Discrete Optimisation for Group-wise Cortical Surface Atlasing

Emma C. Robinson, Ben Glocker, Martin Rajchl, Daniel Rueckert
Department of Computing,
Imperial College London,
South Kensington Campus,
London. SW7 2AZ,
emma.robinson05@imperial.ac.uk

Abstract

This paper presents a novel method for cortical surface atlasing. Group-wise registration is performed through a discrete optimisation framework that seeks to simultaneously improve pairwise correspondences between surface feature sets, whilst minimising a global cost relating to the rank of the feature matrix. It is assumed that when fully aligned, features will be highly linearly correlated, and thus have low rank. The framework is regularised through use of multi-resolution control point grids and higher-order smoothness terms, calculated by considering deformation strain for displacements of triplets of points. Accordingly the discrete framework is solved through high-order clique reduction. The framework is tested on cortical folding based alignment, using data from the Human Connectome Project. Preliminary results indicate that group-wise alignment improves folding correspondences, relative to registration between all pair-wise combinations, and registration to a global average template.

1. Introduction

Spatial normalisation of cortical surfaces is a vitally important processing step for a wide range of neuroimaging studies, including population level comparisons of cortical shape, task activations, or functional and structural connectivity. In particular, surface based processing and co-registration of functional MRI (fMRI signals) has been shown to significantly improve the detection of brain activations by boosting signal-to-noise; impacting on the sensitivity of task-based localisation studies and improving cortical parcellation frameworks.

Current methods for surface atlasing typically perform registration to a single population average template. For example, the popular FreeSurfer framework for spherical projection and alignment of cortical surfaces generates an initial template through selection of one subject to act as an initial reference. This is the target of a stage of highly constrained, cortical folding driven, non-rigid alignment, after which point all surfaces and features are averaged and the process of registration and averaging is iterated until convergence. The FreeSurfer (fsaverage) template is widely adopted, used by [28] within a diffeomorphic alignment framework, and in [26] it is adapted to force hemispheric symmetry. Further, [5, 22] adapt the FreeSurfer framework to align fMRI timeseries and functional connectivity features respectively. In each case the registration target is the population average.

However, there is growing evidence that brain topology is not consistent, even across healthy populations of fixed age. For example, major cortical folding patterns, such as the cingulate, have been shown to vary with regards to the number of folds and degree of branching [26]. Further, the relative placement of different functional activations has been shown to vary [10]. Methods that register to a common target, bias the resulting average to the reference anatomy, and ensure only features that are common across the whole dataset are consistently aligned.

For this reason, several methods have been proposed to reduce bias by performing registration for all pairwise combinations [3, 16, 18, 27]. For example, [3] generate neonatal cortical surface atlases via pairwise alignment of cortical folding data using spherical alignment and average all pairwise transformations, according to the approach in [23]. [17] improves on an initial pairwise spherical template, estimated from spherical harmonic alignment of landmark curves [18], by minimising ensemble entropy across the group. Whereas, [27] performs group-wise registration in two stages, first performing pairwise spectral embeddings of all cortical surface shapes using the approach in [16]; then performing a joint embedding using correspondences from the pairwise stage to anchor the embedding.

Spectral embedding methods are also used for functional alignment by [10, 14]. These methods have many

advantages in terms of being fast and (under certain constraints) diffeomorphic [16]. However, because of the lack of ground-truth understanding of how cortical folding relates to functional connectivity [2] they are typically only applied to one type of feature set at a time. Further, embedding implies no spatial regularisation of the solution; which can lead to spatially distant points on the cortical surface being related through long distance correlations between features. In [27] this is addressed through edge-based smoothing of the pair-wise connections.

In this paper we take a different approach and take inspiration from volumetric image registration frameworks that approach simultaneous group-wise alignment within a spatially constrained deformation model [4, 25]. In particular, adopting the discrete optimisation approach proposed by Sotiras et al [24] in which group-wise alignment of 2D image slices is implemented using a global entropy cost. By contrast, we propose a new group-wise deformation model that seeks to align features representative of the cortical surfaces (such as folding or function), whilst simultaneously minimising the rank of the data set; the goal is to find the most compact representation, even in the face of structural or functional variability across data sets.

