Boundary Mapping through Manifold Learning for Connectivity-Based Cortical Parcellation

Salim Arslan, Sarah Parisot, and Daniel Rueckert

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

Abstract. The study of the human connectome is becoming more popular due to its potential to reveal the brain function and structure. A critical step in connectome analysis is to parcellate the cortex into coherent regions that can be used to build graphical models of connectivity. Computing an optimal parcellation is of great importance, as this stage can affect the performance of the subsequent analysis. To this end, we propose a new parcellation method driven by structural connectivity estimated from diffusion MRI. We learn a manifold from the local connectivity properties of an individual subject and identify parcellation boundaries as points in this low-dimensional embedding where the connectivity patterns change. We compute spatially contiguous and non-overlapping parcels from these boundaries after projecting them back to the native cortical surface. Our experiments with a set of 100 subjects show that the proposed method can produce parcels with distinct patterns of connectivity and a higher degree of homogeneity at varying resolutions compared to the state-of-the-art methods, hence can potentially provide a more reliable set of network nodes for connectome analysis.

1 Introduction

Connectome analysis has recently gained a lot of attention due to its potential to reveal the functional and structural architecture of the human brain, as well as understand its evolution through development, aging, and neurological disorders [14]. Brain connectivity is typically analyzed via graphical models obtained by connecting cortical regions to each other with respect to the similarity between their connectivity profiles, derived from functional MRI (fMRI) or diffusion imaging (dMRI). In a whole-brain connectivity analysis, parcellation of the cortex constitutes an integral part of the pipeline, as the performance of the subsequent stages depends on the ability of the parcels to reliably represent the underlying connectivity [6]. Traditionally, parcellations derived from anatomical landmarks or randomly partitioned subregions have been used for connectome analysis, however such parcellations generally fail to fully reflect the function of the cortical architecture [14]. More recent approaches take into account the connectivity information, generally in association with clustering algorithms [5, 12, 1, 2] in order to group vertices of connectional similarity [16]. Despite promising results, the parcellation problem is still open to improvements. This is primarily due to the fact that the problem itself is ill-posed, thus, obtaining accurate

parcels both depends on the proposed method's fidelity to the given data [12] and its capacity to differentiate vertices with different connectivity profiles [6].

To this end, we introduce a new parcellation method, in which we learn a manifold from local connectivity characteristics of an individual subject and develop an effective way of computing parcels from this manifold. Our approach rests on the assumption that through dimensionality reduction, we can capture the underlying connectivity structure that may not be visible in high-dimensional space [10]. We use the manifold to locate transition points where connectivity patterns change and interpret them as an abstract delineation of the parcellation boundaries. After projecting back to the native cortical space, these boundaries are used to compute non-overlapping and spatially contiguous parcels. We achieve this with a watershed segmentation technique, originally utilized to parcellate resting-state correlations [8]. Nonlinear manifold learning has been formerly used to identify functional networks from fMRI [9, 17] and for surface matching [10], as well as within many other fMRI analysis techniques, such as [15]. Nevertheless, we propose to use such technique in association with dMRI-based structural connectivity and boundary mapping, in order to compute cortical parcellations for individual subjects, which can be used as the network nodes in a whole-brain connectome analysis.

We assess the parcellation quality based on parcel homogeneity [2, 8] and silhouette analysis [5, 6]. Besides the dMRI data, we also evaluate the parcellations with functional connectivity data obtained from resting-state fMRI as a means of external validation [6]. Our method is compared to the state-of-the-art connectivity-based parcellation techniques [5, 12], as well as two parcellation schemes which do not take into account any connectivity information [16]. In addition, we show the extent to which our parcellation boundaries agree with well-established patterns of cortical myelination and cytoarchitecture.

2 Method

We start with preprocessing the dMRI data using probabilistic tractography to estimate a structural connectivity network, which is then reduced in dimensionality through manifold learning. Driven by the boundaries identified in the low-dimensional embedding as points where connectivity patterns change, we utilize a watershed segmentation to achieve the final parcellation (Fig. 1).

