A Continuous Flow-Maximisation Approach to Connectivity-driven Cortical Parcellation

Sarah Parisot, Martin Rajchl, Jonathan Passerat-Palmbach, and Daniel Rueckert

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

Abstract. Brain connectivity network analysis is a key step towards understanding the processes behind the brain's development through ageing and disease. Parcellation of the cortical surface into distinct regions is an essential step in order to construct such networks. Anatomical and random parcellations are typically used for this task, but can introduce a bias and may not be aligned with the brain's underlying organisation. To tackle this challenge, connectivity-driven parcellation methods have received increasing attention. In this paper, we propose a flexible continuous flow maximisation approach for connectivity driven parcellation that iteratively updates the parcels' boundaries and centres based on connectivity information and smoothness constraints. We evaluate the method on 25 subjects with diffusion MRI data. Quantitative results show that the method is robust with respect to initialisation (average overlap 82%) and significantly outperforms the state of the art in terms of information loss and homogeneity.

1 Introduction

Brain connectivity network analysis can provide key insights into the brain's organisation and its evolution through disease and ageing. Building these networks from functional (fMRI) or diffusion (dMRI) MR imaging is a challenge in itself due to the high dimensionality of the data. Therefore, network construction requires an initial parcellation stage of the cortical surface into distinct regions. While anatomical or random parcellations are used in most existing studies, they do not necessarily represent the brain's underlying connectivity accurately and can introduce a strong bias in the constructed network and its subsequent analysis [14].

In order to address the shortcomings of anatomical and random parcellations, connectivity driven brain parcellation has received an increasing amount of attention. In addition to providing a sensible basis for constructing connectivity networks, it can enable the identification of functionally specialised brain regions. The problem is typically cast as a clustering problem of the cortical surface vertices, where the goal is to maximise the correlation between connectivity profiles (dMRI) or time series (fMRI). While common clustering techniques can be used when parcellating a subset of the brain [1, 8], the problem becomes more difficult

when the aim is to parcellate the whole cortical surface. Hierarchical [3, 10] and spectral clustering [5, 13] based methods are the most popular approaches. However, the latter tends to create homogeneous parcel sizes, which can disagree with the brain's structure, while hierarchical clustering implies that regions boundaries are the same at different granularity levels. This can propagate errors from low to high resolution.

A Markov Random Field (MRF) based method for fMRI driven parcellation was proposed in [7]. It maximises the correlation between nodes and a parcel centre subject to smoothness constraints. The method considers all nodes (vertices on the cortical surface) as potential parcel centres and adds a penalty term at the introduction of a new label. The use of all nodes as parcel centres makes the method sensitive to noise and computationally very demanding. At the same time, there is no direct control on the number of parcels obtained.

In this paper, we propose an iterative MRF formulation to the parcellation problem based on the continuous flow-maximisation solver introduced in [15]. Each iteration consists of maximising the correlation between the nodes and parcel centres, as well as a smart update of the centres based on the cluster's homogeneity. A coarse to fine multi-resolution implementation allows to reduce the influence of noise while efficiently exploring the space for updates of the parcel centres. The experimental evaluation is performed on 25 subjects from the Human Connectome Project (HCP) database for dMRI driven parcellation. We show that the method is robust with respect to initialisation, and significantly outperforms existing methods in terms of information loss and parcel homogeneity.

2 Methodology

We consider a subject whose cortical surface is to be parcellated into a set of K parcels. The surface is represented as a mesh graph $\mathcal{S} = \{\mathcal{V}, \mathcal{E}\}$, where $\mathcal{V}$ corresponds to a set of N vertices and $\mathcal{E}$ to the edges between neighbouring vertices. The problem is simplified by inflating the cortical mesh to a sphere on which computations are performed. The graph is associated with an affinity matrix describing how dissimilar two nodes are in terms of connectivity. The correlation ρ between the vertices' connectivity profiles obtained from tractography (dMRI-driven) or times series (fMRI-driven) is a commonly employed measure.

2.1 Iterative Markov Random Field Formulation

We cast the connectivity driven parcellation task as a vertex labelling problem, where each graph node $\mathbf{v}$ is to be assigned a label $l \in [\![1, K]\!]$ based on its affinity with a predefined parcel centre c_l. We adopt a coarse to fine multi-level MRF formulation that iteratively evaluates the label assignments of the vertices and updates the parcel centres based on the current parcellation configuration.

We initialise each resolution level R with the construction of a regular icosahedron $\mathcal{S}^R$ of constant radius RAD and the smooth resampling of the data term D to the new resolution.

