On functional annotation of the human brain by combining resting state connectivity and activation foci

Ching-Yi WU1 , Soumya RAY1 , Kevin P. BARRY2 , Anthony. I. JACK2

1 Department of Electrical Engineering & Computer Science, Case Western Reserve University
2 Department of Cognitive Science, Case Western Reserve University

Introduction

How can we identify and label functional circuits in the human brain? There are at least two extant data sources which can inform this task. If two brain areas are implicated in the same functional comparison, e.g. the same contrast between tasks, then this provides evidence that they share functionality. A second source of evidence that two areas form part of the same functional circuit derives from resting state connectivity data (rsfMRI). Here, we address the problem of comparing and combining resting state connectivity data with activation foci data. For both of these two sources, we represent the functional relationship between two brain regions using two affinity matrices, based on a undirected weighted graph and apply graph-based clustering algorithm to investigate the following three questions.

1. ***Are there subsets of regions of the brain that are reasonably similar, according to the affinity matrices?***

2. ***Does clustering based on these matrices generate similar results?***

3. ***How can we effectively combine these two sources of information to derive a functional annotation of the human brain?***

Methods

Data

We obtained activation foci (n = 26616) from the SUMS database (http://sumsdb.wustl.edu) and resting state connectivity data for 500 individuals downloaded from the NITRC 1000 functional connectomes project (http://www.nitrc.org/docman/?group_id=296). We divided the brain into regions as follows: To establish cortical regions, we first identified their locations ('centers') by transforming a regularly spaced grid (1x1cm spacing) on the PALS (Van Essen, 2005) left and right flat surfaces into volume space. To establish subcortical and cerebellar region centers, the grey-matter volume not covered by the PALS atlas was divided using a regularly-spaced grid (9x9x9mm spacing). This defined 2153 centers that were relatively evenly spaced throughout areas of grey matter. For the resting state data analyses, each region was defined by a sphere (r=6mm) surrounding its respective center. Voxels falling within two such spheres were assigned to the closest center and the region of the other center was truncated, preventing any overlap. For the coactivation data, foci falling with 9.8 mm of a center were mapped to the closest center. We only consider those centers containing two or more foci for further analysis. This results in 2000 regions.

Analysis

A. Construction of Affinity Matrices:

1. Co-activation affinity matrix (CA): We create an "affinity matrix" between the 2000 regions by defining an edge between two regions if and only if each region contains foci from the same statistical comparison in the same study. The weight of this edge is computed according to the formula: (total number of edges between foci in region *i* and region *j*)/(minimum number of foci in region *i* and region *j*)
2. rsfMRI affinity matrix (RS): Using the same 2000 regions, we construct another affinity matrix from the rsfMRI data, by taking the correlation matrix created using the established methods of Fox et al. (2005), divided by two plus one (so creating a matrix of values whose theoretical limits are 0 and 1).

B. Three analyses based on the two affinity matrices, to answer questions above

1. *Find a subset of the 2000 regions which are similar*: We calculate a standard correlation coefficient across different regions in order to discover a subset of regions between CA & RS which are somewhat similar.

2. *Compare clustering solutions on the two matrices:* We apply spectral clustering (Shi and Malik, 2000) to the two affinity matrices. This is a graph clustering method that makes no assumptions about the generative model of the data. The algorithm interprets an affinity matrix as the adjacency matrix of a graph. It partitions the graph so that the ratio of the total weight of edges within a partition divided by the weights on those leaving the partition, summed over all partitions, is maximized. This approach has recently been applied to rsfMRI data (van den Heuvel, 2008 and Venkataraman, 2009). Note that among 2000 regions, we left out 99 regions since they consist of either too few or too many foci, resulting in regions with too small or too large edge weights. These act as outliers for the algorithm, creating either very large or very small clusters. This gives us1901 regions in total for our analysis.

3. Combine two affinity matrices (CA, RS) We then use two approaches to combine the two affinity matrices and generate a single affinity matrix. Then we run spectral clustering in order to find the unified clustering solution on functional brain networks. The two combination methods we use are:
1)We treat each (normalized) affinity matrix as the probability matrix of a Markov random field. A combined affinity matrix is generated by treating the individual matrices as independent sources of information, and combining them via a product.
2)We treat parts of the RS solution as a partial constraint on the CA matrix solution. This is motivated using the observation that the RS solution is able to identify important functional regions such as visual and motor regions very well, and so can be used to "fill in" gaps in the CA solution.

