

Multi-Level Parcellation of the Cerebral Cortex Using Resting-State fMRI

Salim Arslan and Daniel Rueckert

Biomedical Image Analysis Group, Department of Computing,
Imperial College London, London, UK

Abstract. Cortical parcellation is one of the core steps for identifying the functional architecture of the human brain. Despite the increasing number of attempts at developing parcellation algorithms using resting-state fMRI, there still remain challenges to be overcome, such as generating reproducible parcellations at both single-subject and group levels, while sub-dividing the cortex into functionally homogeneous parcels. To address these challenges, we propose a three-layer parcellation framework which deploys a different clustering strategy at each layer. Initially, the cortical vertices are clustered into a relatively large number of supervertices, which constitutes a high-level abstraction of the rs-fMRI data. These supervertices are combined into a tree of hierarchical clusters to generate individual subject parcellations, which are, in turn, used to compute a groupwise parcellation in order to represent the whole population. Using data collected as part of the Human Connectome Project from 100 healthy subjects, we show that our algorithm segregates the cortex into distinctive parcels at different resolutions with high reproducibility and functional homogeneity at both single-subject and group levels, therefore can be reliably used for network analysis.

1 Introduction

Parcellation of the cerebral cortex constitutes one of the core steps to reveal the functional organization of the brain. It is usually followed by network analyses devised to generate graphical models of the connections between the parcellated regions. Such analyses have the potential to uncover the neural mechanisms behind the human behavior and to help understand neurological disorders [8]. It is of great importance to obtain reliable parcellations, since errors at this stage propagate into the subsequent analysis and consequently affect the final results. Among others, there are two notable attributes that define “reliability” in the context of a cortical parcellation: 1) parcellated sub-regions should be functionally consistent and comprise similar vertices, since network nodes are typically represented by a single entity (such as the average time series of the constituent vertices) and 2) both individual subject and groupwise parcellations should be reproducible to some extent, that is, multiple parcellations obtained from different datasets of the same subject as well as groupwise parcellations computed from the subsets of the same population should exhibit functional and structural similarity.

With this motivation, we propose a whole-cortex parcellation framework based on resting-state functional magnetic resonance imaging (rs-fMRI). The brain is still functional in the absence of external stimuli, thus rs-fMRI time series can be utilized to parcellate the cortical surface into functionally homogeneous sub-regions. The majority of the literature on the rs-fMRI driven parcellation techniques consists of methods (i) that aim to discover the cortical networks (e.g. default mode network) [10,5] and (ii) that propose to subdivide the entire cortical surface in order to produce a baseline for connectome analysis [3,2,7]. Hierarchical clustering, independent component analysis (ICA), region growing, spectral graph theory, and k -means are some of the many statistical models proposed to compute cortical parcellations as reviewed in [6]. Although promising solutions exist among them, there still remain challenges to be overcome, especially in order to obtain reliable parcellations that fulfill the aforementioned requirements at both single-subject and group levels.

We address these challenges with a three-layer parcellation framework in which each layer makes use of a clustering technique targeting a specific problem. First, we pre-parcellate the cortical vertices into highly consistent, relatively large number of homogeneous supervertices with a hybrid distance function based on rs-fMRI correlations and geodesic distance. This stage does not only reduce the dimensionality of the data and decrease the computational cost, but also improves the SNR. Second, we build hierarchical trees on top of the supervertices to obtain individual parcellations reflecting the functional organization of the cortex without losing the spatial integrity within the parcels. Third, we compute a graphical model of the parcel stability across the individual parcellations and cluster this graph in order to generate a groupwise representation of the subjects in the population. Our framework is capable of parcellating the cortical surface into varying number of sub-regions (50 to 500 per hemisphere), thus allowing the analysis at multiple scales.

The most closely related work to our approach is a single-subject parcellation method composed of region growing and hierarchical clustering [2]. The major differences to our proposed work are twofold. First, we introduce a new pre-parcellation technique based on supervertices instead of relying on regions derived from stable seeds. Second, our method is capable of generating parcellations at the group level, thus can be used for population network analysis as opposed to [2]. We also compare our approach to another state-of-the-art cortical parcellation method based on spectral clustering with normalized cuts [3] and demonstrate that the proposed framework is more effective than the other approaches at both single-subject and group levels, achieving high reproducibility while preserving the functional consistency within the parcellated sub-regions.

