impairment, late mild cognitive impairment, and AD. We train and validate the network using structural connectivity graphs obtained from diffusion tensor imaging data. Using receiver operating characteristic curves, we show that the GCNN classifier outperforms a support vector machine classifier by margins that are reliant on disease category. Our findings indicate that the performance gap between the two methods increases with disease progression from CN to AD. We thus demonstrate that GCNN is a competitive tool for staging and classification of subjects on the AD spectrum.

Keywords

Graph CNN; convolutional neural network; Alzheimer's disease; classification

1. INTRODUCTION

Alzheimer's disease, a neurodegenerative disorder that currently affects over 30 million people worldwide, is associated with dysfunction of the brain connectivity network and is considered a *connectopathy* or connectivity disorder. A number of recent studies have attempted to diagnose and stage AD using image-derived brain connectivity graphs [1, 2]. Graph-based classifiers are becoming increasingly significant in the field of connectomics

Author Manuscript

^{*}These two authors contributed equally.

which uses advanced neuroimaging techniques to map brain connectivity with the ultimate goal of unraveling mechanistic details of cognition, behavior, and disease [3]. Graphs are generic data representation forms that are associated with irregular, non-Euclidean domains. In recent years, convolutional neural networks (CNNs) have revolutionized data analytics for signals with underlying Euclidean structure e.g. audio and images. CNN-based techniques for data representation and classification have outperformed more traditional machine learning approaches for datasets with Euclidean structure by a staggering margin. The non-Euclidean nature of the domains in which graph signals are defined implies that common concepts such as local neighborhood definitions (at the domain level) and shift invariance or equivariance (at the system level) do not have natural definitions for these datasets. Consequently, the convolution operation, which lies at the crux of CNN architectures, is not well-defined on such non-Euclidean domains. Very recently, there has been a series of endeavors toward deep geometric learning with the goal of extending CNN models to non-Euclidean datasets defined on irregularly-structured domains such as graphs and manifolds [4–6]. In this paper, we implement and validate a multi-class graph CNN (GCNN) classifier for AD. The classifier uses structural connectivity inputs in the form of graph Laplacians derived from diffusion tensor imaging (DTI) to generate the cognitive status category label as its output. Four class labels are used in our work: cognitively normal (CN), early mild cognitive impairment (EMCI), late mild cognitive impairment (LMCI), and Alzheimer's disease (AD). The GCNN implementation details and the DTI data processing pipeline are presented in section 2. Training and validation details and GCNN performance measures are provided in section 3. Section 4 provides a summary of this work and comments on future directions.

2. METHODS

2.1. Graph Convolutions

Two key properties that are leveraged by CNNs and which account for their efficiency are local connectivity and shift in variance. CNN feature extraction is based on *receptive fields* that operate on local neighborhoods. This leads to global parameter sharing across spatial locations thus exploiting translation invariance. Parameter sharing in neural networks significantly reduces the number of unknown parameters or weights to be computed in the training phase and drastically reduces the computing overhead. Though these ideas are fundamental to images and time series, a direct generalization of these for graphs is non-trivial, as real-world graphs lack a grid structure, and not all nodes have the same number of neighbors following the same order. The problem could either be tackled in the spectral domain, or locality may be extended to graph signals by utilizing their irregular neighborhood structure with varying numbers of neighbors per node.

A graph is defined as $\mathscr{G} = (\mathscr{V}, \mathscr{C})$, with a set of nodes (or vertices) $V = \{1, 2, ..., p\}$ and a set of edges $\mathscr{C} \subseteq \mathscr{V} \times \mathscr{V}$ connecting the nodes. In case of an undirected graph, the condition $(i, j) \in \mathscr{C}$ iff $(j, i) \in \mathscr{C}$ holds. Typically, graph representations rely on an adjacency matrix $W \in \mathbb{R}^{|\mathscr{V}| \times |\mathscr{V}|}$, where each element of W is an edge weight, i.e., a connection strength between a pair of nodes. The normalized graph Laplacian is defined as $L = I - D^{-1/2}WD^{-1/2}$, where $D = diag_i [\sum_{j \neq i} w_{ij}]$ is The diagonal degree matrix. In a brain connectivity network,

the nodes typically represent either individual voxels or supersets of voxels representing anatomically meaningful parcellations of the brain. Edges, in a brain network, are indicative of connection strengths. These connections may be *structural* (e.g. brain networks based on DTI) or *functional* (e.g. brain networks based on PET or fMRI).