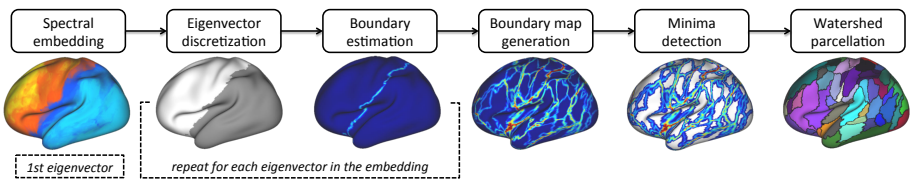

Fig. 1. Parcellation pipeline, summarizing all steps after preprocessing.

Estimating Structural Connectivity We perform whole-brain probabilistic tractography on dMRI data by following the procedures summarized in [12]. We applied an element-wise log transformation to the tractography matrix to reduce the bias towards short connections and sampled 5000 streamlines from each of the cortical vertices. We define a connectivity fingerprint for each vertex v_i by counting the number of streamlines that connect v_i to other vertices. Each subject's structural connectivity network $C \in \mathbb{R}^{N\times N}$ is estimated as the cross-correlations of the fingerprints associated with each vertex, where N is the number of vertices. We excluded the medial wall vertices from further processing as they do not possess reliable information for connectivity analysis.

Learning a Manifold from Connectivity We propose to use *Laplacian eigenmaps* to compute a nonlinear embedding from a connectivity network [3]. This method can reveal the intrinsic geometry of the underlying connectivity by forming an affinity matrix based on how vertices are connected within their neighborhoods. To this end, we transform C into a locality-preserving affinity matrix $W \in \mathbb{R}^{+N\times N}$ by only retaining the correlations of the k nearest neighbors of each vertex. We set $k = 100$ in order to effectively capture the local connectivity structure and to ensure that the affinity matrix is connected and positive-semidefinite (i.e. all $W_{ij} \geq 0$) for each subject. A nonlinear embedding is computed through spectral decomposition of the normalized graph Laplacian, defined as $L = D^{-1/2}(D - W)D^{-1/2}$, where D is a diagonal matrix with each entry $D_{ii} = \sum_j W_{ij}$ representing the degree of v_i. Solving the generalized eigenvector problem [3] with respect to L reveals the eigenvectors $f_0, f_1, \ldots, f_{N-1}$, ordered according to their eigenvalues $0 = \lambda_0 \leq \lambda_1 \leq \ldots \leq \lambda_{N-1}$. After omitting the eigenvector f_0 corresponding to λ_0, we can use the next d eigenvectors to define an embedding that can approximate a low dimensional manifold [3]. Hence, each cortical vertex v_i can be expressed as a row in this spectral embedding, i.e. $i \mapsto (f_1(i), \ldots, f_d(i))$.

Eigenvector Discretization The process of dimensionality reduction preserves local connectivity as well as imposes a natural clustering of the data [3]. Therefore, the parcellation problem can be cast as a graph partitioning problem and one would attempt to subdivide the connectivity graph with spectral clustering, e.g. using the normalized cuts criterion and solving the aforementioned generalized eigenvalue problem [13]. In particular, each of the smallest eigenvectors corresponds to a real valued solution that optimally sub-partitions the graph. These partitions can be approximated by transforming the real valued eigenvectors into discrete forms, ideally by dividing them into two parts with respect to a splitting point [13]. This can further be generalized towards a multiway partitioning with a recursive or simultaneous discretization of the smallest eigenvectors [13], and thus, can be used to obtain a parcellation [5]. However, by definition, our affinity matrix does not impose any spatial constraints, hence such spectral methods cannot guarantee spatial contiguity within the parcels. Instead, we propose a more effective way of deriving parcellations from discrete

eigenvectors and later show that this method can produce more reliable parcellations compared to spatially constrained spectral clustering.

We discretize the eigenvectors using k-means and partition each eigenvector into two subregions. The edge between these subregions potentially provides good separation points towards obtaining a parcellation, as the vertices within the same subregions tend to have similar connectivity properties, whilst the points closer to the boundary attribute to the cortical areas where the connectivity is in transition. For example, Fig. 2(a) shows that connectivity profiles of different vertices may exhibit similar or varying patterns, depending on their relative location to an edge. In order to show that this tendency holds across the whole cortex, we randomly selected vertices from one subregion adjacent to the edge and paired them with their closest neighbors residing in the other subregion. Keeping the distance between the vertices in pairs approximately the same, we selected new pairs of vertices, but this time from within the same subregions. We then measured the average correlation between the paired vertices' connectivity profiles in each set and repeated this for all eigenvectors and subjects. Fig. 2(b) shows that, the similarity between the connectivity profiles of vertices drops by at least 20% if they reside on different sides of a boundary.

