



Individual parcellation of resting fMRI with a group functional connectivity prior[☆]



M. Chong^a, C. Bhushan^a, A.A. Joshi^a, S. Choi^a, J.P. Haldar^a, D.W. Shattuck^b, R.N. Spreng^c,
R.M. Leahy^{a,*}

^a Signal and Image Processing Institute, University of Southern California, Los Angeles, CA, United States

^b Ahmanson-Lovelace Brain Mapping Center, Department of Neurology, UCLA, Los Angeles, CA, United States

^c Laboratory of Brain and Cognition, Human Neuroscience Institute, Department of Human Development, Cornell University, Ithaca, NY, United States

ABSTRACT

Cortical parcellation based on resting fMRI is an important tool for investigating the functional organization and connectivity of the cerebral cortex. Group parcellation based on co-registration of anatomical images to a common atlas will inevitably result in errors in the locations of the boundaries of functional parcels when they are mapped back from the atlas to the individual. This is because areas of functional specialization vary across individuals in a manner that cannot be fully determined from the sulcal and gyral anatomy that is used for mapping between atlas and individual. We describe a method that avoids this problem by refining an initial group parcellation so that for each subject the parcel boundaries are optimized with respect to that subject's resting fMRI. Initialization with a common parcellation results in automatic correspondence between parcels across subjects. Further, by using a group sparsity constraint to model connectivity, we exploit group similarities in connectivity between parcels while optimizing their boundaries for each individual. We applied this approach with initialization on both high and low density group cortical parcellations and used resting fMRI data to refine across a group of individuals. Cross validation studies show improved homogeneity of resting activity within the refined parcels. Comparisons with task-based localizers show consistent reduction of variance of statistical parametric maps within the refined parcels relative to the group-based initialization indicating improved delineation of regions of functional specialization. This method enables a more accurate estimation of individual subject functional areas, facilitating group analysis of functional connectivity, while maintaining consistency across individuals with a standardized topological atlas.

Introduction

Parcellations of the human cerebral cortex representing cyto-, myelo- or chemo-architecture are helpful in understanding the organization and functioning of the brain (Brodmann, 1909; von Economo and Koskinas, 1925; Glasser et al., 2016). While we are not yet able to reliably identify architectonic boundaries with *in vivo* imaging, functional magnetic resonance imaging (fMRI) reveals large-scale correlations in the resting BOLD signal that can be used to identify functional networks (Biswal et al., 1995). Identification of contiguous areas of cortex that exhibit similar network connectivity can be used to define a functional parcellation (Raichle et al., 2001; Sporns et al., 2005; Smith et al., 2013; Kim et al., 2010; Eickhoff et al., 2011, 2015).

One approach to parcellation is to process each subject separately, subdividing the cortex into a set of functionally homogeneous regions

specific to that subject (Blumensath et al., 2013). Limited data available for a single individual can result in instability of the estimated boundaries (Rubinov and Sporns, 2010; Smith et al., 2011). Furthermore, it can be difficult to then determine correspondence between parcels across individuals, as required for group studies of connectivity.

Group analysis is typically performed with respect to a structural coregistration that uses sulcal and gyral features to align the cortices from each subject to a common atlas. Individual parcellations can then be pooled into a single group parcellation with respect to this anatomical correspondence using clustering or other techniques (Yeo et al., 2011; Wig et al., 2014; Craddock et al., 2012; Glasser et al., 2016). Alternative multi-graph K-way clustering methods have been proposed that simultaneously parcellate all subjects in a group (Shen et al., 2010; Arslan et al., 2015). An alternative approach uses group independent component analysis (Beckmann et al., 2009) to directly produce a single parcellation.

[☆] This work supported by NIH grants R01 NS089212, R01 NS074980 and R01 EB009048.

* Corresponding author.

E-mail address: leahy@sipi.usc.edu (R.M. Leahy).

These methods are all based on the implicit assumption that all individuals have a common functional parcellation when mapped to the common atlas. However, it is well known that architectonic boundaries and regions of functional specialization do not align precisely with macroscopic morphological features (Brodman, 1909; von Economo and Koskinas, 1925; Zilles and Amunts, 2010; Amunts et al., 1999; Penfield and Rasmussen, 1950; Honey et al., 2009; Gordon et al., 2015). As a result, a common parcellation will invariably result in errors in defining functionally specialized regions when these areas are mapped back from the atlas to the individual. It can therefore be advantageous to maintain separate parcel boundaries for each subject while still defining correspondence across subjects. Approaches to this joint individual/group parcellation problem include use of shrinkage for each individual towards a group connectivity measure (Shou et al., 2014) and individual parcellation in which connectivity is defined using weighted signal averages over each parcel where the weights are reduced in regions of high inter-subject variability (Wang et al., 2015).

