

GraMPa: Graph-based Multi-modal Parcellation of the Cortex using Fusion Moves

Sarah Parisot^{1*}, Ben Glocker¹, Markus D. Schirmer², and Daniel Rueckert¹

¹ Biomedical Image Analysis Group, Imperial College London, UK

² Stroke Group, Massachusetts General Hospital, Harvard Medical School, USA

Abstract. Parcellating the brain into a set of distinct subregions is an essential step for building and studying brain connectivity networks. Connectivity driven parcellation is a natural approach, but suffers from the lack of reliability of connectivity data. Combining modalities in the parcellation task has the potential to yield more robust parcellations, yet hasn't been explored much. In this paper, we propose a graph-based multi-modal parcellation method that iteratively computes a set of modality specific parcellations and merges them using the concept of fusion moves. The merged parcellation initialises the next iteration, forcing all modalities to converge towards a set of mutually informed parcellations. Experiments on 50 subjects of the Human Connectome Project database show that the multi-modal setting yields parcels that are more reproducible and more representative of the underlying connectivity.

1 Introduction

Advances in neuroimaging have provided a tremendous amount of in-vivo information about the brain's organisation. In particular, studying connectivity networks from a network theoretic perspective has shown great potential and receives growing interest [15]. An essential step for carrying out network analysis is the definition of the nodes of the networks, as the high dimensionality of the data acquired at the voxel or vertex level hinders tractable and meaningful analysis. Nodes are typically obtained by parcellating the brain into a set of subregions. Connectivity driven parcellation is a natural approach, as each network node will comprise regions with similar connectivity profiles. Cortical parcellation can also be approached from the perspective of identifying functionally specified cortical areas, an elusive objective that has been ongoing for over a century [4]. Cortical areas can be defined based on their local microstructure, connectivity and function. This suggests that a single modality is not sufficient for identifying these areas [4]. Brain parcellation has notably been carried out using myelin [5], diffusion MRI (dMRI) and tractography [10, 12] and functional MRI (fMRI) data [2, 3]. Yet, these modalities suffer from important drawbacks

* 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.

that cannot be addressed in a mono-modal setting. dMRI is prone to false negatives and biased with respect to the location of fibre terminations, fMRI can be noisy and prone to false positives, while myelin lacks information outside the motor area and visual cortex.

Exploring parcellations driven by several modalities could provide more robust and accurate cortical delineations. For instance, strong similarities have been observed between myelin maps and resting-state fMRI gradients [5], while functional and structural connectivity are intrinsically linked. Few methods have attempted to combine different modalities. The majority of efforts have aimed to construct a more robust fMRI connectivity matrix informed from tractography, for instance by eliminating functional connections that do not have a structural support [11]. This kind of approach however assumes a strong reliability of dMRI data and a global agreement between structural and functional connectivity. Markov Random Field (MRF) models have been applied successfully to dMRI and fMRI driven cortical parcellation tasks [6, 13]. Their main advantage is their versatility, in the sense that no restriction is made on the data term driving the parcellation scheme. As a result, the same framework can be used for parcellation tasks using different kinds of input data.

In this paper, we exploit this idea and extend the mono-modal MRF models to the multi-modal setting. We propose an iterative approach where each iteration computes a set of parcellations driven by a single modality. These parcellations are subsequently merged based on each modality’s local reliability using fusion moves [9]. The merged parcellation initialises the next iteration, forcing the different modalities to converge towards a set of mutually informed parcellations. The method was tested on the Human Connectome Project (HCP) database using myelin maps, and fMRI and dMRI data. Focusing on fMRI parcellation, our experiments show that the multi-modal setting yields parcels that are more reproducible and more representative of the underlying connectivity.

2 Methods

As illustrated in Fig. 1, our method alternates between generating a set of N modality specific parcellations (Sec. 2.1), and the construction of a joint parcellation that initialises the next iteration (Sec. 2.2 and 2.3).