Boundary Map Generation and Cortical Parcellation To locate the connectivity transition points and construct a boundary map, we first transfer the discrete eigenvectors back to the native high-dimensional space. We then calculate the gradients of each eigenvector across the cortical surface and combine them into a boundary map. This map constitutes a more robust substitution for the boundary maps based on gradients directly calculated from the spatial correlations [8], since it can adjust for possible spurious gradients. In addition, the traditional boundary mapping requires a considerable amount of data in order to effectively model the brain function at the individual level [11] and only becomes reliable for parcellation when averaged across many subjects/datasets [8]. In order to obtain the final parcellations from the boundary map, we use a marker-controlled watershed algorithm [8]. We define a set of markers on the

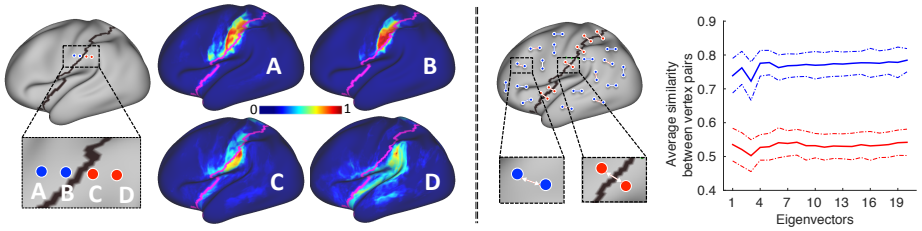

(a) (b)

Fig. 2. (a) Connectivity profiles of vertices from different sides of a boundary. (b) *Left:* Illustration of the vertex selection procedure. *Right:* Average similarity (correlation) between paired vertices for each eigenvector. Dotted lines show the standard deviations.

boundary map where each marker corresponds to an estimated parcel position and then grow these markers until a boundary is reached or two ridges touch each other in the flooding process of the watershed. The marker definition is typically performed by defining a threshold on the boundary map. We set this threshold to the 25th percentile of the boundary map intensities, since in many empirically tested cases, this effectively revealed approximate parcel locations to be used as ideal markers for a watershed transformation.

3 Experiments

Data Experiments are conducted on a set of 100 randomly selected adults (54 females, age 22-35) from the Human Connectome Project (HCP) S500 release[1]. All data have been acquired and preprocessed following the HCP minimal preprocessing pipelines [7]. For each subject, the gray-matter voxels have been registered onto the 32k triangulated mesh at 2 mm spatial resolution, yielding a standard set of cortical vertices per hemisphere.

Evaluation We assess the quality of the parcellations using two validation techniques: parcel homogeneity [2, 8] and silhouette analysis [5, 6]. The former expresses the degree of homogeneity that a parcellation exhibits by calculating average cross-correlations within each parcel. Silhouette analysis combines parcel homogeneity with inter-parcel separation and measures how vertices within a parcel are similar to each other, compared to the vertices in the nearest parcels [6]. The goodness-of-fit is estimated based on the structural connectivity data from which the parcellations have been derived. In addition, we evaluate parcellations by measuring their extent to reflect the underlying connectivity estimated from resting-state fMRI, which can provide an external data source for validation [6]. We compare our parcellations to the ones obtained by hierarchical clustering applied to the low-dimensional embedding (HC-Low), hierarchical clustering driven by the connectivity profiles in the high-dimensional space (HC-High), multi-scale spectral clustering (M-Scale) [12], normalized cuts (N-Cuts) [5], random parcellations by Poisson disk sampling, and geometric parcellations, i.e. k-means clustering of the vertex coordinates [16]. All methods are spatially constrained to ensure the contiguity of parcels. M-Scale and HC-High are based on an initial connectivity-based over-parcellation of the cortex to compensate for the noise, and thus, to obtain higher accuracy (1000, 2000 and 3000 regions for M-Scale; 3000 regions for HC-High). Random and geometric parcellations do not account for any connectivity information, therefore provide a baseline for the assessment [16].