Here we describe GPIIP (*Group Prior Individual Parcellation*), a novel cortical parcellation method based on resting fMRI that performs group-based parcellation but also allows individual variability across subjects in the boundaries of these parcels. As a result, we can retain a single anatomical correspondence across subjects but functional regions and their boundaries can vary from individual to individual with respect to this anatomical correspondence. We use a Bayesian formulation that combines two priors, one on the group connectivity between parcels and one that influences individual parcel boundaries. We use a Markov Random Field (MRF) (Besag, 1986, Geman and Geman, 1984) prior in the form of a Potts model (Boykov et al., 2001) to define a discrete labeling on cortex that represents the set of parcels. This spatial interaction model is defined for each individual. Group data comes into play in the prior on connectivity. We use a group sparsity measure (Cotter et al., 2005), first used in the context of fMRI connectivity estimation by Ng et al. (2012), to represent connectivity between parcels. The measure is applied to the concentration matrix, which is the inverse of the covariance in a multivariate Gaussian model of connectivity between parcels. By using group sparsity we favor a connectivity structure that is common across subjects but allows the strength of this connectivity to vary between subjects. Combining these priors with a Gaussian model of interaction between the mean time-series of each parcel we define a posterior density. Maximization of this posterior using a two-steps iteration process produces our final parcellation. We describe the method below and evaluate performance using a combination of simulation and *in vivo* data from the Human Connectome Project (Van Essen et al., 2013; Glasser et al., 2013; Smith et al., 2013; Barch et al., 2013; Woolrich et al., 2001).

Methods

Individual functional parcellation using a group sparsity prior

We assume that the cerebral cortex for each subject can be partitioned into M contiguous functional regions or parcels. This partition is based on a surface representation of cerebral cortex for all subjects, extracted and co-registered using cortical anatomy to a standard atlas with FreeSurfer (Fischl, 2012) or BrainSuite (Shattuck and Leahy, 2002). While each subject has the same M parcels with a common spatial topology, we allow the size and boundaries of these parcels to differ across subjects with respect to the common atlas space.

We represent the functional connectivity between parcels by an M dimensional Gaussian model whose likelihood is computed from the resting fMRI time series averaged over the set of vertices in each of the M parcels. The connection strength between each pair of parcels is then defined as the corresponding element of a sparsified estimate of the inverse covariance or concentration matrix (Whittaker, 2009). We use a group sparsity model to reflect our expectation that the presence or absence of connectivity between parcels is common across subjects

although the strength of these connections can vary. A Bayesian framework is used to model spatial contiguity of the cortical parcels and the group sparsity property. GPIIP computes a maximum a posterior (MAP) estimate of the individualized parcellations and their respective connectivity as described below.

To initialize GPIIP we use a common parcellation with respect to the surface atlas. We show results below for initialization based on a relatively coarse parcellation with 17-networks (Yeo et al., 2011) as well as a more fine-grained atlas with 360 regions (Glasser et al., 2016). The GPIIP algorithm then alternates between refining the parcel boundaries for each subject and updating the covariance matrices that describe connectivity between these parcels. We next describe the elements of this model and the optimization algorithm in more detail.

Denote the resting fMRI data as $D = \{D_1, D_2, \dots, D_N\}$ for a group of N subjects with $D_i = \{D_i(a, t), a = 1 \dots V, t = 1 \dots T\}$ representing the spatio-temporal data for subject i , where V denotes the number of vertices in the surface and T the number of time samples. We subdivide the vertices of the surface tessellation for each subject into M contiguous subsets to obtain parcellations $L = \{L_1, L_2, \dots, L_N\}$ which assigns a label $p \in [1 \dots M]$ to each surface vertex such that $L_i(a) = p$ for $a = 1 \dots V$.