We have optimized a GCNN paradigm adapted to brain networks [7] for application to AD datasets. In this construction of CNNs, the adjacency matrix is utilized to create eight edge-to-edge layers, one edge-to-node, and one node-to-graph layer (in that order). It has been shown that an edge-to-edge filter over the graph \mathcal{G} is equivalent to a filter over a line graph \mathcal{G} with a *k*-hop factor of 1, which in turn, may be expressed as a generalized convolution [8]. The to-node filter creates a unitary output at each node by accumulating the convolutions of the adjacency matrix with one-dimensional spatial row and column filters. The graph-to-node filter is a one-dimensional spatial filter that furnishes a single response from all nodes.

2.2. GCNN Implementation

We implemented a GCNN multi-class classifier using the Caffe deep learning suite. The implementation was based on the architecture described in Fig. 1. The network contained eleven layers including nine convolutional layers (denoted Conv), two fully connected layers (the first one with 32 units and the second one with 10 units), each followed by a rectified linear unit (denoted ReLU), and a softmax output layer computing class-membership probabilities.

2.3. Data Description

DTI data for this study was derived from the Alzheimer's Disease Neuroimaging Initiative (ADNI) [9], a public repository containing images and clinical data from 2000+ human datasets. The primary labeled 4-class (CN, EMCI, LMCI, and AD) datasets used used for training and validation were obtained from ADNI. Subject demographics are provided in Table 1.

2.4. Data Augmentation

We utilize the adaptive synthetic sampling (ADASYN) approach [10]. This approach extends the synthetic minority oversampling technique (SMOTE) [11] by adaptively generating minority samples that are harder to learn than samples that are easier to learn. All 4 data classes (12 original datasets per class) were augmented to a size of 132 datasets per class. Of these, 100 were used for training, and the remaining were used for validation.

2.5. Data Processing

Structural networks were computed from a series of diffusion MR images with anatomical reference information obtained from corresponding Tl-weighted MR scans. A sample diffusion image and the corresponding Tl-weighted MR are shown in Fig. 2A and B respectively. White matter fiber tracts were reconstructed from the diffusion MR images via tractography using the software MedlNRIA. Prior to tractography, distortions and possible misregistrations of diffusion MR image due to eddy currents and/or motion are corrected via the FSL toolbox. The final step of the pipeline is the computation of adjacency matrices via

fiber counting. To do so, first, the Tl-weighted anatomical reference images were segmented using an atlas after deformable registration to match the FreeSurfer anatomical template shown in Fig. 2C. This step generated a segmented image volume based on the Freesurfer Desikan-Killiany atlas (Fig. 2D) containing 112 parcellated anatomical regions-of-interest (ROIs). The segmented T1 MR was then registered to the diffusion images. Fiber counting was performed on the segmented diffusion image volumes to derive pairwise inter-region connection strengths thereby yielding an adjacency matrix. A series of 112×112 adjacency matrices were computed. Fig. 2E shows sample fiber maps generated by tractography, and a sample adjacency matrix is shown in Fig. 2F.

3. RESULTS

3.1. Evaluation Metric and Reference Approach

To evaluate the classifier performance, we resorted to the receiver operating characteristic (ROC) curve, which is a graphical plot of the true positive rate (TPR = sensitivity) vs. the false positive rate (FPR = 1 - specificity) for a classifier as the discrimination threshold is varied. ROCs are commonly used for binary classifiers. For our multi-class classifier, we computed separate ROCs for each class which reflect the overall accuracy for a given class relative to all other classes. As a reference approach to determine relative accuracy of classification, we rely on a multi-class support vector machine (SVM) classifier. SVM performs classification by determining the hyperplane that best separates two classes after nonlinearly mapping the data to a high-dimensional projection space. Unlike representation learning techniques, SVMs are reliant on a predefined set of input features. Here node degrees and clustering coefficients computed from the graphs were used as SVM inputs.

3.2. Classification Performance

The GCNN was trained and tested using the Caffe platform on a Thinkmate VSX R5 540V4 workstation with an NVIDIA 1080 Ti 11 GB graphics card. During training, an L^2 loss function was minimized using the stochastic gradient descent (SGD) algorithm. The algorithm parameters were set to the following: learning rate 0.0003, batch size 10, learning momentum 0.9, weight decay 0.0001, and 200 training epochs. The final ROC curves comparing the GCNN classifier and the SVM classifier are shown in Fig. 3. The two classifiers showed near-perfect performance for the CN category with ample data samples. For the EMCI, LMCI, and AD categories, GCNN greatly outperformed SVM. The performance gap between the two methods increased with increase in levels of disease progression from CN to AD. For robust assessment, we adopted a 10-fold cross-validation strategy. Average accuracy, maximum accuracy, and standard deviation of accuracy across runs are reported in Table 2.

