A Machine Learning Case for a Higher Order Control Plexus in the Frontal Pole Cortex

Nishant Zachariah[1], Zhihao Li[2,3], Jason Langley[2], Shiyang Chen[2], Mark Davenport[1], Justin Romberg[1], and Xiaoping Hu[2]

[1]Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, United States, [2]Department of Biomedical Engineering, Emory University and Georgia Institute of Technology, Atlanta, GA, United States, [3]Institute of Affective and Social Neuroscience, Shenzhen University, Shenzhen, Guangdong, China

Introduction: In this paper we examine fMRI data in order to determine if a higher order cognitive plexus exists within the human brain. We define a higher order cognitive plexus as a functional group of less than 4 voxels whose time course activity can be used to accurately (≥ 90%) distinguish between several higher order cognitive tasks. Our findings demonstrate that cognitive plexuses of 3-4 voxels do exist for every subject that was examined in this study. Remarkably, we are able to use a linear classifier with only 3-4 voxel time courses (out of 902,629 voxels) to accurately distinguish between 7 higher order cognitive + resting state with (≳ 90%) accuracy. Furthermore, when examined across a population, we found that the most common voxel location for these control plexuses was in the frontal pole cortex. Given that the frontal pole cortex is one of the least understood regions of the frontal lobe, this suggests a novel functional role that the frontal pole cortex plays in modulating higher order cognitive tasks.

Methods: *(i) Data*: We use fMRI data from the HCP (Human Connectome Project) because it has the highest spatial and temporal resolution of all publicly available datasets at present [1]. For each of the 10 subjects examined in this study, time series for 902,629 voxels were as part of the HCP protocol. The duration of the time series varied between cognitive tasks with a combined 3141 time points across all tasks. Resting state activity along with 7 higher order cognitive task activity was collected for each patient. Task based activity included paradigms which tested working memory, relational processing, social cognition, language processing, risk evaluation, cognition in emotion evoking tasks, and tasks which required processing of motor cues. The final feature matrix (voxels as columns and time series as rows) for each subject had more than 2.8343 trillion entries when all the task and resting state data was pooled for analysis. We randomly partition the time series for each subject into a 90% training and 10% testing set (hold out set) to prevent over-fitting and ensure generalizability of our results. We did not apply any form of additional pre-preprocessing to the post MNI atlas registered data as provided by the Human Connectome Project.

(ii) Feature (voxel) selection: Given the enormous size of the data set, traditional forward feature selection methods in which each feature is picked (one at a time) in order to maximize performance, is computationally intractable. Thus, in order to reduce the size of the feature (voxel) space that we have to examine for forward feature selection, we first extract voxels that are considered statistically significant across a broad range of metrics (Battacharya test, t-test, Entropy ranking and the Wilcoxon Mann Whitney test). This range of statistical metrics are used to overcome the assumptions that would otherwise be a limitation of a single method (Ex. data normality assumption for t-test). The extracted voxel list forms a much smaller subset (~40,000 voxels) which then allows for the application of traditional forward feature selection methods for voxel ranking based on prediction accuracy.

(iii) Classifier: We intentionally use a linear classifier with no alteration to the raw features (voxel time series) in order to ensure interpretability of the final result. A linear discriminant analysis (LDA) classifier was chosen for this task as the decision boundaries of an LDA classifier are a function of both the subspace spanned by the features (voxel time series) and the covariance between features (voxel time series). Other linear classifiers such as the support vector machine and logistic regression do not consider correlation between features (voxels) in the computation of a decision boundary. As such an LDA classifier is an ideal candidate for detection of a control plexus as it considers both the time course of a single voxel and that voxel's interactions with the additional voxel time series being considered. This form of correlation between voxels is integral to the model as it is what we would expect from a control plexus voxel which regulates the activity of spatially variant brain regions.

Results: (i) *Proof of Concept*: We begin by applying the feature selection method to each subject. The resulting classifier accuracy is plotted in Fig. 1(A). While the perfect prediction accuracy of an LDA classifier with 2-3 voxels is remarkable, the results are biased as a result of over-fitting (given that the test set was used as part of the forward feature selection method). Thus Fig. 1(A) while strongly indicative that higher order cognitive control plexuses do exist, it is not conclusive.