Registration is implemented using a modification of the discrete spherical alignment approach MSM [20], inspired by DROP [8] a discrete optimisation approach for Free-Form Deformations of volumetric image grids. Like DROP, MSM has advantages in terms of: reduced sensitivity to local minima [9]; and a modular optimisation framework which means that any combination of similarity term and regularisation framework can be used. This has advantages for cortical alignment as it means any type of features or combination of features, that can be projected to the surface, can be used to drive the alignment; provided an appropriate data cost can be found. Accordingly MSM has been applied to a wide variety of feature sets including, cortical folding [3], myelin [20] resting state functional MRI (fMRI) [20], retinotopic maps [1].

The full group-wise framework incorporates three terms: a pairwise cross-correlation term to encourage alignment of common features on the surfaces of subjects whose topology is constant; a quartet-based global low rank penalty to encourage correspondences between topologically inconsistent feature sets; and a triplet-based smoothness term to encourage smoothness of the warp. In the next section we describe the discrete framework. As proof of concept, the proposed model is tested on cortical folding driven alignment of the 10 unrelated subjects available from the Human Connectome Project.

2. Group-wise Discrete Registration Model

Within the proposed group-wise surface alignment framework there exist n cortical surfaces S_i each projected to a sphere through the procedures outlined in [7]. These surfaces are associated with I_i feature sets which can be any combination of features describing cortical folding, functional activations, correlates of cytoarchitecture, or structural connectivity. The optimisation seeks transformations T_i, which simultaneously map datasets $\{S_i{:}I_i\}$ onto a common reference frame. The reference mesh is represented by a regularly sampled sphere S_0, with counterpart low resolution reference control grid G_0.

The multi-resolution deformation model is built upon use of n sets of triangulated, spherical, control-point grids G_i (one for each subject: Figure 1). Displacements of the individual control point grids follow the framework described in [20]. At each iteration, the points $\mathbf{p}_i$ are constrained to move to a new, discrete set of locations on the surface. These are determined through regular sampling of the surface following the scheme in [20]. The associated displacements are described by rotation matrix $\mathbf{R}^{l_{\mathbf{P}_i}}$, where end point is indexed using a label $l(\mathbf{p_i})$, and $\mathbf{R}^{l_{\mathbf{P}_i}}$ rotates each control point $\mathbf{p}_i$ into position over the label. From this the optimal deformation is then solved as a multilabeling problem using discrete optimisation [20, 8]

Registration is optimised so as to simultaneously improve similarity for all pairwise comparisons between feature sets I_i^k, I_j^k. These represent patches of data at the respective position k on the low resolution reference G_0, where all features are compared following barycentric resampling to the reference grid S_0. Transformations for each point $\mathbf{x}_i$ from S_i are composed as $T_i^t(\mathbf{x}_i) = \mathbf{x}_i + T_i^{t-1}(\mathbf{x}_i) + D_i(\mathbf{x}_i)$ where $D_i(\mathbf{x}_i) = \sum_{\mathbf{p}_i \in F_i(\mathbf{x_i})} \eta(|\mathbf{x}_i - \mathbf{p}_i|)\mathbf{R}^{l_{\mathbf{P}_i}}\mathbf{p}_i$, such that $\eta(.)$ is a barycentric interpolation function that constrains each data point to move with the control-point grid triplet $F_i(\mathbf{x}_i)$.