2.1 Modality Specific Markov Random Field Formulation

Considering a set of N aligned modalities, we represent the brain’s cortical surface as a triangular mesh $\mathcal{M} = \{\mathcal{V}, \mathcal{E}\}$, where each vertex is associated with modality specific data (e.g. fMRI timeseries). We cast the parcellation task as a labelling problem where, for each modality, we aim to assign a label $l_{\mathbf{v}}^{mod}$ to each vertex \mathbf{v} , $l_{\mathbf{v}}^{mod} \in \mathcal{L} = \llbracket 1, K \rrbracket$, where K is the number of sought parcels, and each label corresponds to a parcel assignment. This can be done in the mono-modal setting following the coarse to fine iterative approach proposed in [13]. Given an initial parcellation, each update consists of: 1) defining a parcel specific property

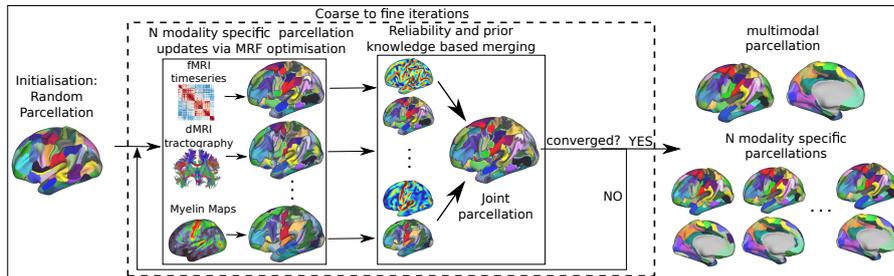


Fig. 1: Overview of the proposed iterative method. Each iteration updates an initial parcellation into a set of modality specific parcellations using an MRF model. The parcellations are then merged based on the modalities relative influences into a multi-modal parcellation, which initialises the next iteration.

based on the current parcellation and 2) updating the parcellation by optimising an MRF model. This corresponds to minimising the following energy:

$$E(\mathbf{l}^{mod}) = \sum_{\mathbf{v} \in \mathcal{V}} D_{\mathbf{v}}^{mod}(\mathbf{l}_{\mathbf{v}}^{mod}) + \beta \sum_{\mathbf{v} \in \mathcal{V}} \sum_{\mathbf{w} \in \mathcal{N}(\mathbf{v})} U_{\mathbf{v}, \mathbf{w}}(\mathbf{l}_{\mathbf{v}}^{mod}, \mathbf{l}_{\mathbf{w}}^{mod}) \quad (1)$$

The pairwise term $U_{\mathbf{v}, \mathbf{w}}(\mathbf{l}_{\mathbf{v}}^{mod}, \mathbf{l}_{\mathbf{w}}^{mod})$ acts as a smoothing prior and is designed for all modalities as a Potts model, that penalises assigning different labels to neighbours with a constant cost. The unary cost $D_{\mathbf{v}}^{mod}(\mathbf{l}_{\mathbf{v}}^{mod})$ defines the likelihood of assigning vertex \mathbf{v} to a specific parcel. We consider three different kinds of modalities that can be used to drive the parcellation scheme and associate each modality to a specific unary cost $D_{\mathbf{v}}^{mod}(\mathbf{l}_{\mathbf{v}}^{mod})$ and parcel property.

Diffusion MRI. We adopt the unary cost used in [13] which showed good performance for tractography data. The unary cost computes the Pearson’s correlation coefficient between the vertices’ tractography connectivity profiles and the parcel centres. Parcel centres are the vertices that have the highest correlation to the rest of their currently assigned parcel.

Functional MRI. Each vertex is associated with fMRI timeseries. Similarly to dMRI, we obtain a parcel centre by maximising the within parcel timeseries correlation. Due to the low SNR of fMRI data, we follow the approach proposed in [6] and compute a representative average timeseries for each label. We average the timeseries of the N closest vertices (in terms of shortest path) to the parcel centre that are in the same parcel. The unary cost for a given vertex is then the correlation between the vertex’s timeseries and the parcel’s average timeseries.