The connectivity between the parcels is governed by an M dimensional joint Gaussian distribution for each subject with connection strength given by the elements of the concentration (inverse covariance) matrices $C = \{C_1, C_2, \dots, C_N\}$. For the case where the data are first normalized to zero mean and unit variance, as we do below, the concentration matrix is identical to the matrix of partial correlations between parcels, with a flipped sign on off-diagonal elements (Whittaker, 2009).

The posterior probability for the parcellation and functional connectivity is given by Bayes theorem as

$$P(L, C|D) = \frac{P(D|L, C)P(L, C)}{P(D)} = \frac{\prod_{i=1}^N P(D_i|L_i, C_i)P(L_i)}{P(D)}P(C), \quad (1)$$

where $P(D|L, C)$ represents the probability of the resting fMRI data given L and C and $P(L, C)$ represents the joint prior probability for the group parcellation and functional connectivity. In (1) we assume that the spatial priors for parcellation of each subject, $P(L_i)$, are independent of each other and that each is independent of the connectivity prior. The connectivity prior, $P(C)$, is based on a group sparsity model as defined below and cannot be decoupled into independent terms. We also assume that the resting fMRI data, D_i , are conditionally independent across subjects given the individual parcellations, L_i , and connectivities, C_i .

The parcellation and connectivity can then be computed by maximizing the log posterior:

$$\mathcal{L}(L, C|D) = \sum_{i=1}^N \ln P(D_i|L_i, C_i) + \sum_{i=1}^N \ln P(L_i) + \ln P(C) + \text{const}. \quad (2)$$

We now specify each term in $\mathcal{L}(L, C|D)$. For each subject, we assume a zero-mean multi-variate normal likelihood

$$P(D_i|L_i, C_i) = \frac{1}{\sqrt{(2\pi)^M |C_i^{-1}|}} \exp \left\{ -\frac{1}{2} \text{Trace}(Z_i^T C_i Z_i) \right\}, \quad (3)$$

where Z_i is an $M \times T$ matrix representing the mean signal averaged over all vertices in each of the M parcels in L_i , as described. The log likelihood function can be expressed as

$$\ln[P(D_i|L_i, C_i)] = \frac{1}{2} \ln(\det C_i) - \frac{1}{2} \text{Trace}(Z_i^T C_i Z_i) + \text{const}. \quad (4)$$

The group functional connectivity prior $P(C)$ in (2) is designed to favor a sparse concentration matrix whose non-zero entries determine the connectivity between parcels. We assume that subjects will exhibit similar patterns of connectivity but the connection strengths between