Results As there is no known optimal number of parcels, we evaluate the proposed method at different scales, determined by the number of eigenvectors incorporated into the boundary map. We present results for $d = 10$, 15, and 20

[1] `http://www.humanconnectome.org/documentation/S500/`

eigenvectors per hemisphere, which on average, yield parcellations with around 180, 230, and 280 regions for each subject, respectively. Our experiments with fewer eigenvectors resulted in very coarse parcellations that may not be ideal for network analysis, whereas using $d > 30$ eigenvectors led to noisy boundary maps, generating many unreliable parcels. For a fair comparison, other methods are tuned to use the same number of parcels as inferred by our models. Validation measures were calculated for each subject-parcellation pair and then averaged across all subjects. We present the results based on structural and functional connectivity data in Fig. 3 and Fig. 4, respectively.

Fig. 3 shows that our method surpasses other approaches at all resolutions in terms of silhouette analysis and performs equally effective as HC-Low with respect to homogeneity. This may indicate that, spectral embedding, which drives both methods, can successfully reveal the intrinsic geometry of the underlying connectivity, and hence, provides a more robust set of features towards parcellating the cortical surface. In addition, the way we utilize discrete eigenvectors for deriving parcellations help obtain more distinct parcels compared to the others. This can be deduced from silhouette coefficients, where we especially perform better than HC-Low, which directly applies a traditional clustering approach to the spectral coordinates. In addition, considering the results obtained by HC-High, we can infer that nonlinear dimensionality reduction can identify local connectivity patterns which may not be directly detected in the high dimensional space. On the other hand, M-Scale and N-Cuts can obtain reliable parcellations only to some extent. These spectral approaches solely consider the immediate neighbors for the construction of their affinity matrices. Therefore, they may fail to fully capture the underlying connectivity.

The difference in performance between our approach and the others becomes more prominent with the resting-state functional connectivity results (Fig. 4). Both homogeneity and silhouette analysis indicate that, the proposed method can effectively subdivide the cortical surface into functionally coherent subregions, hence can better reflect the underlying function. Although, other methods can generate homogeneous parcels to some degree, they fail to separate vertices with different signals from each other, as indicated by silhouette coefficients.

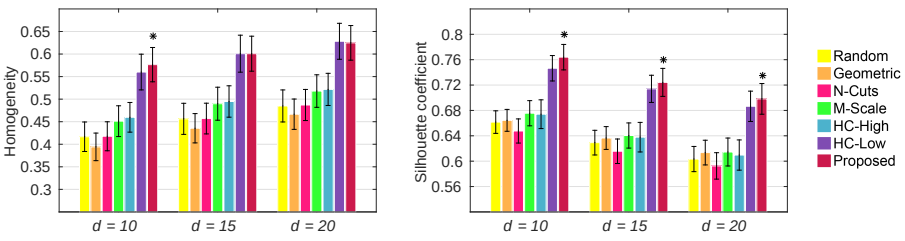

Fig. 3. Quantitative results based on structural connectivity estimated from dMRI. Error bars represent the variability across subjects. Stars (*) indicate statistical significance between the winner and the runner-up with $p < 0.01$.

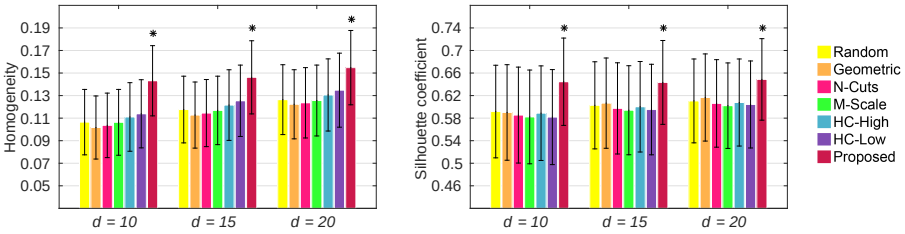

Fig. 4. Quantitative results based on resting-state functional connectivity.

Finally, visual assessment of parcellations shows some alignment with Brodmann's cytoarchitectural areas and highly myelinated cortical regions (see Supplementary Material). Dice-based overlapping measures [4] indicate that this observation is substantially consistent across subjects, especially for the motor (BA[1,3,4]) and visual cortex (BA17), with average Dice scores of 0.81 ($\pm$0.05) and 0.82 ($\pm$0.05), respectively.