The optimisation has three terms: a pairwise similarity term and two penalty terms seeking a) minimisation of the rank of the data, through reduction of nuclear norm for the full group data matrix ($\mathbf{M}^k$); and b) regularisation of each control grid warp by penalising the strain of the deformation for each control point grid face triplet. The full framework is described in equation 1, with each term being described in more detail below:

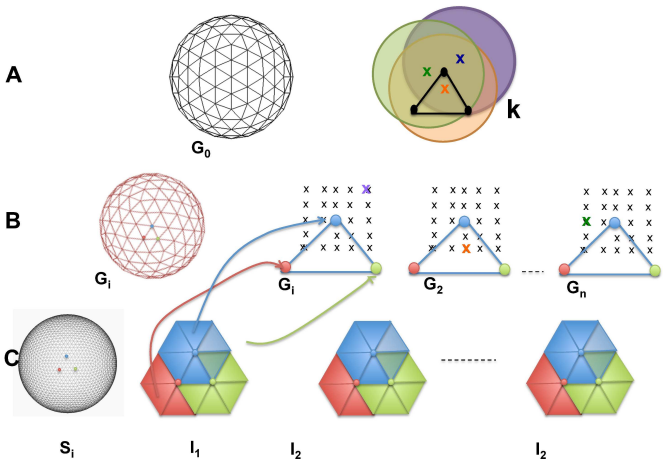

Figure 1. An overview of the group registration framework: A) Data I_i associated with each high resolution spherical mesh S_i is subdivided according to the control point face it falls within. B) Each control-point (red, green and blue dots) samples data from all neighbouring faces, and chooses from a finite choice of end points (black crosses) for each displacement. C) Between-subject control-point neighbourhoods are found by tracking closest points relative to each reference grid position k.

$$E(T_1 \circ G_1, ...T_n \circ G_n) = \underbrace{\sum_{i \in n} \sum_{\mathbf{p} \in G_i} \sum_{\substack{\mathbf{q} \in \\ (N(\mathbf{p}) \setminus G_i)}} V_{\mathbf{pq}}(l_\mathbf{p}, l_\mathbf{q})}_{\text{pairwise between mesh similarity term}}$$
$$+ \underbrace{\sum_{i \in n} \sum_{p \in G_i} \sum_{\substack{\mathbf{qr} \in \\ (N(\mathbf{p}) \cap G_i)}} \lambda V_{\mathbf{pqr}}(l_\mathbf{p}, l_\mathbf{q}, l_\mathbf{r})}_{\text{triplet warp regularisation term}}$$
$$+ \underbrace{\sum_{\mathbf{p} \in G_i} \sum_{\substack{\mathbf{qrs} \in \\ (N(p) \setminus G_i)}} \alpha V_{\mathbf{pqrs}}(l_\mathbf{p}, l_\mathbf{q}, l_\mathbf{r}, l_\mathbf{s})}_{\text{low rank quartet similarity term}} \quad (1)$$

2.1. Pairwise Similarity Comparisons

The first term in the objective function seeks to optimally match surface features between all pairwise combinations of surfaces. It is estimated from cross-correlations between overlapping patches of data I_i, I_j relative to position k in the reference space:

$$V_{\mathbf{p}_i^k \mathbf{q}_j^k} = \sum_{\mathbf{x}_i \in F(\mathbf{p})} \sum_{\mathbf{x}_j \in F(\mathbf{q})} \Phi(|T_i(\mathbf{x}_i) - T_j(\mathbf{x}_j)|) \times \frac{(I_i(T_i(\mathbf{x}_i)) - \bar{I_i}^k)(I_j(T_j(\mathbf{x}_j)) - \bar{I_j}^k)}{\sigma_i^k \sigma_j^k} \quad (2)$$

Each patch I_i^k represents the overlap between surface points $(\mathbf{x}_i)$ and control-point-grid triplets $F(\mathbf{p_i})$. Between mesh neighbours $\mathbf{q} \in N(\mathbf{p}_i) \setminus G_i)$ are found by tracking nearest points relative to the template mesh G_0 grid position k (Figure 1). $\Phi(|T_i(\mathbf{x}_i) - T_j(\mathbf{x}_j)|)$ represents a Dirac-driven function whose role is to define which pixels correspond to the same position at the reference pose defined as in [24].

2.2. Global Low Rank Cost

Let $\mathbf{M_{pqrs}}$ be a matrix whose columns represent patches of features for each subject centred at position k in the reference. We expect these feature sets when aligned to exhibit high linear correlation. Thus in the case of perfect alignment of patches with common topology, we expect the matrix to have single rank. On the other hand if the population can be expressed in terms of a series of subgroups, each with different topology then the data may have rank $1 < r << n$. This term therefore supplements pairwise comparisons in the case where feature sets are inconsistent.