2 Methodology

2.1 Data Acquisition and Preprocessing

We evaluate our algorithm using data from the Human Connectome Project (HCP). We conducted our experiments on the rs-fMRI datasets, containing scans

from 100 different subjects (54 female, 46 male adults, age 22-35). The data for each subject was acquired in two sessions, divided into four runs of approximately 15 minutes each. Data was preprocessed and denoised by the HCP structural and functional minimal preprocessing pipelines [4]. The outcome of the pipeline is a standard set of cortical time series which were registered across subjects to establish correspondences. This was achieved by mapping the gray matter voxels to the native cortical surface and registering them onto the 32k standard triangulated mesh at a 2 mm spatial resolution. Finally, each time series was temporally normalized to zero-mean and unit-variance. We concatenated the time series of 15-minute scans acquired in the same sessions, obtaining two 30-minute rs-fMRI datasets for each of the 100 subjects and used them to evaluate our approach.

2.2 Initial Parcellation via Supervortex Clustering

We start the parcellation process by clustering the cortical vertices into a set of functionally uniform sub-regions. This stage does not only enforce spatial continuity within clusters, but is also beneficial for delineating more reproducible parcellations. To this end, we propose a k -means clustering algorithm inspired by SLIC superpixels [1]. Differently from the classical k -means, we limit the search space for each cluster to reduce the number of distance calculations and we define a hybrid distance function¹ which is capable of grouping highly correlated vertices, yet ensuring spatial continuity within clusters.

The cortical surface is represented as a smooth, triangulated mesh with no topological defects. Initially, k seeds are selected as the singleton clusters by uniformly sub-sampling the mesh. The algorithm iteratively associates each vertex with a cluster by computing their similarity using an Euclidean function in the form of $\sqrt{(d_c/N_c)^2 + (d_g/N_g)^2}$, where d_c and d_g correspond to the functional and spatial distance measures, respectively. Functional similarity between two vertices is measured by the Pearson’s distance transformation of their corresponding time series. This transformation ensures the distance between highly correlated vertices being close to zero, thus increases their likelihood of being assigned to the same cluster. Spatial proximity is measured by the geodesic distance along the cortical surface, which is approximated as the length of the shortest path between the nodes in the mesh graph. N_c and N_g refer to the normalization factors, which are set to their corresponding maximal values in a cluster². The algorithm converges when none of the clusters change between two consecutive iterations. Clustering decreases the dimensionality of the dataset to the number of supervertices, thus reduces the computational cost of the subsequent stages. Each supervortex is now represented by the average time series of the constituent vertices, minimizing the effects of noisy signals throughout the dataset.

¹ Please note that our distance function is not a metric, since it does not necessarily satisfy the triangle inequality.

² We straightforwardly set N_c to 2, since the Pearson’s distance values fall within the range $[0, 2]$. Similarly, N_g is set to the predefined local search limit, since the maximum distance within a cluster cannot exceed it.

2.3 Single-Level Parcellation via Hierarchical Clustering

The supervertices over-parcellate the cortical surface, therefore a second stage should be deployed in order to obtain a reasonable number of parcels without losing the ability to represent the functional organization of the cortex and having non-uniform functional patterns within sub-regions. Towards this end, we join supervertices into non-overlapping parcels using agglomerative hierarchical clustering. This approach builds a hierarchy of clusters using a bottom-up strategy in which pairs of clusters are merged if their similarity is the maximal among the other pairing clusters. We only join adjacent clusters into a higher level in order to ensure the spatial continuity throughout the parcellation process. The algorithm is driven by Ward’s linkage rule and the similarity between pairing clusters is computed by Pearson’s distance. The algorithm iteratively constructs a dendrogram, in which the leaves represent the supervertices and the root represents an entire hemisphere. Cutting this tree at different levels of depth produces parcellations with the desired precision. We investigate the effect of different granularities on the parcel reproducibility and functional consistency in the following section.

2.4 Groupwise Parcellation via Spectral Clustering

Connectome analyses usually require a reliable groupwise representation for identifying common functional patterns across groups of healthy or disordered subjects and compare how the connectivity changes, for example, through aging. To this end, we deploy a final clustering stage in order to identify the group parcellations. We compute a graphical model of the parcel stability across the whole population [9], in which an edge between two vertices is weighted by the number of times they appear in the same parcel across all individual parcellations. Notably, the spatial integrity of the parcels is automatically guaranteed, since only vertices sharing the same cluster membership can have a correspondence between each other. The graph is subdivided by spectral clustering with normalized cuts [3] into pre-defined number of sub-regions, thus similar to individual parcellations, allowing analysis of the connectome at different levels of detail.