Non connectivity data (e.g. myelin maps). Each vertex is associated with a scalar value. We use unary costs and parcel properties inspired from MRF-based image segmentation. The parcel centre is defined as the geometric centre of the parcel (obtained by erosion of the parcel), and is assigned the average vertex value within the parcel. The unary cost is the shortest path on the mesh \mathcal{M} between the centre and the vertex \mathbf{v} , where the edges are weighted by the

absolute difference between the values of the parcel centre and the vertex. The use of the shortest path allows the parcels to be contiguous.

This set up yields independent mono-modal parcellations for a set of modalities, where each modality is parcellated independently and subject to modality specific noise. In particular, parcellations may not be coherent, making comparisons difficult. Forcing all parcellations to agree (e.g. by using inter-modality pairwise costs) would not be adequate as the modalities are not expected to provide information that agrees across the whole cortex [4] and modalities should not have the same local importance. We propose to force the parcellations to converge towards a set of coherent parcellations by initialising each iteration with a multi-modal parcellation computed by merging the individual parcellations.

2.2 Merging Modalities with Fusion Moves

Modelling the merging problem as an MRF of energy $E^m(\mathbf{l}) = \sum_{\mathbf{v} \in \mathcal{V}} U_{\mathbf{v}}^m(l_{\mathbf{v}}) + \sum_{\mathbf{v} \in \mathcal{V}} \sum_{\mathbf{w} \in \mathcal{N}(\mathbf{v})} U_{\mathbf{v}, \mathbf{w}}(l_{\mathbf{v}}, l_{\mathbf{w}})$ allows to tailor the model according to the modalities considered, their interaction, and the quality of the data. We propose to define the unary cost $U_{\mathbf{v}}^m(l_{\mathbf{v}})$ based on the mono-modal labellings and how important and reliable a modality is locally. For instance, $U_{\mathbf{v}}^m(l_{\mathbf{v}})$ can give lower costs to labels selected by the most reliable modalities. Reliability could be defined based on prior knowledge, or in a data-driven way using segmentation uncertainties obtained from the mono-modal MRFs [7]. $U_{\mathbf{v}}^m(l_{\mathbf{v}})$ could also integrate parcel boundaries obtained from a different source (e.g. registered atlases or expert annotations). An example of unary costs is presented in section 2.3.

Each mono-modal solution can be seen as a suboptimal solution to the multi-modal problem, where the modality is given more importance compared to the others. This makes the concept of fusion moves particularly well suited to solving our problem. Fusion moves [9] cast the task of combining an ensemble of suboptimal solutions as a set of simple binary MRF subproblems. In each binary optimisation (a fusion move), we consider two suboptimal labellings \mathbf{l}_1 and $\mathbf{l}_2 \in \mathcal{L}^{\mathcal{V}}$. We seek to label the cortical mesh \mathcal{M} with a binary label $\mathbf{b} \in \{0, 1\}^{\mathcal{V}}$ so as to obtain a combination \mathbf{l}^c defined as $l^c(\mathbf{b}) = \mathbf{l}_1 \circ (1 - \mathbf{b}) + \mathbf{l}_2 \circ \mathbf{b}$, where \circ is the Hadamard product. The fusion move is carried out by minimising the binary MRF energy $E^b(\mathbf{b}) = E^m(l^c(\mathbf{b}))$. The next fusion move considers a new modality and the current combined parcellation \mathbf{l}^c as the suboptimal labellings to merge. Fusion moves are repeated for all modalities until the combined labelling is no longer updated. The joint parcellation then initialises the next resolution.