The global cost is estimated through simultaneous displacement of quartets of control points $\mathbf{qrs} \in N(\mathbf{p}) \setminus G_i$. This is motivated by the assumption that simultaneous optimisation of displacements across groups of surfaces will enable faster convergence. However, increasing clique sizes increases computation overhead as more get added to the pairwise QPBO optimisation. For this reason, quartet terms were considered a reasonable trade off between run time and accurate approximation of rank at time t.

Rank is approximated through the nuclear norm, or sum of singular values $||\mathbf{M}_{pqrs}||_* := \sum_i Tr(\sqrt{\mathbf{M}_{pqrs}^T \mathbf{M}_{pqrs}})$ of the feature matrix $\mathbf{M_{pqrs}}$:

$$V_{\mathbf{pqrs}}(l_\mathbf{p}, l_\mathbf{q}, l_\mathbf{r}, l_\mathbf{s}) = \alpha ||\mathbf{M}_{pqrs}||_* \quad (3)$$

Where, use of the nuclear norm as a relaxation of the number of eigenvalues (or rank) of the matrix is commonly used approximation in convex optimisation methods such as [19]. Here $\mathbf{M}_{pqrs}$ is formed from the overlap between patches $I_1^k \cap I_2^k \cap I_2^k ... \cap I_n^k$ relative to reference position k. Sample displacement of data associated with control points $\mathbf{qrs}$ for labels $l_\mathbf{p}, l_\mathbf{q}, l_\mathbf{r}, l_\mathbf{s}$ at time t is compared with data resampled from the remaining subjects at time $t-1$.

2.3. Regularisation

Taking inspiration from [13] we propose a penalty term that constrains the strain-energy density (W) of locally affine warps $\mathbf{F_{pqr}}$ defined between the triangular faces for the original and deformed control-point grid meshes:

$$V_{\mathbf{pqr}}(l_\mathbf{p}, l_\mathbf{q}, l_\mathbf{r}) := \lambda W_{pqr}^2 = \lambda \frac{1}{2}(\mu(I_1^* - 3) + \kappa(J-1)^2)^2 \quad (4)$$

Here μ is the shear modulus and κ is the bulk modulus. These were set to a ratio of 2:1 based on parameter optimisation over pairwise registration combinations for 6 independent subjects.

Here, $I_1^* = I_1 I_3^{-1/3}$. I_1 and I_3 are strain invariants given by:

$$I_1 = trace(\mathbf{F}_{pqr}^T.\mathbf{F}_{pqr}) \quad (5)$$

$$I_3 = J^2 = det(\mathbf{F}_{pqr}^T.\mathbf{F}_{pqr}) \quad (6)$$

2.4. Optimisation

In order to account for triplet terms we adopt the approach of [11, 12]. This allows reduction of higher order terms either by: A) addition of auxiliary variables (for example by reducing a triplet to three pairwise terms [11]; or B) by reconfiguration of the polynomial form of the MRF energy, until the high-order function can be replaced by a single quadratic [12]. We present results using the latter version, known as Excludable Local Configuration (ELC). This has the advantage that, provided an ELC can be found, there is no increase in the number of pairwise terms, which has some impact computational time. In each instance the reduction is passed to the QBPO solver for optimisation [21].

3. Experiments

We test the framework using data from the Human Connectome Project. As proof of concept we run the framework on the 10 unrelated subjects, and estimate group-wise alignment for cortical folding features. We choose cortical folding as these are the most straightforward to visually interpret, and are unaffected by noise, which, by contrast, is a large contributor to the inter-subject variances differences observed for functional activations. However, the framework can be directly applied to all forms of cortical surface data.

Three different frameworks for template generation were compared: 1) Registration to a group average template; 2) Pairwise registration between all surface combinations, followed by transformation averaging and resampling to a reference grid; 3) Group-wise registration using the proposed framework.