3 Results

We assess the parcellation performance in two ways: (a) reproducibility and (b) functional consistency. Reproducibility is measured with a two-pass Dice score-based method suggested in [2]. In the first pass, overlapping sub-regions are matched and given the same label based on their Dice scores. Functional inconsistency due to low SNR usually results in a higher number of parcels than desired. To eliminate its effect on the performance, a second-pass is applied to the parcellations and over-segmented sub-regions are merged. A group of parcels is considered as over-segmented if at least half of it overlaps with a single parcel in the other parcellation. After locating and merging over-segmented parcels, the

average Dice score is used to assess the accuracy of the resulting parcellations. Functional consistency of a parcellation is assessed with respect to homogeneity and silhouette width [3]. Homogeneity is computed as the average pairwise correlations within each parcel after applying Fisher’s z-transformation. Silhouette width combines homogeneity with inter-parcel separation³, in order to assess the effectiveness of the parcellation algorithm in terms of generating functionally consistent sub-regions⁴. Algorithms were run on the left and right hemispheres separately and performance measurements were computed at different scales, starting from 500 parcels per hemisphere and in decrements of 50. Subsequently, results for the left and right hemispheres were averaged for the final plots.

We present single-subject reproducibility and functional consistency results obtained by our approach (3-LAYER), region growing (RG-HC), and normalized cuts (NCUT) in Fig. 1. Each subject has two rs-fMRI datasets, therefore reproducibility can be computed by matching parcellations derived from them separately. Functional homogeneity of the first parcellation was computed using the second rs-fMRI dataset, and vice versa, to avoid bias that may have emerged during the computation of the parcellations. Both 3-LAYER and RG-HC were run with ~ 2000 initial clusters, each cluster having ~ 30 cortical vertices on average. Results indicate that RG-HC and 3-LAYER perform similarly in terms of computing functionally segregated parcels; the former having slightly better performance at the lower resolutions, whereas the latter is able to generate single-subject parcellations with higher homogeneity, having 1-5% better scores. This is primarily due to the fact that we represent the cortical surface by functionally uniform supervertices and incorporate a flexible spatial constraint into our distance function, enabling any vertex to be assigned to a cluster if they exhibit high correlation and are spatially close (but not necessarily neighbors) as opposed to the region growing, which is based on stable seed points and a more strict definition of spatial proximity. However, region growing shows moderately better performance in terms of reproducibility when the same seed set (RG-HC2) is used at both parcellations, but this has no impact on the functional homogeneity. Although NCUT generates highly reproducible parcellations at low resolutions, it produces the least homogeneous sub-regions and shows a poor performance in silhouette width. Its high reproducibility can be attributed to the bias emerging from the structure of the cortical meshes that were used to drive the algorithm, which tends to generate evenly shaped parcels at the expense of losing functional segregation ability as also discussed in [2] and [3].

Functional consistency of the parcellations was also qualitatively measured by inspecting the connectivity profiles of the parcellated sub-regions thoroughly. We identified that sharp transitions across functional patterns are more significantly aligned with our parcellation boundaries compared to the other algorithms.

³ Inter-parcel separation is computed as the average of the correlations between the vertices constituting a parcel and the remaining vertices across the cortex.

⁴ Silhouette width is defined as $\frac{(H-S)}{\max\{H,S\}}$, where H is the within-parcel homogeneity and S is the inter-parcel separation for a given parcel. Obtaining silhouette widths of close to 1 indicates a highly reliable and functionally consistent parcellation.




Fig. 1. Single-subject reproducibility (*left*), functional homogeneity (*middle*), and functional segregation (*right*) results obtained by our approach (3-LAYER), region growing with different (RG-HC1) and same seeds (RG-HC2), and normalized cuts (NCUT). See the text for the experimental setup.

Further examinations on the parcellations revealed that our boundaries matched well with the borders in cytoarchitectonic cortical areas, especially within the somato-sensory and vision cortex, showing agreement with the findings in [2].

