Deep Graph Convolutional Neural Networks Identify Frontoparietal Control and Default Mode Network Contributions to Mental Imagery

Michael M Craig (mmc57@cam.ac.uk) Division of Anaesthesia & Department of Clinical Neuroscience University of Cambridge

Ram Adapa (ra342@cam.ac.uk) Division of Anaesthesia University of Cambridge

Ioannis Pappas (ip322@cam.ac.uk) Division of Anaesthesia & Department of Clinical Neuroscience University of Cambridge

David K Menon (dkm13@cam.ac.uk) Division of Anaesthesia University of Cambridge

Emmanuel A Stamatakis (eas46@cam.ac.uk) Division of Anaesthesia & Department of Clinical Neuroscience University of Cambridge

Abstract:

Brain network connectivity has been characterized during a variety of tasks and neurological/psychiatric diseases. Recently, various machine learning methods, including artificial neural networks, have used connectivity measures to predict cognitive or disease state. One area where these methods could be useful is in the prognosis of patients with disorders of consciousness (DOC). Previous work has used mental imagery tasks to assess DOC patient volitional ability, however no work has focused on incorporating machine learning methods to automatically detect awareness in these patients. The present study aims to establish a baseline for these methods in classifying mental imagery states. We developed a graph convolutional network classifier that can distinguish between mental imagery states in healthy subjects using only functional connectivity data. Furthermore, we examined whether certain large scale brain networks were more predictive than others, and found that frontoparietal control and default mode networks were most predictive of whether a participant was performing a mental imagery task or resting. These results demonstrate that graph convolutional networks could be developed to aid in detection of awareness in DOC patients and show that changes in connectivity patterns in frontoparietal control and default mode networks underlie alterations in mental imagery.

Keywords: Frontoparietal Control Network, Default Mode Network, Mental Imagery, Graph Convolutional Networks

Introduction

Changes in brain network connectivity using functional MRI have been characterized during a wide variety of disease states and behavioural paradigms (Bassett and Bullmore, 2009; Bullmore and Sporns, 2009). The past few years have seen increasing use of graph theoretical measures as a way to model these changes both during tasks and during resting state. Most recently, graph convolutional neural networks have been used to diagnose various neurological/psychiatric disorders using functional brain connectivity (Ira Ktena et al., 2017). Concurrently, mental imagery paradigms have been used to aid in the diagnosis and prognosis of patients with disorders of consciousness (DOC; Owen et al., 2006; Monti et al., 2010).

The aim of the present study is to merge these two lines of research by applying deep graph convolutional neural networks (DGCNN) to classify brain network connectivity data of healthy participants while they perform a mental imagery task used in the assessment of patients with DOC. To achieve this we used the full brain connectivity graphs as well as individual canonical large scale brain networks as inputs to the network. This work will establish a baseline for DGCNN performance, with the aim of applying this method to patient groups in the future.

Methods

Participants and Paradigm Specifications

Data was collected at the Wolfson Brain Imaging Centre at Addenbrooke's Hospital, Cambridge, UK. Ethics for this study were obtained from the Cambridgeshire 2 Regional Ethics committee. 25 healthy participants were recruited for the study (9 female; mean age = 35; range = 19-35 years). 3 participants were removed due to excessive head motion during scanning.

The behavioural paradigm consisted of five 30 second blocks where participants alternated between resting state and a motor imagery paradigm. In the motor imagery paradigm participants were instructed to imagine being on a tennis court and swinging their arm to hit a tennis ball to an instructor on the opposite side of the net (Monti et al., 2010). The beginning of the motor imagery paradigm was cued with the word "Tennis" appearing on the screen, and the beginning of the resting state paradigm began with the word "Rest" appearing on the screen.

Data Acquisition and Preprocessing

Data was acquired using a 3T Tim Trio Siemens system (Erlangen, Germany). We used a 12-channel head matrix transmit-receive coil with a fast-sparse 32 slice axial oblique sequence (TR=2000ms, TE=30ms, flip angle=78°, voxel size = 3.0 x 3.0 x 3.0 mm^3, matrix size 64 x 64, field of view 192mm x 192mm, slice thickness = 3.0mm). We collected 150 EPI images in each subject's run. T-1 weighted MPRAGE high-resolution structural images were also acquired with 1.0 x 1.0 x 1.0 mm^3 resolution (TR = 2250ms, TI = 900ms, TE= 2.99ms, flip angle = 9°).

Preprocessing was performed with Statistical Parametric Mapping 12 (SPM12; http://www.fil.ion.ucl.ac.uk/spm/) and MATLAB version 2017a (http://www.mathworks.co.uk). Following motion correction in the functional dataset, the participant's high-resolution structural images were coregistered to the mean EPI and segmented into grey matter, white matter and cerebrospinal fluid masks. Next the images were normalised to Montreal Neurological Institute (MNI) space resampled to a resolution of 2 mm^3. Functional images were smoothed with a 6mm FWHM Gaussian kernel. To reduce residual movement-related and physiological artifacts, data underwent de-spiking with a hyperbolic tangent squashing function. Next the aCompCor technique was used to remove the first 5 principal components of the signal from the white matter and cerebrospinal fluid masks, as well as 6 motion parameters and their first order temporal derivatives and a linear de-trending term. Functional images were then highpass filtered to remove low frequency fluctuations associated with scanner noise (0.009 Hz < f).