2.3 Application to Multi-modality Informed rs-fMRI Parcellation

We propose to apply the multi-modal framework to increase the robustness of resting-state fMRI (rs-fMRI) driven parcellations. rs-fMRI provides useful information all around the cortex. However, it suffers from a low SNR which can significantly impact the obtained parcels and their reproducibility. Introducing multi-modal information in the parcellation scheme could be a way of addressing this issue. We consider combining rs-fMRI, dMRI and myelin maps due to

their expected similarities. We define the merging unary costs based on how informative the modalities are, and from prior knowledge of their weaknesses: $U_v^m(l_v) = \min_{mod \in [1, N]} (1 - \alpha^{mod} \delta(l_v, l_v^{mod}))$, where $\alpha^{mod} \in [0, 1]^V$ are costs that describe the local reliability of all modalities, $\delta(., .)$ is the Kronecker delta function and l^{mod} are the labellings obtained from mono-modal parcellations. All α^{mod} costs are rescaled between 0 and 1.

Because of rs-fMRI’s low SNR, the joint parcellation should be influenced by the other modalities when they are reliable. We therefore assign a uniform reliability $\alpha^{fMRI} = 0.5$ to rs-fMRI of across the whole cortex. Myelin maps should influence the merged parcellation in regions where strong variations of myelination are observed. We therefore define the myelin cost as the gradient of the pre-smoothed myelin maps (see Fig. 3b). Finally, dMRI tractography suffers from a gyral bias: tractography streamlines tend to terminate preferentially in gyri [17]. This bias influences the boundaries of dMRI driven parcellations that tend to align with cortical folding. To evaluate which vertices are impacted by this bias, we compute for each vertex \mathbf{v} the ratio of the number of fibres that terminate at \mathbf{v} over the number of connections obtained by sending streamlines from \mathbf{v} . As shown in Fig. 3a, this measurement supports the gyral bias theory as the resulting map agrees with cortical folding patterns. Using this measure as a unary cost prevents the vertices affected by the bias to influence the joint parcellation. In this setting, dMRI will have little influence on parcel boundaries and essentially act as a smoothing prior, indicating which vertices should be in the same parcel. As a result, we expect the converged joint parcellation to be similar to the rs-fMRI parcellation.

3 Results

Evaluation of cortical parcellations is challenging due to the absence of ground truth. Our proposed evaluation has two main objectives: evaluate i) whether multi-modality increases the robustness of the parcellation method, ii) how well the parcellations reflect the underlying connectivity. Since our application is tailored to construct more reliable rs-fMRI parcellations, we focus our evaluation on this modality. we evaluate the impact of multi-modal information by comparing the mono-modal rs-fMRI driven parcellation to the joint and individual rs-fMRI parcellations obtained using our Graph-based Multi-modal Parcellation (GraMPa) method. We tested GraMPa on 50 randomly selected subjects (left hemisphere) of the HCP database (S500 release) and used the HCP’s preprocessed fMRI and dMRI data, and myelin maps. dMRI tractography connectivity profiles are obtained using FSL’s bedpostX and probtrackX [1]. 5000 streamlines are sampled from each mesh vertex. We perform rs-fMRI driven parcellation using timeseries from a 30 minutes acquisition. Evaluation is performed on a second independent 30 minutes acquisition to test the method’s robustness. The MRF’s smoothness parameter β is set heuristically to 0.3. Modality specific MRFs were optimised using fastPD [8] due to its speed, while fusion moves were optimised using QPBO [14] because of asymmetric pairwise costs. Several MRF optimisation

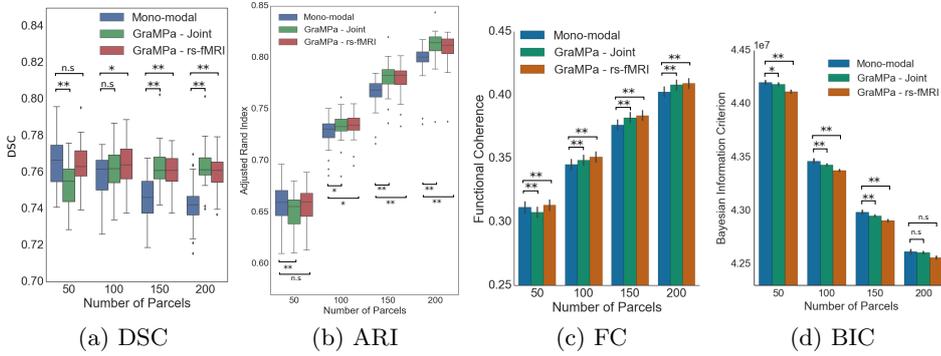


Fig. 2: Quantitative evaluation measures. From left to right in each figure, we compare the mono-modal parcellations to the merged and the multi-modality guided rs-fMRI parcellations. Lower BIC values (d) are better. Paired t-test results are shown as non-significant (n.s), $p < 0.05$ (*) and $p < 0.001$ (**).