4. CONCLUSION

We have implemented and validated a multi-class GCNN classifier for classification of subjects on the AD spectrum. We trained and validated the network using structural connectivity graphs based on DTI. ROC analysis results show that the GCNN classifier outperforms SVM by margins that are reliant on disease category. Our findings show that the

performance gap between the two methods increases with disease progression from CN to AD. Our result is important from a technical standpoint because it elucidates the potential for the GCNN classifier to yield high performance under low-sample-size settings. The initial implementation relies on simple L^2 loss function. As future work, we will extend this classifier to larger datasets, test alternative loss functions, and test various GCNN architectures and implementations.

Acknowledgments

This work was supported in part by the NIH grants K01AG050711, RF1AG052653, and P41EB022544.

5. REFERENCES

- Yang Jiarui, Hu Chenhui, Guo Ning, Dutta Joyita, Vaina Lucia M., Johnson Keith A., Sepulcre Jorge, El Fakhri Georges, and Li Quanzheng, "Partial volume correction for PET quantification and its impact on brain network in Alzheimer's disease", Sci. Rep, vol. 7, pp. 13035, 2017. [PubMed: 29026139]
- [2]. Prasad Gautam, Joshi Shantanu H, Nir Talia M, Toga Arthur W, Thompson Paul M, Alzheimer's Disease Neuroimaging Initiative (ADNI), et al., "Brain connectivity and novel network measures for Alzheimer's disease classification", Neurobiol. Aging, vol.36, pp.S121–S131, 2015. [PubMed: 25264345]
- [3]. Sporns Olaf, "The human connectome: a complex network: The human connectome", Annals of the New York Academy of Sciences, vol.1224, no. 1, pp. 109–125, 4 2011. [PubMed: 21251014]
- [4]. Defferrard Michael, Bresson Xavier, and dergheynst Pierre Van, "Convolutional neural networks on graphs with fast localized spectral filtering", in Advances in Neural Information Processing Systems, 2016, pp. 3844–3852.
- [5]. Thomas N Kipf and Max Welling, "Semi-supervised classification with graph convolutional networks", in International Conference on Learning Representations, 2017.
- [6]. Levie Ron, Monti Federico, Bresson Xavier, and Bronstein Michael M., "CayleyNets: Graph convolutional neural networks with complex rational spectral filters", arXiv:1705.07664 [cs], 5 2017, arXiv: 1705.07664.
- [7]. Kawahara Jeremy, Brown Colin J., Miller Steven P., Booth Brian G., Chau Vann, Grunau Ruth E., Zwicker Jill G., and Hamarneh Ghas-san, "BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment", Neuroimage, vol. 146, pp. 1038–1049, 2 2017. [PubMed: 27693612]
- [8]. Shuman DI, Narang SK, Frossard P, Ortega A, and Vandergheynst P, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains", IEEE Signal Process. Mag, vol. 30, no. 3, pp. 83–98, 5 2013.
- [9]. "Alzheimer's Disease Neuroimaging Initiative (ADNI)", http://adni.loni.usc.edu.
- [10]. He Haibo, Bai Yang, Garcia Edwardo A, and Li Shutao, "ADASYN: Adaptive synthetic sampling approach for imbalanced learning", in IJCNN IEEE, 2008, pp. 1322–1328.
- [11]. Chawla Nitesh V, Bowyer Kevin W, Hall Lawrence O, and Kegelmeyer W Philip, "SMOTE: Synthetic minority oversampling technique", J. Artif. Intell. Res, vol. 16, pp. 321–357, 2002.



Fig. 1.

Architecture of the proposed graph convolutional neural network consisting of several convolutional layers (Conv) alternating with rectified linear units (ReLu), one or more fully-connected layers (FullyConn), and a softmax output layer.



Fig. 2.

Sample brain images (transverse views): (A) High-resolution Tl-weighted anatomical MR scan, (B) Diffusion MR scan, (C) Anatomical template in native (FreeSurfer) space, (D) FreeSurfer atlas with 112 anatomical regions, (E) Fiber bundles reconstructed via tractography, and (F) Adjacency matrix computed via fiber counting.





ROC curves comparing GCNN and SVM performance for each of four classes: CN, EMCI, LMCI, and AD.

Subject demographics

Class	п	Age (y): mean/s.d.	Sex: M/F
CN	12	77.4/7.8	2/10
EMCI	12	78.8/6.8	6/6
LMCI	12	72.4/6.4	5/7
AD	12	72.0/10.2	6/6

Table 2.

Performance comparison

Classifier	Average accuracy	Maximum accuracy	Standard deviation
GCNN	0.89	1	0.06
SVM	0.65	0.94	0.13