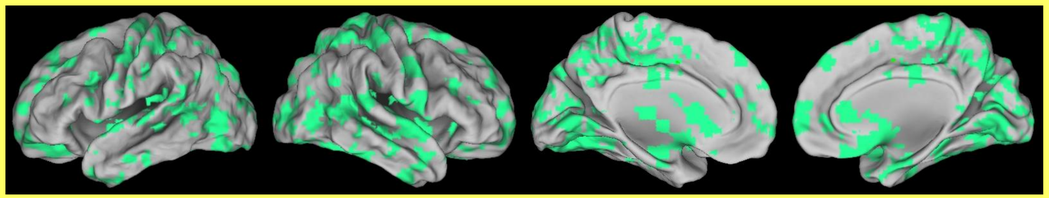

Fig. 1. Out of 2000 regions, 614 regions with weakly similar affinities between co-activation affinity matrix and rsfMRI affinity matrix

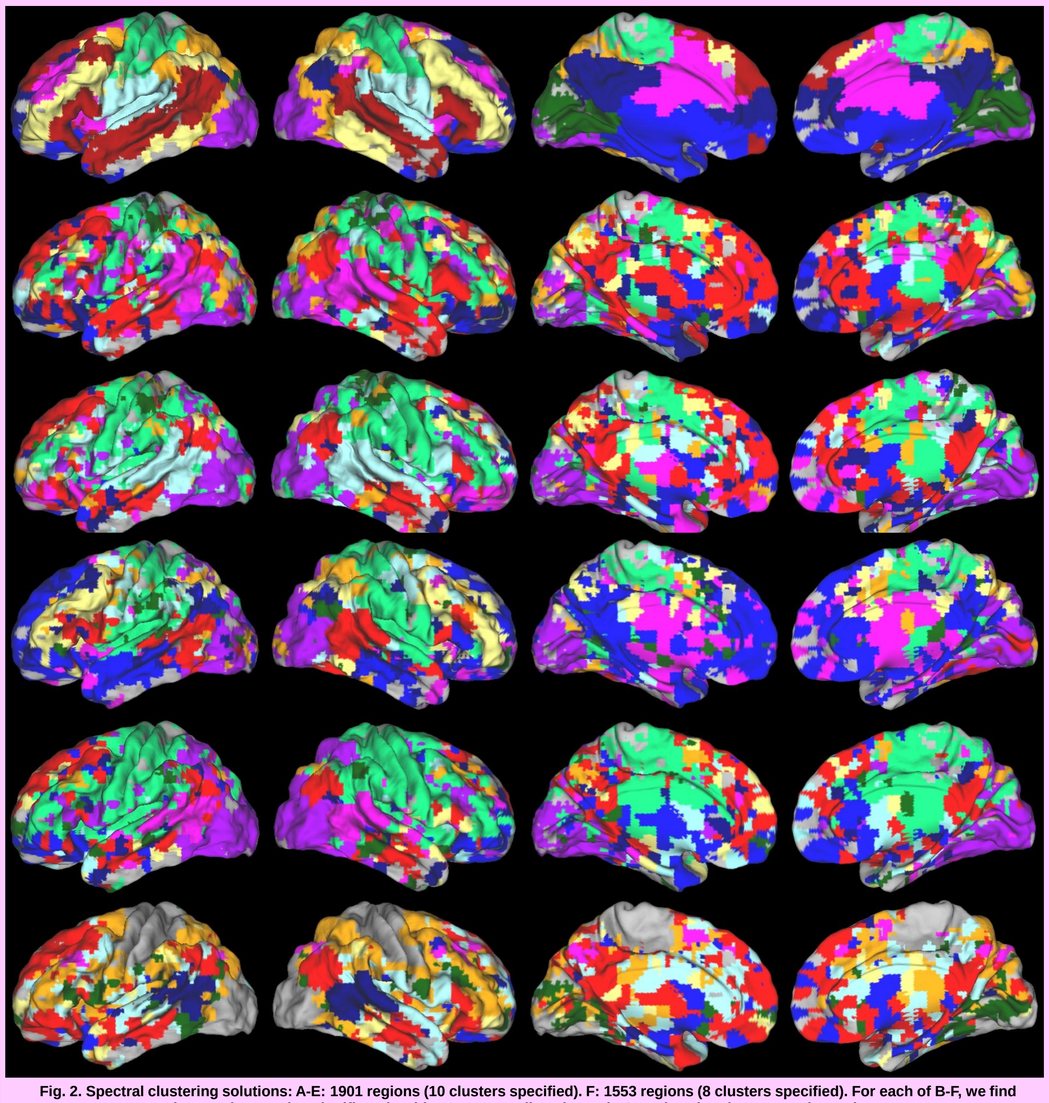

Fig. 2. Spectral clustering solutions: A-E: 1901 regions (10 clusters specified). F: 1553 regions (8 clusters specified). For each of B-F, we find clusters that overlap significantly with a corresponding cluster in A, and assign the same color to these.

Result : Many Dissimilar Regions in Resting State and Co-activation Matrices

Among 2000 regions, we found 614 regions between CA and RS with weakly similar affinities ($p < 0.3$), and only 92 regions with strongly similar affinities ($p < 0.05$). Fig. 1 shows 614 weakly similar regions (green regions) between CA and RS. This is different than the result outlined in Smith et al. (2009), where it was observed that RS and CA data overlapped significantly in terms of indicating functionality. In our results, although the two sources of information (co-activation and rsfMRI) share some common aspects, they differ significantly in many regions as well. We infer that they may behave quite differently when inferring functionality.

Result: Spectral Clustering Solutions for Co-activation and rsfMRI are Different (Fig. 2, A & B)

We run spectral clustering on CA and RS respectively (n = 1901), with the predefined number of clusters equal to 10 (this number was chosen based on Smith et al. 2009). Fig. 2 (**A** and **B**) shows the results. For the rsfMRI data, clusters reflected a high degree of local connectivity and less pronounced long-range connectivity, as compared with the co-activation data. The clustering solution derived from RS (**A**) appears generally better behaved than from the CA (**B**), which may be due to a 'larger' and more even sample. However the co-activation data provides a 'truer' measure of functional correspondence between regions. While some limited correspondence between rsfMRI and co-activation data has been found, the extent to which this method can be relied upon as a guide to function remains to be determined. As predicted from Result 1, the spectral clustering solutions on these two matrices yield quite different results. This motivates us to investigate possible approaches to integrate the two matrices in order to discover a unified solution.