Assignment Stage: Given a set of parcel centres $C^t = \{\mathbf{c}_1^t, \cdots, \mathbf{c}_K^t\} \subset \mathcal{V}$, the first stage of each iteration t consists of minimising the following energy:

$$\begin{aligned} E^t(\mathbf{l}) &= \sum_{i=1}^{K} \int_{\mathcal{S}} D(\mathbf{c}_i^t, \mathbf{v}) dx + \alpha(\mathbf{v}) \sum_{i=1}^{K} |\partial \mathcal{S}_i| \\ \text{s.t.} \quad & \cup_{i=1}^{K} \mathcal{S}_i = \mathcal{S}, \quad \mathcal{S}_i \cap \mathcal{S}_j = \emptyset, \forall i \neq j \end{aligned} \tag{1}$$

where $\mathcal{S}_i = \{\mathbf{v} \in \mathcal{V} | l(\mathbf{v}) = i\}$ corresponds to the subdomain of $\mathcal{S}$ assigned to label l_i and $D(\mathbf{c}_i, \mathbf{v})$ is the distance between a vertex and a parcel centre. The first term of the equation assigns nodes to the parcel based on how similar they are to its centre, while the second term enforces smoothness of the parcellation by minimising the length of the boundary of each subdomain. The strength of the smoothness term is defined by the position dependent cost $\alpha(\mathbf{v})$.

Centre Update Stage: For each parcel $\mathcal{S}_i$ obtained during the assignment stage, the centre is defined based on the correlation between nodes within the parcel. The parcel centre should have strong connectivity information and be the most similar to the rest of the parcel:

$$\mathbf{c}_i^{t+1} = \arg\max_{\mathbf{v} \in \mathcal{S}_i} \sum_{\mathbf{w} \in \mathcal{S}_i} \rho(\mathbf{v}, \mathbf{w}) \tag{2}$$

This smart centre update enables to guide the parcels' position based on connectivity information. The coarse to fine approach enables to explore the centre space efficiently and to avoid local minima due to the presence of noise. Centres are propagated from one resolution level to the next by minimising the geodesic distance d_g between the parcel centre and the new resolution's vertices on the inflated mesh: $\mathbf{c}_i^{R+1} = \arg\min_{\mathbf{v} \in \mathcal{S}^{R+1}} d_g(\mathbf{v}, \mathbf{c}_i^R)$.

2.2 Continuous Max-Flow Optimisation

We minimise the assignment stage's MRF energy using the recently proposed Continuous MaxFlow (CMF) algorithm [15] that we adapt to triangular spherical meshes. This solver offers two significant advantages over discrete MRF optimisation approaches. First, it is highly paralellisable, and as a result very scalable with respect to the mesh size and number of desired parcels. Second, it provides a continuous labelling, which results in smoother parcel boundaries and a better exploration of the solution space.

Similarly to discrete s-t-mincut/maxflow MRF optimisation methods [4], the continuous maxflow model constructs a directed graph through the introduction of two terminal nodes (source s and sink t). The state of a mesh vertex $\mathbf{v}$ is given by three flows: the source $p^s(\mathbf{v})$ and sink flows $p^t(\mathbf{v})$ connecting $\mathbf{v}$ to both terminal nodes, and a spatial flow field $\mathbf{q}(\mathbf{v})$ connecting neighbouring vertices. In the multi-label setting, K copies of the surface mesh are constructed in parallel, each associated with label-wise spatial and sink flows $\mathbf{q}_i(\mathbf{v}), p_i^t(\mathbf{v}), i \in [\![1, K]\!]$, while $p^s(\mathbf{v})$ remains unique. The multi-label configuration is illustrated in Fig.1a.

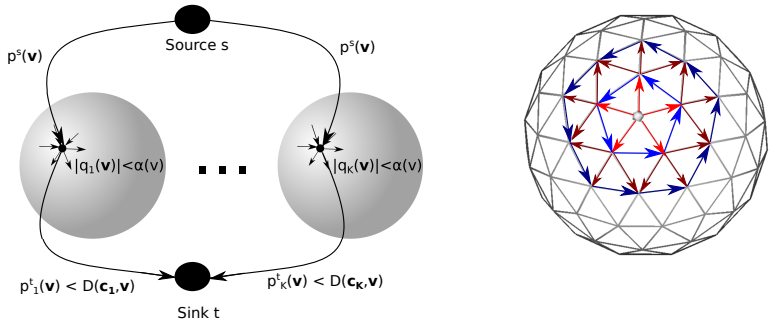

(a) (b)

Fig. 1: (a) Configuration of the spherical CMF algorithm with K labels. (b) Spatial flow computation: from a starting vertex (marked), alternation between longitudinal (red arrows) and latitudinal (blue arrows) computation.