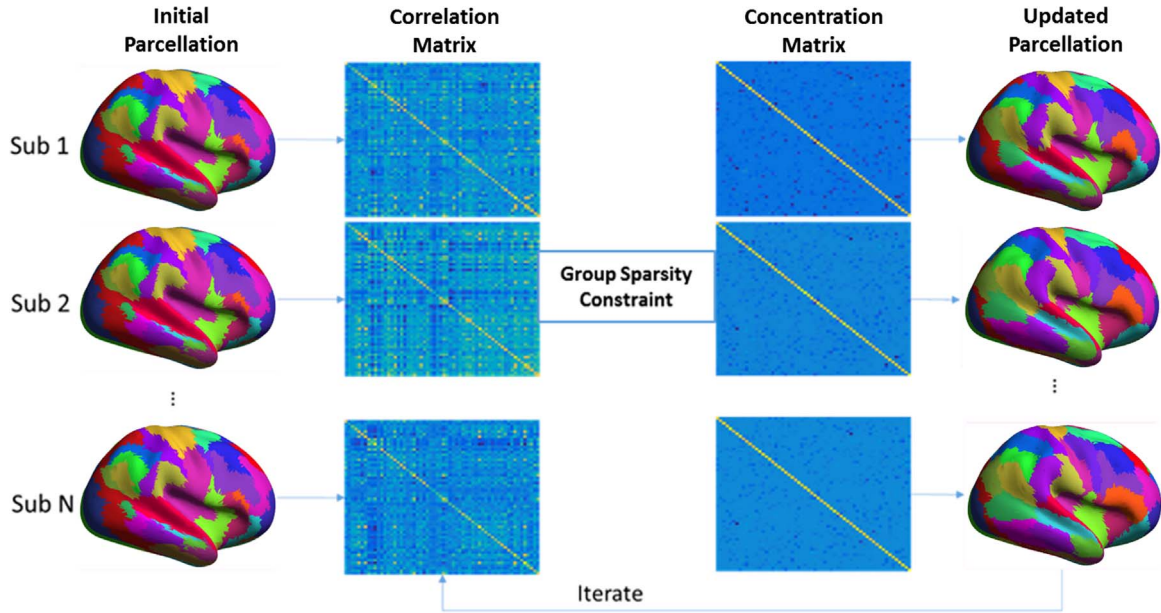


Fig. 1. Illustration of the iterative optimization approach. Leftmost column shows the initial parcellation which is identical for all subjects. We use this to compute parcel-wise sample correlation matrices as shown the second column. We then estimate parcel-wise concentration matrices under the group sparsity constraint in the third column. In the right most column we then refine each individual's parcellation. Iterative sequential updates of the concentration matrices and parcellations are then repeated to convergence.

parcels can vary. To achieve this we use the $l_{2,1}$ norm to represent group sparsity (Cotter et al., 2005; Ng et al., 2012):

$$\ln [P(C)] = -\lambda \sum_{p=1}^M \sum_{q=1}^M \sqrt{\sum_{i=1}^N C_i^2[p, q]} + \text{const.} \quad (5)$$

The final term in the model is the spatial prior $P(L_i)$ on the parcellation labels defined on the cerebral cortex. We use a Potts Markov Random Field (MRF) model (Boykov et al., 2001) for this purpose. We denote the labels of each vertex $L_i(a) = p \in \{1 \dots M\}$, for $a = 1 \dots V$ and subject $i \in \{1 \dots N\}$. We add a penalty each time a vertex label $L_i(a)$ differs from that of one of its neighbors N_a using the Kronecker delta $\delta(x)$. The joint density (Besag, 1986) has the form:

$$P(L_i) = \frac{1}{Z_\beta} \exp \left\{ -\beta \sum_{a=1}^V \sum_{b \in N_a} (1 - \delta[L_i(a) - L_i(b)]) \right\}, \quad (6)$$

where β is the Gibbs hyperparameter and Z_β the partition function. The neighborhood N_a for each vertex consists of the set of vertices on the surface tessellation that share an edge with vertex a .

Combining (1)–(6) we can express the solution to the maximum a posterior estimation problem as the minimizer of the negative log posterior density as follows:

$$\{L, C\}_{\text{GPIP}} = \underset{L, C}{\text{argmin}} \left\{ \sum_{i=1}^N \left(-\frac{1}{2} \ln(\det C_i) + \frac{1}{2} \text{Trace}(Z_i^T C_i Z_i) \right) + \beta \sum_{a=1}^V \sum_{b \in N_a} (1 - \delta[L_i(a) - L_i(b)]) \right\} + \lambda \sum_{p=1}^M \sum_{q=1}^M \sqrt{\sum_{i=1}^N C_i^2[p, q]}. \quad (7)$$

Optimization method

We use a pre-defined cortical parcellation to initialize all subjects. Here we use those from Yeo, et al. (2011) and Glasser et al. (2016), but it is also possible to use other atlas or data driven parcellations, e.g. Blumensath et al. (2013) or Craddock, et al. (2012), for initialization. The resolution of the initial atlas will determine the density of the final

parcellation so its selection should be based on the intended application. We assume that an invertible mapping is defined from each subject to the common cortical atlas. As parcellations become increasingly fine, we would expect this assumption to break down and in those cases, the method described here would not be able to correctly identify all functional subdivisions. However, for studies involving larger functional parcels or networks we would expect a reasonable degree of consistency across groups (Kim et al., 2010; Eickhoff et al., 2011). The algorithm below can be applied to data mapped to the common space, with the final results mapped back to define the parcellation with respect to each individual's cortical anatomy.

We use an alternating method, with guaranteed monotonic decrease of the cost function, to sequentially update the concentration matrices and parcellations:

$$C^{k+1} = \underset{C}{\text{argmin}} \left\{ \sum_{i=1}^N \left(-\frac{1}{2} \ln(\det C_i) + \frac{1}{2} \text{Trace}(Z_i^{kT} C_i Z_i^k) \right) + \lambda \sum_{p=1}^M \sum_{q=1}^M \sqrt{\sum_{i=1}^N C_i^2[p, q]} \right\}, \quad (8)$$

$$L^{k+1} = \underset{\{L_i: i=1 \dots N\}}{\text{argmin}} \left\{ \sum_{i=1}^N \left(\frac{1}{2} \text{Trace}(Z_i^T C_i^{k+1} Z_i) \right) + \beta \sum_{a=1}^V \sum_{b \in N_a} (1 - \delta[L_i(a) - L_i(b)]) \right\}. \quad (9)$$

In the first step the concentration matrices for each subject are jointly updated given the current parcellation. In the second step, we update the parcellations using the current concentration matrix estimates. Note that in this second step, the parcellations decouple into separate optimization problems, one per subject. Fig. 1 illustrates this iterative approach.