In each instance registrations were initiated through an affine alignment to the HCP group average sulcal template, and, for fair comparison, parameters of each registration framework were optimised to best match of the folds, subject to warp distortions distributions being equal. Here, distortions are calculated as $log_2 \frac{A_2}{A_1}$, where A_2 is mesh face area following registration and A_1 is mesh face area prior to registration. This parameter optimisation led to a choice of: 1) $\lambda = 0.001$; 2) $\lambda = 0.0005$; 3) $\lambda = 0.00001$, $\alpha = 0.0001$.

Results are compared through comparisons between pairwise similarities between all cortical folding feature maps, following warping and resampling to the group average surface; and variances of resulting cortical folding features following registration and resampling to the template space.

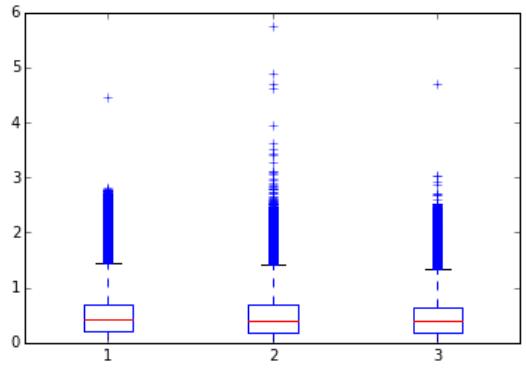

Figure 2. Areal distortion distributions for: 1) Registration to a single reference template; 2) Pairwise registrations; 3) The proposed group-wise registration framework

4. Results

Figure 2 shows box plots of the absolute values mesh face distortions across all warped surfaces. In general these meet the (maximum 3 fold) range of regional variation observed across humans and non-primates in nature [26]. Nevertheless registrations have also been optimised for folding feature similarity, and the high tails of the distributions suggest some unnaturally high distortions indicating that perfect matching would not be feasible, even in this small data set, due to the extent of variability of folding in the data.

Figure 3 demonstrates significant improvement in between subject sulcal depth map correlations for the proposed group-wise registration framework, relative to registration to the HCP sulcal depth average template, and pairwise alignments. Note pairwise correlations are estimated directly for each independent pairwise registration, prior to transformation averaging and resampling.

This improvement is also reflected in a significant improvement in sharpness for the group average template (Figure 4: top) and reduction in variance (Figure 4: bottom) of sulcal depth maps, as measured following resampling of the warped feature sets to the reference space. The group-wise average also displays greater deviation from the HCP template, particularly in areas of known folding variation (white arrows).

Table 1 shows the percentage variance shared by each of the first 5 singular values of the full ($features \times subjects$) data matrix, for each of the 3 methods, and relative to affine alignment. In each case the first eigenvalue takes the largest share of the variance, and this is significantly improved upon following group registration. However, the significant variance share of the first eigenvalue following affine alignment suggests that this particular data set does not support the existence of multiple subsets within this small group.

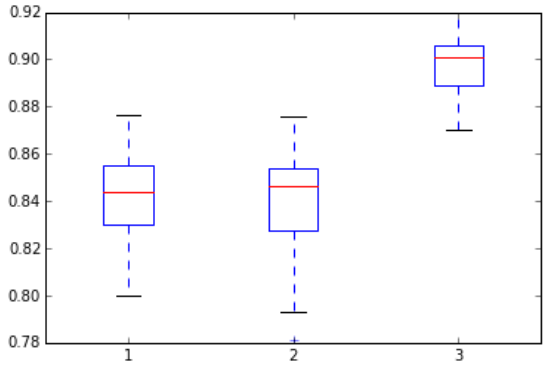

Figure 3. Between subject correlation of sulcal depth maps following: 1) Registration to a single reference template; 2) Pairwise registrations; 3) The proposed group-wise registration framework

Table 1. Singular value distributions

AFFINE	SINGLE TARGET	PAIRWISE	GROUP-WISE
0.318	0.453	0.387	0.510
0.090	0.072	0.079	0.064
0.085	0.070	0.077	0.060
0.080	0.065	0.074	0.0564
0.077	0.062	0.065	0.0550

5. Discussion

Motivated by the observation, that cortical folding patterns and functional activations display topological variations, even across healthy adult brains; this paper has proposed a preliminary framework for group-wise registration of cortices, that seeks to minimise the rank of the surface feature set. The goal is to find optimal spatial correspondence across a population even when the population can be subdivided into groups with conflicting feature topologies.