Because deep learning methods need a large amount of training data, we augmented our data using dynamic functional connectivity. Each condition contained 75 volumes per run. We used 50 volume-sliding windows across each condition to generate 25 correlation matrices per condition, per subject.

Functional connectivity was calculated by computing the correlation coefficient between time series from each of the 118 cortical regions from the Lausanne parcellation. Each correlation matrix was binarized, keeping the top 20% of connections. We used the entire 118 region graph, as well as 6 large scale brain networks as input to the DGCNN. This allowed us to determine whether a specific large scale network is most important for classifying cognitive state. These networks included the auditory (Aud), default mode network (DMN), frontoparietal control network (FPCN), salience network (SN), somatomotor network (SM) and visual network (Vis). Node assignment to each network was calculated by overlapping each ROI with a network mask from Smith et al., 2009.

Deep Graph Convolutional Neural Network

The DGCNN is adapted from Zhang et al., 2018 (https://github.com/muhanzhang/DGCNN) and consists of three sequential stages. 1) Graph convolutional layers to extract node connectivity features; 2) a SortPooling layer to sort node features and equate input features size for, 3) a series of classical convolutional and fully connected neural network layers to read the sorted graph representations and make predictions (Zhang et al., 2018).

The graph convolutional layers aggregate node information from neighboring nodes to extract multiscale graph substructures important for classification. Input to each graph convolutional layer includes $\mathbf{A}$, an adjacency matrix, $\mathbf{D}$, a diagonal degree matrix, $\mathbf{X}$, a node information matrix with nodes in the rows and c node features in the columns, and $\mathbf{W}$, a matrix of trainable parameters. Each layer contains four separate steps. First a linear transformation is applied to the node information matrix $\mathbf{XW}$. This maps the c feature channels to c' features in the next layer. The second step propagates node information to neighboring nodes. Step three normalizes each node's feature vector. The fourth step applies a nonlinear activation function and outputs the graph convolution.

The SortPooling layer sorts the output of the graph convolutional layer so that node feature vectors are pooled together and outputted in a consistent order. This is important because the final 1D convolutional step is most effective at classification when features are presented in a consistent order. This final step consists of two layers of classical 1D convolutional layers, each with a convolutional layer, followed by a rectified linear unit activation function and maxpooling layer. This is followed by a fully connected layer and a softmax layer for classification. Each classification analysis was trained for 200 epochs with a learning rate of 0.0001. Of the 25 subjects in the dataset, 22 were used in the final analysis, due to excessive movement during scanning. The complete dynamic functional connectivity data from 4 subjects were randomly allocated to a test set so as to not overfit the classifier during training.

Results

We used a DGCNN to determine whether a specific large scale brain network was most predictive of cognitive state in healthy participants performing a mental imagery task. To assess the results of each binary classificaiton we used the following metrics: Precision (True positives/True Positives + False Positives), recall (True Positives/True Positives + False Negatives), F1 Score (Harmonic mean of Precision and Recall) and the area under the receiver operating characteristic curve (ROC-AUC). The ROC curve is a plot of True Positive Rate against False Positive Rate for different cutoffs of a diagnostic test, and is a measure of tradeoff between sensitivity (true positive rate) and specificity (1 – false positive rate). As our analysis has balanced classes (i.e. an equal number of examples for each cognitive state) the ROC-AUC was considered the most important metric (David & Goadrich, 2006). A comprehensive list of results can be found in Table 1.

	Precision	**Recall**	**F1**	**ROC AUC**
Auditory	0.64	0.635	0.637	0.667
Somatomotor	0.54	0.525	0.532	0.487
Visual	0.51	0.525	0.517	0.475
Salience	0.54	0.545	0.542	0.562
DMN	0.7	0.7	0.700	0.793
FPCN	0.81	0.805	0.807	0.885
FPCN-DMN	0.7	0.7	0.7	0.75
Full Network	0.75	0.76	0.755	0.595

Table 1: Classification results for held out test data for each network input.

We first examined whether we could accurately classify between Tennis Imagery and Resting states using a full cortical graph. This analysis resulted in a ROC-AUC of 0.595 on held out test data. We then used each of the large scale networks as input, and found that FPCN (ROC-AUC = 0.885) and DMN (ROC-AUC = 0.793) had much higher ROC-AUC scores than the other networks, suggesting they alter their connectivity pattern the most between the mental imagery and resting state conditions (Figure 1). We also combined these two networks, including both within and between network connections as input. This however resulted in a slightly poorer classification with a ROC-AUC of 0.75.