Result: Spectral Clustering Solutions for a Combined Matrix (Fig. 2, C & D)

The results above indicate that the RS and CA matrices contain (at least in several sub-regions) differing information about functionality. We therefore investigate integrating the information in the two matrices. The first approach we use is based on the idea that for spectral clustering, one can interpret the (normalized) affinity matrix as the transition probability matrix of a Markov random field (von Luxburg, 2007). Thus if we treat these matrices as independent sources of information, one way to combine them is to take the product of the two (normalized) matrices. The result obtained by clustering this product affinity matrix is shown in Fig. 2 (**C**) . It is notable that motor and early visual areas form parts of highly extended clusters in this solution, an issue we address in the constrained matrices below.

The simple product described above gives equal weight to the RS and CA matrices. As we have observed, however, the RS matrix tends to be more coherent than the CA matrix we have, over many regions of the brain. Thus it seems intuitive to give the RS matrix higher weight during the integration. This can be done as follows: we create the combined matrix as the product $RS^{\lambda}CA^{(1-\lambda)}$. If $\lambda=1$, this corresponds to the pure RS matrix, if zero, the pure CA matrix. If λ is between zero and 0.5, the product is closer to the CA matrix, otherwise the product is closer to the RS matrix. Thus this allows us to choose any operating point between the RS and CA matrices, depending on the strength of our belief in the validity of their information. Fig. 2, (**D**) demonstrates the results for clustering a combined matrix generated with $\lambda=0.7$. This solution appears better behaved, and may reflect patterns of co-activation, as observed by scientists familiar with findings from multiple studies, better than either solution from RS (**A**) or CA (**B**) alone. For instance, particularly in the left hemisphere, the default network (comprising precuneus/posterior cingulate, medial prefrontal, lateral parietal and superior frontal) (Van Dijk et. Al, 2009) seems to emerge as more consistent in this solution than in other solutions. In the RS solution **(A)** the default network appears to be split into three distinct clusters.