The group level parcellations were assessed by dividing the subjects into two equally sized subgroups by random permutation and then computing a group parcellation for each subset. Reproducibility was measured by matching two group parcellations using the same Dice score method. Functional consistency of the groupwise parcellations was computed by measuring the parcel homogeneity and silhouette width based on each individual subject’s rs-fMRI time series and then averaging them within the subgroups. All measurements were repeated for 10 times, each time forming subgroups with different subjects. Performance scores computed in each repetition were then averaged and plotted in Fig. 2. Groupwise parcellations were computed by using the stability graphs obtained from each method’s single-level parcellations. In addition, we also computed a groupwise parcellation by averaging the rs-fMRI datasets of all subjects (after applying Fisher’s z-transformation) and parcellating this average dataset by spectral clustering (MEAN) [3]. In general, the group-level results exhibit a similar tendency with those of the individual subjects. MEAN and NCUT show higher reproducibility at low resolutions, however 3-LAYER outperforms them for increasing number of parcels, obtaining upto 15% better scores. It generates the most functionally consistent parcellations at almost all resolutions achieving up to 4% higher homogeneity and 3% better functional segregation than the other approaches, except for 50-150 parcels, where RG-HC performs slightly better. These results may indicate that our parcellations can effectively reflect common functional characteristics within the population, being minimally affected by the functional and structural variability across different subjects.

Finally, for visual review, we present the parcellations of an arbitrary subject and the groupwise parcellations of the population obtained by the proposed algorithm in Fig. 3. We also provide the Dice scores computed between the single-level parcellations and their respective group representations in order to show




Fig. 2. Group-level reproducibility (*left*), functional homogeneity (*middle*) and functional segregation (*right*) results obtained by our approach (3-LAYER), region growing (RG-HC), normalized cuts (NCUT) and averaging rs-fMRI datasets (MEAN). See the text for the experimental setup.




Fig. 3. (a) Parcellations obtained by the proposed method (3-LAYER) for 200 parcels. Single-level and groupwise parcellations in the first two columns were derived from different rs-fMRI datasets of the same subject and from different subgroups of the population, respectively. The third column shows the differences between the first and second parcellations in each row. (b) Average Dice scores computed between the single-level parcellations and their respective group representations.

the robustness of the proposed method in terms of coping with the functional variability within the population, especially at high resolutions.

4 Conclusions

We presented a new three-layer clustering approach to parcellate the cerebral cortex using resting-state fMRI. Our experiments at the single-subject and group levels demonstrated that the proposed algorithm can produce reliable parcellations, with higher reproducibility and functional consistency compared to state-of-the-art approaches, therefore can be reliably and effectively used for network analysis. The three-layer method is in general more successful in grouping correlated vertices together, thus fulfills a critical requirement of connectome studies.

Having these promising results using solely rs-fMRI, we are now working on a multi-modal approach to improve the precision of the parcellation borders by incorporating task-fMRI into the parcellation framework. Another challenge in the parcellation problem is finding an optimal number of parcels. Our initial experiments showed that, the functional transitions in the connectivity profiles can be utilized to drive an objective function at the hierarchical clustering stage, thus can be used for this purpose.

Acknowledgments. Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657). The research leading to these results received funding from the European Research Council under the European Unions Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 319456.

References

1. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Susstrunk, S.: SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE T. Pattern Anal.* 34(11), 2274–2282 (2012)
2. Blumensath, T., Jbabdi, S., Glasser, M.F., Van Essen, D.C., Ugurbil, K., Behrens, T.E., Smith, S.M.: Spatially constrained hierarchical parcellation of the brain with resting-state fMRI. *NeuroImage* 76, 313–324 (2013)
3. Craddock, R.C., James, G., Holtzheimer, P.E., Hu, X.P., Mayberg, H.S.: A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum. Brain Mapp.* 33(8), 1914–1928 (2012)
4. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M.: The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* 80, 105–124 (2013)
5. Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., Petersen, S.E.: Functional network organization of the human brain. *Neuron* 72(4), 665–678 (2011)
6. de Reus, M.A., van den Heuvel, M.P.: The parcellation-based connectome: Limitations and extensions. *NeuroImage* 80, 397–404 (2013)
7. Shen, X., Tokoglu, F., Papademetris, X., Constable, R.T.: Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. *NeuroImage* 82, 403–415 (2013)
8. Sporns, O., Tononi, G., Ktetter, R.: The Human Connectome: A structural description of the human brain. *PLoS Comput. Biol.* 1(4), e42 (2005)
9. van den Heuvel, M., Mandl, R., Hulshoff Pol, H.: Normalized cut group clustering of resting-state fMRI data. *PLoS One* 3(4), e2001 (2008)
10. Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zollei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L.: The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* 106(3), 1125–1165 (2011)