The three flows are subject to capacity and conservation constraints:

$$\begin{aligned} \forall i \in [\![1, K]\!], \quad |\mathbf{q}_i(\mathbf{v})| < \alpha(\mathbf{v}), \quad p_i^t(\mathbf{v}) < D(\mathbf{c}_i, \mathbf{v}) \\ \forall i \in [\![1, K]\!], \quad (\nabla \cdot \mathbf{q}_i - p^s + p_i^t)(\mathbf{v}) = 0 \end{aligned} \tag{3}$$

As a result, the CMF algorithm consists of maximising the total flow $p^s(\mathbf{v})$ from the source under the aforementioned constraints. The optimisation is based on the augmented Lagrangian method through the introduction of the indicator function $u_i(\mathbf{v}), i \in [\![1, K]\!]$, that have been shown to correspond to the optimal continuous labelling functions [15].

One of the main challenges associated to the new mesh space is to adapt the spatial flow computation task from a rectangular grid (image pixels and neighbours) to a triangular mesh with different neighbourhood configurations. We split the spatial flow into two components (see Fig.1b): a latitudinal and a longitudinal one (with varying number of components) that are computed alternatingly. This ensures that the flow can be computed in parallel in a consistent flow direction. The CMF optimisation is performed at each iteration until convergence. After convergence, the parcellation is obtained by selecting for each node the label that maximises the labelling function u_i. This iterative process is then repeated until convergence of the labelling function (minimal label update).

3 Results

We evaluated our method on 25 different subjects from the latest release of HCP, which were preprocessed following the HCPs minimum processing pipeline [6]. The cortical surfaces are represented as a triangular mesh of 32k vertices per hemisphere. Vertices corresponding to the medial wall are excluded from parcellations due to the lack of reliable connectivity information in this region. We tested the method on structural connectivity based parcellation obtained from

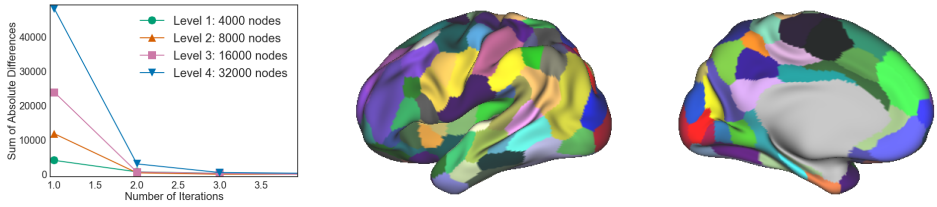

(a) (b)

Fig. 2: (a) Convergence of the labelling update at each resolution level: sum of absolute differences between the continuous labelling functions u_i of two consecutive iterations. (b) Associated parcellation after convergence.

20 different random initialisations via Poisson disc sampling, and four different number of labels (50, 100, 150 and 200 labels). We set the smoothness parameter $\alpha(\mathbf{v})$ to a constant to 0.1 over all vertices $\mathbf{v}$. Structural connectivity information is obtained from dMRI and tractography. The data term D is defined as $D(\mathbf{c}_i, \mathbf{v}) = 1 - \rho(\mathbf{c}_i, \mathbf{v})$, where ρ is the correlation between the vertices' connectivity profiles. The tractography matrix is obtained using FSLs bedpostX and probtrackX methods [2, 9] which estimate the fibres orientation at each voxel with a ball and stick model, and perform probabilistic tractography respectively. Following [9], we fitted three fibre compartments per voxel. 5000 streamlines were sampled from each of the vertices. Each entry of the tractography matrix counts the number of streamlines sampled from vertex $\mathbf{v}$ that reach vertex $\mathbf{q}$.

Our framework consists of four resolution levels (32k, 16k, 8k and 4k nodes). The convergence at each level was monitored by computing the sum of absolute differences between two consecutive labellings functions u_i. As illustrated in Fig.2 for a randomly selected experiment, all four resolution levels converge through the different centres updates. As expected, the convergence rates vary with respect to the quality of the initialisation and number of nodes.