Result: Spectral clustering solutions for a constrained matrix (Fig. 2, E & F)

An alternative method for combining the information in the two matrices is to treat one solution as a constraint on the other. Since in our data, the quality of the CA solution (**B**) is inferior to the RS solution (**A**), we use the RS solution (**A**) as a partial constraint on the CA solution (**B**). More specifically, we were concerned that the solution derived from the CA (**B**) data might be biased by pragmatic constraints which would influence how the data is sampled. Scanner studies generally employ a visual stimulus and a motor output. Co-activation data may therefore produce a solution in which the connectivity of early visual and motor cortices to other areas is exaggerated.
To examine and correct for this, we selected the early visual and motor networks from the RS solution (**A**). We then merge each of these networks into a single node. This forces the CA solution (**B**) to include all regions in each merged network in the same cluster. We investigate two variants of this procedure. In the first case, we add rows in the CA matrix corresponding to the new, merged regions and remove all of the rows corresponding to the original networks (the number of the rows of the matrix for running algorithm: 1555; the corresponding number of brain regions analyzed: 1901). In the second case, we simply remove the original networks, but do not add the merged regions (n = 1553). In the first case**,** we see that the visual and motor networks grow in size (**E**). This is consistent with the presence of the hypothesized artifact distorting the co-activation data. In the second case, we can view the CA solution with this source of distortion removed (**F**). This produces a better behaved solution, in which the default network falls largely in a single cluster, and the correlated dorsal attention and fronto-parietal control networks are seen to fall into a separate single cluster.

Future directions and outstanding issues

Based on our analysis, rsfMRI and co-activation data share certain commonalities in brain functions, but lots of regions are quite different. This brings up several interesting questions for our future investigation:

1.Why do we see different solutions between rsfMRI and co-activation data, while Smith et al see similarity? Will the overlapped regions between co-activation and rsfMRI increase if we have more activation foci?
2.What does rsfMRI tell us about brain function?
3.To what degree will use of Statistical Parametric Maps (i.e. full volume contrast images) instead of activation foci improve the solution for co-activation data?
4.What further sources of evidence (e.g. Neurotransmitter density, DTI, cytoarchitecture) might be integrated with co-activation and rsfMRI data to form a better overall solution to the problem of identifying function?

References

Van Essen D. (2005). A population-average, landmark-and surface-based (PALS) atlas of human cerebral cortex, *Neuroimage*, 28-3, 635-662.
Fox, M. D., Snyder, A. Z., Vincent, J. L., Maurizio C., Van Essen, D. C. & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *PNAS*, 102 (27), 9673-9678.
Shi, J., & Malik, J.(2000). Normalized cuts and image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 22, 888-905.
van den Heuvel M., Mandl R.& Hulshoff Pol H. (2008). Normalized Cut Group Clustering of Resting-State fMRI Data. *PLoS ONE* 3(4): 1-11.
Venkataraman A., Van Dijk K. R.A. ,Buckner R.L. & Golland P. (2009). Exploring functional connectivity in fMRI via clustering. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 441 – 444.
Smith S.M., Fox P.T., Miller K.L., Glahn D.C., Fox, P.M., Mackay C.E., Filippini N., Watkins K.E., Toro R., Laird A.R. & Beckmann C.F. (2009). Correspondence of the brain's functional architecture during activation and rest. *PNAS*, 6, 31, 13040-13045.
von Luxburg U. (2007). A tutorial on spectral clustering. *Statistics and Computing*. 17(4). 395-416.
Van Dijk K.R. A., Hedden T., Venkataraman A., Evans K.C., Lazar S.W. & Buckner R.L. (2009). Intrinsic Functional Connectivity As a Tool For Human Connectomics: Theory, Properties, and Optimization, J. of Neurophysiology. 103: 297-321