algorithms were tested with very little impact on the obtained parcellations. We tested the reproducibility with respect to initialisation using 10 random initialisations constructed using Poisson Disc Sampling. Parcellations were computed for four different resolutions (50, 100, 150 and 200 labels). All measures are computed for all initialisations and subjects. Reproducibility is evaluated using the Adjusted Rand Index (ARI) [4] and the modified Dice Score Coefficient (DSC) [13] that allows merging very similar parcels. ARI is a measure from probability theory that assesses the statistical dependence between two clustering solutions. It takes values between -1 and 1, where 1 means the clusterings are identical. Figures 2a and 2b show comparative boxplots of the two measures between GraMPa and the mono-modal approach. We can see that most configurations are more reproducible. Results are significant ($p < 0.001$) for the two largest resolutions. The lower performance for 50 parcels could indicate that it is difficult to obtain large smooth parcels while agreeing with all reliability maps. Our parcellations' agreement with the underlying structure is evaluated by i) computing the average functional coherence (FC) [6] and ii) evaluating the agreement with task fMRI activation maps (obtained using FSL's standard tools) using the Bayesian Information Criterion (BIC) [16]. FC evaluates the average correlation between a parcel's average timeseries and the timeseries of all vertices in the same parcel. In order to avoid introducing a size bias, very small parcels are ignored from the computation. For each parcel, BIC evaluates how well it is possible to fit a probabilistic model of the concatenated task activation maps of all 50 subjects. As shown in Fig. 2c and 2d, GraMPa yields better results for both measures. Results are significant ($p < 0.001$) for most configurations. Finally, Fig. 3c-g visually compares parcels boundaries with Brodmann and myelin maps and the average task activation maps over all 50 subjects. We can see that GraMPa parcellations have a stronger agreement with task activations boundaries.

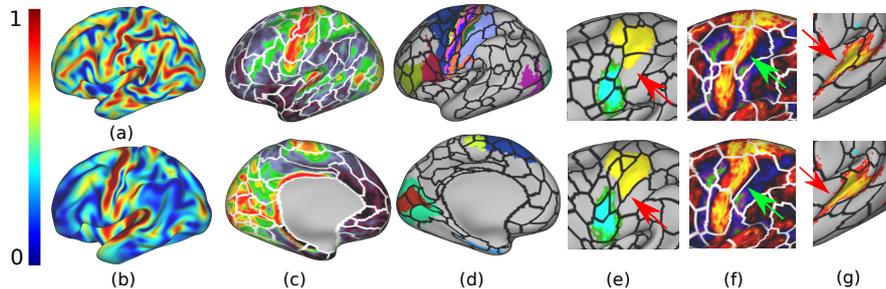


Fig. 3: Visual results for randomly selected subjects. (a) dMRI and (b) Myelin reliability maps. (c,d) Overlap between the boundaries of the multi-modal parcellation and (c) myelin maps and (d) Brodmann areas. (e,f,g) Comparative overlap between rs-fMRI parcellations boundaries and t-fMRI activation maps. Top row: mono-modal parcellations, bottom row: GraMPa rs-fMRI parcellations. (e,f) Motor task, (g) Language task. Coloured arrow indicate striking examples.