The reproducibility with respect to initialisation is evaluated by computing the Dice similarity coefficient (DSC) between same subject parcellations. Depending on the initialisation, certain parcels can be split in several subparcels but have the same overall boundary. To find correspondences between parcellations, we consider a pair of parcellations P_1 and P_2 obtained from different initialisations. Each parcel of P_1 is assigned to the parcel in P_2 that has the highest overlap. Parcels in P_1 are merged if they are assigned to the same parcel in P_2 (split parcel case). The same process is repeated for P_2 with respect to the updated parcellation P_1^*. The DSC between the matched parcellations is computed for all subjects, numbers of labels and initialisations. The values are compared to the DSC between the random initialisations as a baseline. For all numbers of labels, we consistently obtain an average DSC of 0.82 ± 0.01, while

Table 1: Mean Kullback Leibler divergence (a) and parcel homogeneity (b) for all tested labels and methods. The best value is shown in bold. Statistically significant results: $^{*}p < 0.01$.

(a) Kullback Leibler divergence

Method	50 labels	100 labels	150 labels	200 labels
Random	2.67 ± 0.08	2.33 ± 0.08	2.09 ± 0.08	1.96 ± 0.08
Hierarchical	2.55 ± 0.08	2.15 ± 0.08	1.92 ± 0.07	1.80 ± 0.06
Spectral	2.54 ± 0.08	**2.10*** ± 0.08	1.89 ± 0.07	1.80 ± 0.07
Proposed	**2.50*** ± 0.09	2.16 ± 0.1	**1.87** ± 0.09	**1.74*** ± 0.07

(b) Homogeneity

Method	50 labels	100 labels	150 labels	200 labels
Random	0.11 ± 0.005	0.16 ± 0.008	0.19 ± 0.011	0.22 ± 0.015
Hierarchical	0.11 ± 0.006	0.16 ± 0.009	0.21 ± 0.013	0.24 ± 0.016
Spectral	0.11 ± 0.006	0.16 ± 0.009	0.20 ± 0.013	0.24 ± 0.015
Proposed	**0.12*** ± 0.006	**0.18*** ± 0.011	**0.22*** ± 0.015	**0.26*** ± 0.017

random initialisation only reaches 0.64 ± 0.02 on average. Figure 3 shows the comparison between example parcellations after matching, and the local average reproducibility of parcels. We can see that parcels boundaries are very similar despite different initialisations. Figure 3b shows that some regions have consistently matching boundaries while others are more variable across initialisations. This could be due to the fact that the connectivity information in these regions is too weak or that the differences are too subtle to drive the parcellation task.

The quality of the parcellation is evaluated through computation of the average intra-cluster connectivity homogeneity [5], and the Kullback Leibler divergence (KLD) between the normalised tractography matrix and its approximation obtained after parcellation. The KLD evaluates the information loss caused by such approximation. KLD and homogeneity values are compared to the ones obtained by Poisson disc sampling (random parcellations), multi-scale spectral clustering (SC) [11], and hierarchical clustering (HC). Both spectral and hierarchical clustering are spatially constrained and based on an initial connectivity-driven oversegmentation of the cortical surface to adjust for the noise (2000 regions for HC; 500, 1000 and 2000 regions for the multi-scale SC). For each subject, we compute the average KLD for parcellations obtained from the different initialisations. As illustrated in Table 1a, we can observe that the average KLD outperforms spectral and hierarchical clustering at most resolutions. Paired T-tests of our proposed method with the three different methods (random, SC, HC) show that we obtain significantly better results ($p \leq 0.01$) for most label configurations. Furthermore, homogeneity measures show that we consistently obtain significantly better values for all label configurations ($p \leq 0.01$). Average homogeneity and KLD results are shown in Table 1b. Correspondences between our obtained parcel boundaries and cortical myelination are shown in Fig. 3c.

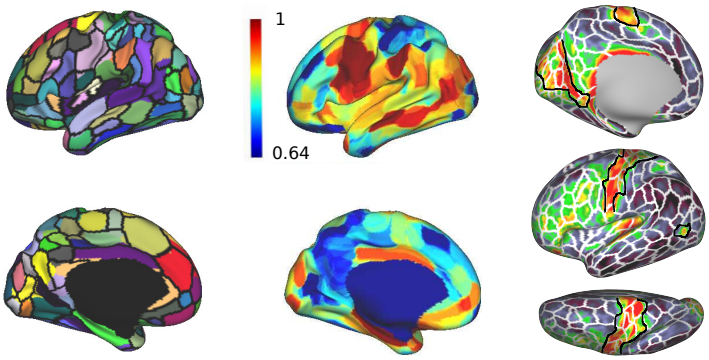

(a) (b) (c)

Fig. 3: (a,b) Reproducibility between initialisations: (a) A parcellation result is superimposed to the boundaries of a parcellation obtained with a different initialisation. (b) Local average DSC over all initialisations for the same subject. (c) Comparison of parcellation borders with cortical myelination. Parcel boundaries matching high myelin variations are highlighted.