4 Conclusions

In this paper, we introduced a new connectivity-driven parcellation approach based on dMRI. The proposed method models the local connectivity characteristics with manifold learning and describes an effective use of this manifold to identify locations where connectivity patterns change. Particularly, these transition locations are interpreted as an abstraction of the parcellation boundaries, and hence, used to derive distinct parcels at different scales. We showed that our parcellations can more reliably capture the underlying connectivity of the brain compared to a set of other approaches. This paper focuses on developing a complete framework for computing subject-specific parcellations, which can be used in many application areas, such as for driving a registration process based on brain connectivity. In addition, a planned future work is to explore the variability across individual parcellations towards generating a connectivity-based cortical atlas, which can allow performing population level connectome studies.

Acknowledgments Authors would like to thank Markus Schirmer for providing the random parcellations. Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54-MH091657). The research leading to these results received funding from the European Research Council under the European Unions Seventh Framework Programme (FP/20-2013) / ERC Grant Agreement no. 319456.

References

1. Arslan, S., Parisot, S., Rueckert, D.: Joint spectral decomposition for the parcellation of the human cerebral cortex using resting-state fMRI. In: Ourselin, S.,

Alexander, D., Westin, C.F., Cardoso, M. (eds.) IPMI 2015. LNCS. vol. 9123, pp. 85–97. Springer, Heidelberg (2015)
2. Arslan, S., Rueckert, D.: Multi-level parcellation of the cerebral cortex using resting-rtate fMRI. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS. vol. 9351, pp. 47–54. Springer, Heidelberg (2015)
3. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. Neural Comput. 15(6), 1373–1396 (2003)
4. Bohland, J.W., Bokil, H., Allen, C.B., Mitra, P.P.: The brain atlas concordance problem: Quantitative comparison of anatomical parcellations. PLoS One 4(9), e7200 (2009)
5. 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)
6. Eickhoff, S.B., Thirion, B., Varoquaux, G., Bzdok, D.: Connectivity-based parcellation: Critique and implications. Hum. Brain Mapp. 36(12), 4771–4792 (2015)
7. 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)
8. Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M., Petersen, S.E.: Generation and evaluation of a cortical area parcellation from resting-state correlations. Cereb. Cortex 26(1), 288–303 (2016)
9. Langs, G., Sweet, A., Lashkari, D., Tie, Y., Rigolo, L., Golby, A.J., Golland, P.: Decoupling function and anatomy in atlases of functional connectivity patterns: language mapping in tumor patients. NeuroImage 103, 462–475 (2014)
10. Langs, G., Golland, P., Ghosh, S.S.: Predicting activation across individuals with resting-state functional connectivity based multi-atlas label fusion. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS. vol. 9350, pp. 313–320. Springer, Heidelberg (2015)
11. Laumann, T.O., Gordon, E.M., Adeyemo, B., Snyder, A.Z., Joo, S.J., Chen, M.Y., Gilmore, A.W., McDermott, K.B., Dosenbach, N.U., Schlaggar, B.L., Mumford, J.A., Poldrack, R.A., Petersen, S.E.: Functional system and areal organization of a highly sampled individual human brain. Neuron 87(3), 657–670 (2015)
12. Parisot, S., Arslan, S., Passerat-Palmbach, J., Wells, W.M., Rueckert, D.: Tractography-driven groupwise multi-scale parcellation of the cortex. In: Ourselin, S., Alexander, D., Westin, C.F., Cardoso, M. (eds.) IPMI 2015. LNCS. vol. 9351, pp. 600–612 (2015)
13. Shi, J., Malik, J.: Normalized cuts and image segmentation. IEEE Trans. Pattern Anal. Mach. Intell. 22(8), 888–905 (2000)
14. Sporns, O.: The human connectome: A complex network. Ann. N. Y. Acad. Sci. 1224(1), 109–125 (2011)
15. Thirion, B., Dodel, S., Poline, J.B.: Detection of signal synchronizations in resting-state fMRI datasets. NeuroImage 29(1), 321–327 (2006)
16. Thirion, B., Varoquaux, G., Dohmatob, E., Poline, J.B.: Which fMRI clustering gives good brain parcellations? Front. Neurosci. 8(167) (2014)
17. Wang, D., Buckner, R.L., Fox, M.D., Holt, D.J., Holmes, A.J., Stoecklein, S., Langs, G., Pan, R., Qian, T., Li, K., Baker, J.T., Stufflebeam, S.M., Wang, K., Wang, X., Hong, B., Liu, H.: Parcellating cortical functional networks in individuals. Nat. Neurosci. 18(12), 1853–1860 (2015)