4 Discussion

In this paper, we proposed a general graph-based framework which provides modality specific coherent parcellations, as well as a multi-modal parcellation that merges modalities based on their reliabilities. We propose an application to the construction of more reliable rs-fMRI parcellations through the introduction of multi-modal information from structural connectivity and myelin maps. Our experiments show that GraMPa’s parcellations are more robust and more representative of the underlying structure. One of the main advantages of the proposed framework is its flexibility. It can be tailored for a specific set of modalities and issues associated with a particular acquisition process. It is also easy to integrate prior knowledge both in designing the modalities’ reliability maps and through the introduction of known reliable boundaries defined as a new locally reliable modality. Another possibility is to design a fully data-driven fusion move step. Local segmentation uncertainties could be estimated for each modality after each MRF optimisation using min-marginal energies [7].

Furthermore, our model alleviates the need to match the different modalities’ unary costs and does not limit the number of modalities considered. The method could be extended to other multi-modal segmentation tasks. It could prove particularly well-suited to group-wise parcellation, where each subject would be assimilated to a modality and the fusion would be driven by group consistency measures. It would have the potential of handling very large groups, as subjects don’t have to be considered simultaneously. The method could similarly be used to merge MRF parcellations obtained from a large set of initialisations. Many challenges remain associated with the multi-modal parcellation task. fMRI and dMRI are currently the best way of measuring in vivo connectivity, but remain very indirect measurements and can be unreliable. In addition, multi-modal analyses would benefit from a stronger knowledge of the modalities interactions and

similarities. Finally, using our parcellations in a clinical context requires the development of robust methods for analysing the obtained connectivity networks, while parcellation of diseased subjects may be associated with new challenges.

References

- 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)
- 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)
- 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)
- Eickhoff, S.B., Thirion, B., Varoquaux, G., Bzdok, D.: Connectivity-based parcellation: Critique and implications. *Hum brain Mapp* 36(12), 4771–4792 (2015)
- Glasser, M.F., Van Essen, D.C.: Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI. *J Neurosci* 31(32), 11597–11616 (2011)
- 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)
- Kohli, P., Torr, P.H.: Measuring uncertainty in graph cut solutions. *Computer Vision and Image Understanding* 112(1), 30–38 (2008)
- Komodakis, N., Tziritas, G.: Approximate labeling via graph cuts based on linear programming. *IEEE Trans. Pattern Anal. Mach. Intell.* 29(8), 1436–1453 (2007)
- Lempitsky, V., Rother, C., Roth, S., Blake, A.: Fusion moves for markov random field optimization. *IEEE Trans. PAMI* 32(8), 1392–1405 (2010)
- Moreno-Dominguez, D., Anwender, A., Knösche, T.R.: A hierarchical method for whole-brain connectivity-based parcellation. *Hum Brain Mapp* 35, 5000–5025 (2014)
- Ng, B., Varoquaux, G., Poline, J.B., Thirion, B.: Implications of inconsistencies between fmri and dmri on multimodal connectivity estimation. In: *MICCAI*, pp. 652–659. Springer (2013)
- Parisot, S., Arslan, S., Passerat-Palmbach, J., Wells III, W.M., Rueckert, D.: Tractography-Driven Groupwise Multi-scale Parcellation of the Cortex. In: *IPMI*. pp. 600–612. Springer (2015)
- Parisot, S., Rajchl, M., Passerat-Palmbach, J., Rueckert, D.: A continuous flow-maximisation approach to connectivity-driven cortical parcellation. In: *MICCAI*, pp. 165–172. Springer (2015)
- Rother, C., Kolmogorov, V., Lempitsky, V., Szummer, M.: Optimizing binary mrfs via extended roof duality. In: *CVPR*. pp. 1–8. IEEE (2007)
- Sporns, O.: The human connectome: a complex network. *Ann NY Acad Sci* 1224, 109–125 (2011)
- Thirion, B., Varoquaux, G., Dohmatob, E., Poline, J.B.: Which fmri clustering gives good brain parcellations? *Frontiers in neuroscience* 8(167), 13 (2014)
- Van Essen, D.C., Jbabdi, S., Sotiropoulos, S.N., Chen, C., et al.: Mapping Connections in Humans and Non-Human Primates: Aspirations and Challenges for Diffusion Imaging. *Diffusion MRI* pp. 337–358 (2013)