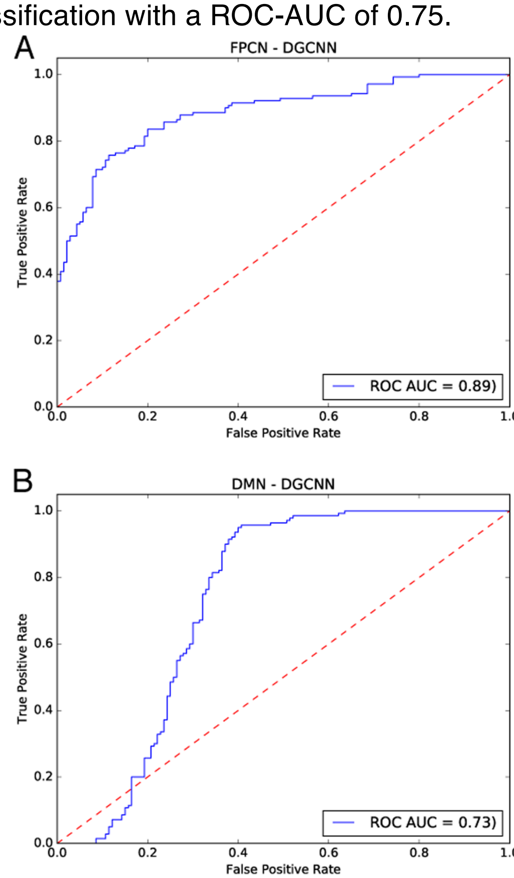

Figure 1:ROC Curve for **A** FPCN and **B** DMN.

Discussion

The present study used functional connectivity fMRI, a mental imagery task and a DGCNN to determine whether specific large scale brain networks were more predictive of mental imagery state.

We first found that the full cortical graph could predict mental imagery state with modest accuracy. Interestingly, the full graph had a relatively high precision and recall, but a relatively low ROC-AUC. This is likely due to the fact that ROC-AUC is calculated using prediction probabilities (values ranging between 0 and 1) and not binary classification measure (either 0 or 1). Essentially, this means that the full graph classification was often correct, however it was less confident in its decision.

We also examined the predictive capacity of several large scale cortical brain networks. We found that FPCN and DMN were most predictive of whether a participant was performing the Tennis mental imagery task or resting. Previous work has shown that these networks dynamically interact during mental imagery tasks involving future planning (Gerlach et al., 2014). We also found that using both within and between DMN and FPCN connections as input features did not improve the ROC-AUC, suggesting that within network connections for each of these networks are the most important for classification. Notably, a previous study by Spreng et al. (2012) showed that the typically anticorrelated DMN and SN are modulated by FPCN, which flexibly couples to either network depending on whether the task involves internally or externally directed attention. Our findings are therefore in line with these results in that they show FPCN and DMN are most predictive of a mental imagery task involving the switching between two states where attention is directed internally.

The present study has focused only on healthy participants, however mental imagery tasks similar to the one used here have been used in the diagnosis and prognosis of patients with DOC (Owen et al., 2007; Monti et al., 2010). This work provides a method that could potentially be used to aid in the automated detection of awareness in these patients. Future work in this area could use more sensitive mental imagery tasks like movie-watching and incorporate more extensive preprocessing steps and deeper neural network architectures that could overcome the complexities of working with data from patients with brain damage.

References

Bassett, D.S., & Bullmore, E.T. (2009). Human Brain Networks in Health and Disease. *Curr Opin Neurol* 22(4): 340-7

Bullmore, E.T., & Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*. 10(3): 186-98

Behzadi, Y., Reston, K., Liau, J., Liu, TT. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage*. 37(1): 90-101

Davis, J., Goadrich, M. (2006). The relationship between Precision-Recall and ROC curves. *International Conference on Machine Learning*. 244-40

Gerlach, K., Spreng, R.N., Madore, K.P., Schacter, D.L. (2014). Future planning: default network activity couples with frontoparietal control network and reward-processing regions during process and outcome simulations. *Social Cognitive and Affective Neuroscience*. 12(1): 1942-51

Ira Ktena, S., Parisot, S., Ferrante, E., Rajchl, M., Lee, M., Glocker, B., Rueckert. (2017). Distance Metric Learning using Graph Convolutional Networks: Application to Functional Brain Networks. *International Conference on Medical Image Computing and Computer-Assisted Intervention.* 469-477

Monti, M.M., Vanhaudenhuyse, A., Coleman, M.R., Boly, M., Pickard, J.D., Tshibanda, L., Owen, A.M., Laureys, S. (2010) Willful Modulation of Brain Activity in Disorders of Consciousness. *New England Journal of Medicine*. 362: 579-589

Owen A.M., Coleman, M.R., Boly, M., Davis, M.H., Laureys, S., Pickard, J.D. (2006). Detecting Awareness in the Vegetative State. *Science*. 313: 1402

Spreng, R.N., Sepulcre, J., Turner, G.R., Stevens, W.D., Schacter, D.L. (2013). Intrinsic architecture underlying the relations among the default, dorsal attention, and frontoparietal control networks of the human brain. *Journal of Cognitive Neuroscience*. 25(1): 10.1162

Zhang, M., Zicheng, C., Neumann, M., Chen, Y. (2018). An End-to-End Deep Learning Architecture for Graph Classification. *Association for the Advancement of Artificial Intelligenc*e

Acknowledgements

The funding for this study was provided by the Wellcome Trust.