4 Discussion

In this paper, we propose a continuous MRF model for connectivity-driven cortex parcellation. We develop a coarse to fine approach that iteratively updates parcel centres and the remaining nodes assignments. The inference is performed using a Continuous MaxFlow algorithm that is adapted to spherical meshes. We demonstrate the method's robustness with respect to initialisation and show that it significantly outperforms state-of-the-art clustering methods both in terms of information loss and homogeneity for most tested label configurations. The method is generalisable and can be applied to both fMRI and dMRI data. Furthermore, it is very flexible with respect to the definition of the cost function and smoothness term. For instance, it is straightforward to implement data-driven or inter-subject pairwise costs for groupwise parcellation. Other applications could also be considered, such as MRF based cortical registration [12]. The continuous solver could prove very useful in this kind of set up for finding a smooth deformation field. In order to show the robustness of the method, we have presented parcellation results obtained with a poor initialisation that has no correlation with the underlying connectivity. A smart centre initialisation would provide faster convergence, and potentially better consistency in less reproducible regions where the information from the tractography matrix is not strong enough. Identifying the regions that are the most reproducible also enables to identify where the connectivity information can be relied upon. This could provide interesting insight into multi-modal analysis of functional and structural data.

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 has received funding from the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 319456.

References

1. Anwander, A., Tittgemeyer, M., von Cramon, D.Y., Friederici, A.D., Knösche, T.R.: Connectivity-based parcellation of Broca's area. Cereb Cortex 17(4), 816–825 (2007)
2. Behrens, T., Berg, H.J., Jbabdi, S., Rushworth, M., Woolrich, M.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? NeuroImage 34(1), 144–155 (2007)
3. 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)
4. Boykov, Y., Funka-Lea, G.: Graph cuts and efficient ND image segmentation. International Journal of Computer Vision 70(2), 109–131 (2006)
5. Craddock, R.C., James, G.A., Holtzheimer, P.E., Hu, X.P., Mayberg, H.S.: A whole brain fMRI atlas generated via spatially constrained spectral clustering. Hum brain Mapp 33, 1914–1928 (2012)
6. 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., Essen, D.C.V., Jenkinson, M.: The minimal preprocessing pipelines for the Human Connectome Project. NeuroImage 80, 105 – 124 (2013)
7. Honnorat, N., Eavani, H., Satterthwaite, T., Gur, R., Gur, R., Davatzikos, C.: GraSP: Geodesic Graph-based Segmentation with Shape Priors for the functional parcellation of the cortex. NeuroImage 106, 207–221 (2015)
8. Jbabdi, S., Woolrich, M.W., Behrens, T.E.: Multiple-subjects connectivity-based parcellation using hierarchical Dirichlet process mixture models. NeuroImage 44, 373–384 (2009)
9. Jbabdi, S., Sotiropoulos, S.N., Savio, A.M., Graa, M., Behrens, T.E.J.: Model-based analysis of multishell diffusion MR data for tractography: How to get over fitting problems. Magn Reson Med 68(6), 1846–1855 (2012)
10. Moreno-Dominguez, D., Anwander, A., Knösche, T.R.: A hierarchical method for whole-brain connectivity-based parcellation. Hum Brain Mapp 35, 5000–5025 (2014)
11. Parisot, S., Arslan, S., Passerat-Palmbach, J., Wells III, W.M., Rueckert, D.: Tractography-Driven Groupwise Multi-scale Parcellation of the Cortex. In: Information Processing in Medical Imaging. pp. 600–612. Springer (2015)
12. Robinson, E.C., Jbabdi, S., Glasser, M.F., Andersson, J., Burgess, G.C., Harms, M.P., Smith, S.M., Van Essen, D.C., Jenkinson, M.: MSM: A new flexible framework for Multimodal Surface Matching. Neuroimage 100, 414–426 (2014)
13. Shen, X., Tokoglu, F., Papademetris, X., Constable, R.: Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. NeuroImage 82, 403 – 415 (2013)
14. Sporns, O.: The human connectome: a complex network. Ann NY Acad Sci 1224, 109–125 (2011)
15. Yuan, J., Bae, E., Tai, X.C., Boykov, Y.: A continuous max-flow approach to potts model. In: K. Daniilidis, P. Maragos, and N. Paragios: ECCV 2010, Part VI, LNCS 6316, pp. 379-392. Springer, Heidelberg (2010)