(ii) Unbiased Top Voxel Based Prediction: In Fig. 1(B) we extract relevant voxels across the entire population. Once the relevant subset of voxels was extracted, forward selection was used to generate a ranked voxel list for each subject. We then train and test classifier performance for all subjects except the subject that was used to generate the voxel ranking. By doing so, we eliminate any form of selection bias / test set contamination. The plot seen in Fig. 1(B) demonstrates that even with a single voxel's time series, we are able to perform almost 4 times better than random chance in predicting the exact type of higher order cognitive task. With just 4 of the top voxels, we are able to predict with ≳90% accuracy between all 7 cognitive tasks + resting state.

(iii) Location of Top Voxels: Fig. 1(B) makes a strong case for the existence of a control plexus of 4 voxels in each subject that was examined as part of this study. However, since the location of the top voxel differ from subject to subject, we need to examine the location of these voxels across patients to examine if a population trend exists. When the location of the top 4 voxels are examined on the MNI atlas, the region with the largest number of top voxels across all subjects is the frontal pole cortex.

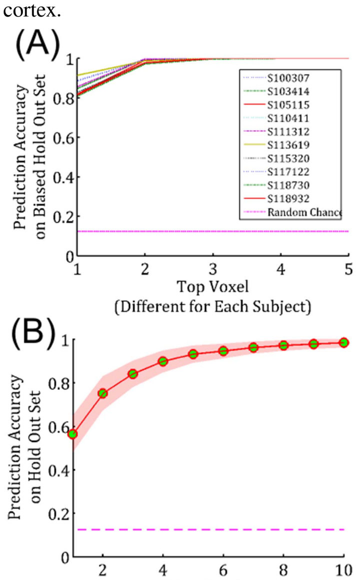

Discussion and Conclusion: In this study, we demonstrate that when using a purely linear classifier, with as few as the time series of 4 voxels, we are able to accurately (≥ 90%) distinguish between 7 higher order cognitive tasks along with resting state activity. Furthermore, when we examine the location of these top 4 voxels across all subjects, we find that the most common location of one of the top 4 voxels is in the frontal pole cortex (FPC). The FPC has generated much interest given its stark distinction in size and volume between humans and lower order primates. Recent single cell electrode studies in the FPC of monkeys, suggest that the region is involved in modulating activity based on feedback [2, 3]. Other studies suggest that the FPC in humans plays a role in modulating behavior for improved learning outcomes for goal based activity [3]. The results of our study demonstrate a novel and pivotal role that the frontal pole cortex plays in higher order cognitive tasks that is independent of external feedback. Furthermore, our study implies that all resting and task based connectivity profiles must include frontal pole activity given the modulatory nature of the FPC that we have uncovered here. Finally, this work demonstrates a new methodological approach for discovering the location of pivotal control plexuses within the human brain.

References: [1] Van Essen, David C., et al., Neuroimage, Vol 80, pg: 62-79, 2013. [2] Tsujimoto, Satoshi et al, Nature Neuroscience, Vol 13 (1), pg. 120-126, 2010. [3] Tsujimoto, Satoshi, et al, Trends in Cognitive Sciences, Vol 15 (4), pg. 169-176, 2011.

Figure 1: (A) Voxel selection (forward feature selection) applied to each subject on a hold out data set. Different subjects are plotted in different colors. The broken line represents performance based on random chance. Since the performance is plotted for the same hold out set on which the voxels (features) are picked, this is a classical case of over fitting. However, this plot suggests that each subject has a small group of voxels can play a pivotal role in higher order cognitive function (B) Prediction accuracy without test set bias as shown in (B). The median accuracy across all subjects with bound of +/- 1 standard deviation are plotted. The broken line represents the performance of a classifier which randomly assigns labels. With as few as 4 top voxels, the LDA classifier is able to predict with ≳90% accuracy. This makes a strong case for the existence of a control plexus of 3-4 voxels for each subject