The framework has been tested on ten subjects from the HCP, where cortical folding patterns were chosen to drive alignment on account of folding maps being most straightforward to visually interpret. The proposed framework shows strong improvements in terms of pairwise similarities between the feature maps, as well as reduction in variance of cortical folding across subjects, following alignment.

One surprising outcome of this initial result is that pairwise registration with transformation averaging appears to generate results with greater variance (Figure 4). One likely reason for this is that regularisation for each method was constrained to enforce comparable levels of areal distortion. With more relaxed regularisation pairwise alignment would allow a sharper result. Further improvements in shape driven cortical alignment may be offered by diffeomorphic based approaches such as [6]. However, this is explicitly not the goal of this approach. Rather, the point being made is that outside of what common topology can be found, there exists a high degree of natural folding and functional variation in brain structure. Indeed, there is arguably no common shape to deform to.

At this time, there is not sufficient evidence to definitively support the inclusion of the global low rank term within the group registration framework, nor the choice of implementation as a quartet term. This is largely due to limited validation performed on a small data set, for which the eigenvalue distribution of the chosen feature set supports a single rank solution. The framework would need to be tested on larger populations in which clear topological differences in folding can be found.

Unfortunately, the relatively small number of subjects used within this paper was necessitated by the increased computational overhead of the patch resamplings needed for the quartet term. Neighbourhood search and interpolation is a greater overhead for cortical surface (as opposed to volumetric image grid) resampling on account of the use of irregular grids. In future this may be improved though code parallelisation [15]. Nevertheless, more extensive validation would need to be performed on the relative benefits of using a quartet term, for global cost, rather than a unary one.

The predominant limitation of the current global-rank approach is that the high degree of inter-subject variation is likely to mean that feature sets violate the assumption of linear correlation following alignment. This is likely to worsen with extrapolation to larger data sets, and functional data which has significantly more artefactual noise. In the computer vision literature, this problem has been addressed for face alignment by using a sparse and low rank decomposition approaches [19]. In the context of face recognition the sparse principal component analysis approach allows generation of a low rank solution where features of the data, that are not common across any of the data sets, are passed to a sparse error term. Optimisation minimises the sum of matrix nuclear norm and regularised sparse penalty terms. Future work will incorporate this approach.

In conclusion this framework shows potential for improving the global compactness of any group-wise cortical surface atlasing approach. Much greater validation of the approach is necessary, specifically with regards to comparison against other methods and appropriate experimental justification of use of high-order costs. In future the framework has the potential to improve understand of population variability in brain topology. A significant goal is to assess whether clusters of brain structures exist, in which sub-populations share specific variants of cortical shape and functional activation topologies and clear differences can be identified between groups.

References

[1] R. O. Abdollahi, H. Kolster, M. F. Glasser, E. C. Robinson, T. S. Coalson, D. Dierker, M. Jenkinson,

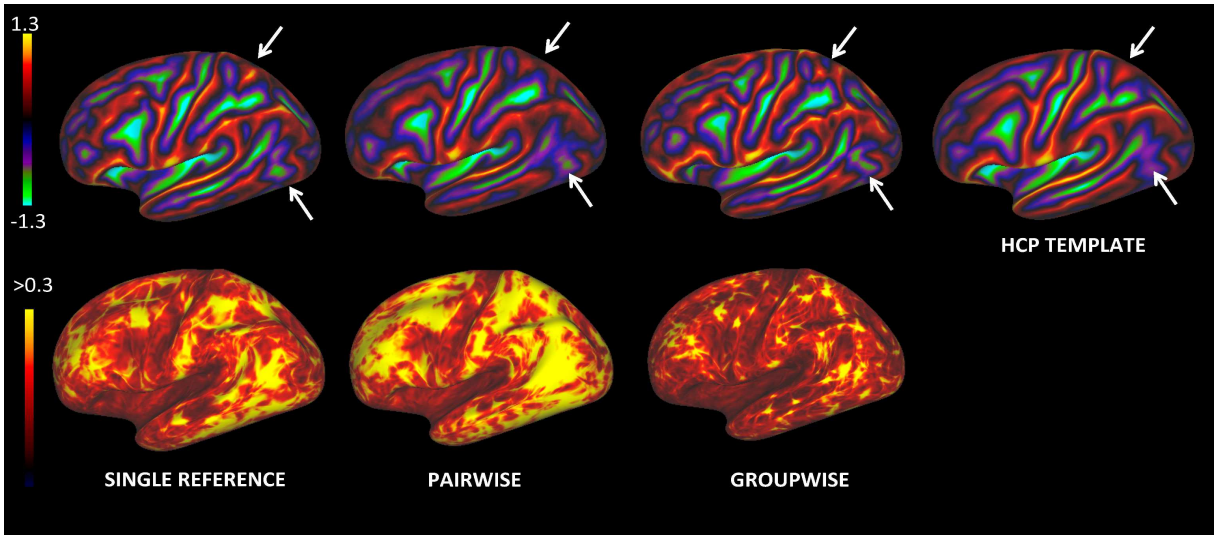

Figure 4. Template averages and variances of sulcal depth maps, following: 1) Registration to single reference target (variance distribution $\bar{\sigma}_1 = 0.195 \pm 0.096$) ; 2) Registration for all pairwise combinations, followed by transformation averaging and resampling ($\bar{\sigma}_2 = 0.237 \pm 0.118$) ; 3) The proposed group-wise registration framework ($\bar{\sigma}_3 = 0.149 \pm 0.079$)

D. C. Van Essen, and G. A. Orban. Correspondences between retinotopic areas and myelin maps in human visual cortex. *Neuroimage*, 99:509–524, 2014. 2

[2] K. Amunts, A. Schleicher, and K. Zilles. Cytoarchitecture of the cerebral cortex - more than localization. *NeuroImage*, 37(4):1061–1065, 2007. 2

[3] Authors. Construction of a neonatal surface atlas using multimodal surface matching. In *International Symposium on Biomedical Imaging (ISBI)*, Prague, April 2016. IEEE. 1, 2

[4] K. K. Bhatia, J. V. Hajnal, B. K. Puri, A. D. Edwards, and D. Rueckert. Consistent groupwise non-rigid registration for atlas construction. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 908–911. IEEE, 2004. 2

[5] B. R. Conroy, B. D. Singer, J. S. Guntupalli, P. J. Ramadge, and J. V. Haxby. Inter-subject alignment of human cortical anatomy using functional connectivity. *NeuroImage*, 81:400–411, 2013. 1

[6] S. Durrleman, X. Pennec, A. Trouvé, and N. Ayache. Statistical models of sets of curves and surfaces based on currents. *Medical image analysis*, 13(5):793–808, 2009. 5

[7] B. Fischl, M. I. Sereno, and A. M. Dale. Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, 9(2):195–207, 1999. 2

[8] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios. Dense image registration through MRFs and efficient linear programming. *Medical image analysis*, 12(6):731–741, 2008. 2

[9] B. Glocker, A. Sotiras, N. Komodakis, and N. Paragios. Deformable medical image registration: Setting the state of the art with discrete methods*. *Annual review of biomedical engineering*, 13:219–244, 2011. 2

[10] J. V. Haxby, J. S. Guntupalli, A. C. Connolly, Y. O. Halchenko, B. R. Conroy, M. I. Gobbini, M. Hanke, and P. J. Ramadge. A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron*, 72(2):404–416, 2011. 1

[11] H. Ishikawa. Higher-order clique reduction in binary graph cut. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 2993–3000. IEEE, 2009. 4

[12] H. Ishikawa. Higher-order clique reduction without auxiliary variables. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1362–1369. IEEE, 2014. 4

[13] A. K. Knutsen, Y. V. Chang, C. M. Grimm, L. Phan, L. A. Taber, and P. V. Bayly. A new method to measure cortical growth in the developing brain. *Journal of biomechanical engineering*, 132(10):101004, 2010. 3

[14] G. Langs, Y. Tie, L. Rigolo, A. Golby, and P. Golland. Functional geometry alignment and localization of brain areas. In *Advances in neural information processing systems*, pages 1225–1233, 2010. 1

[15] V. Lempitsky, C. Rother, S. Roth, and A. Blake. Fusion moves for markov random field optimiza-

tion. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(8):1392–1405, 2010. 5

[16] H. Lombaert, J. Sporring, and K. Siddiqi. Diffeomorphic spectral matching of cortical surfaces. In *Information Processing in Medical Imaging*, pages 376–389. Springer, 2013. 1, 2

[17] I. Lyu, S. H. Kim, J.-K. Seong, S. W. Yoo, A. Evans, Y. Shi, M. Sanchez, M. Niethammer, and M. A. Styner. Robust estimation of group-wise cortical correspondence with an application to macaque and human neuroimaging studies. *Frontiers in neuroscience*, 9, 2015. 1

[18] I. Lyu, S. H. Kim, J.-K. Seong, S. W. Yoo, A. C. Evans, Y. Shi, M. Sanchez, M. Niethammer, and M. A. Styner. Cortical correspondence via sulcal curve-constrained spherical registration with application to macaque studies. In *SPIE Medical Imaging*, pages 86692X–86692X. International Society for Optics and Photonics, 2013. 1

[19] Y. Peng, A. Ganesh, J. Wright, W. Xu, and Y. Ma. Rasl: Robust alignment by sparse and low-rank decomposition for linearly correlated images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(11):2233–2246, 2012. 3, 5

[20] E. C. Robinson, S. Jbabdi, M. F. Glasser, J. Andersson, G. C. Burgess, M. P. Harms, S. M. Smith, D. C. Van Essen, and M. Jenkinson. Msm: A new flexible framework for multimodal surface matching. *Neuroimage*, 100:414–426, 2014. 2

[21] C. Rother, V. Kolmogorov, V. Lempitsky, and M. Szummer. Optimizing binary mrfs via extended roof duality. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007. 4

[22] M. R. Sabuncu, B. D. Singer, B. Conroy, R. E. Bryan, P. J. Ramadge, and J. V. Haxby. Function-based intersubject alignment of human cortical anatomy. *Cerebral Cortex*, 20(1):130–140, 2010. 1

[23] A. Serag, P. Aljabar, G. Ball, S. J. Counsell, J. P. Boardman, M. A. Rutherford, A. D. Edwards, J. V. Hajnal, and D. Rueckert. Construction of a consistent high-definition spatio-temporal atlas of the developing brain using adaptive kernel regression. *Neuroimage*, 59(3):2255–2265, 2012. 1

[24] A. Sotiras, N. Komodakis, B. Glocker, J.-F. Deux, and N. Paragios. Graphical models and deformable diffeomorphic population registration using global and local metrics. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009*, pages 672–679. Springer, 2009. 2, 3

[25] C. Studholme. Simultaneous population based image alignment for template free spatial normalisation of brain anatomy. In *Biomedical Image Registration*, pages 81–90. Springer, 2003. 2

[26] D. C. Van Essen. A population-average, landmark-and surface-based (PALS) atlas of human cerebral cortex. *NeuroImage*, 28(3):635–662, 2005. 1, 4

[27] R. Wright, A. Makropoulos, V. Kyriakopoulou, P. Patkee, L. Koch, M. Rutherford, J. Hajnal, D. Rueckert, and P. Aljabar. Construction of a fetal spatio-temporal cortical surface atlas from in utero mri: Application of spectral surface matching. *NeuroImage*, 120:467–480, 2015. 1, 2

[28] B. T. Yeo, M. R. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, and P. Golland. Spherical demons: fast diffeomorphic landmark-free surface registration. *Medical Imaging, IEEE Transactions on*, 29(3):650–668, 2010. 1