To update the covariance estimates, at the k th iteration we first compute the average time series in each cortical patch for each subject as the set of matrices $Z_i^k \in \mathcal{R}^{M \times T}$. We then update the concentration matrix estimates using the group-sparsity regularized cost function in (8). This is a convex problem and can be optimized using the Alternating Direction Method of Multipliers (ADMM) (Boyd et al.,

2011). We use the ADMM approach described by (Ng et al., 2012).

To update the parcellations we use a graph theoretic approach as described in detail in the Appendix A. By inspection, the optimization in (9) decouples into separate problems, one per subject. We use a graph-cut formulation of the problem of GPIP parcellation over a Markov random field (Boykov et al., 2001). We update parcel labels using an iterative process in which, rather than consider the entire cortex, at each iteration we instead form a graph that contains nodes that correspond only to those vertices in the surface tessellation that lie at the boundaries of the current parcellation. The update procedure then partitions this graph into M subgraphs each corresponding to one parcel label. Each sub-iteration allows the boundary vertices in each parcel to move by at most one vertex, with a guaranteed reduction of the cost function in (9). Iterative repetition of the graph cutting procedure, each time using a new graph based on the latest parcel boundaries, allows the boundaries to move to the positions that jointly minimize the cost function in (9). The purpose of using only boundary vertices is to substantially reduce the time and memory complexity of the problem to an acceptable scale at each iteration. The spatial Markov property we assume and the manner in which the graphs are constructed at each iteration ensures that at convergence the method will select an optimal partition with respect to the entire cortical surface.

The two-steps iteration is repeated until effective convergence at which no more than one vertex is changed in each parcel during the most recent update or 40 iterations have been performed.

In vivo data

All of the *in vivo* results below used the minimally preprocessed (ICA-FIX denoised) resting fMRI data from 40 randomly selected subjects (all right handed, age 26–30, 16 male and 24 female), which are publicly available from the Human Connectome Project (HCP) (Van Essen et al., 2013; Glasser et al., 2013; Smith et al., 2013; Barch et al., 2013; Woolrich et al., 2001). These data were acquired for four independent resting fMRI sessions (LR_1, LR_2, RL_1, RL_2) of 15 min each (TR=720 ms, TE=33.1 ms, 2 mm×2 mm×2 mm voxels) with the subjects asked to relax and fixate on a projected bright cross-hair on a dark background. The phase encoding direction was reversed between (LR_1, LR_2) and (RL_1, RL_2). Here we use the LR_1 sessions to obtain the individual parcellations and the LR_2 sessions to evaluate performance.

HCP's minimal preprocessing corrects the fMRI data for acquisition artifacts, resamples the data on the cortical surface and performs a non-aggressive spatio-temporal cleanup (Glasser et al., 2013; Smith et al., 2013). The data are then resampled for each subject onto a common standard cortical surface mesh representation (32 K Conte-69) in the native subject space (Smith et al., 2013; Van Essen et al., 2012). Additional spatial-temporal processing is then used to remove scanner, motion and other artifacts (Smith et al., 2013). Our only additional preprocessing prior to application of the GPIP parcellation algorithm was to normalize the resting fMRI time series at each vertex to zero mean and unit variance.

The functional parcellation results presented below are compared to task-localizer data also made available by the HCP. Task-based fMRI data were obtained for the same 40 subjects and included seven major task domains: emotion processing, reward & decision making, language processing, motor strip mapping, higher order relation processing, social cognition and working memory. We used HCP's pre-processed and analyzed task-fMRI data resampled onto the cortical surface, with 4mm Gaussian smoothing, which yielded a total of 15 statistical task-pair activation maps as described in detail by Barch et al. (2013) and Woolrich et al. (2001). To identify regions involved in emotion processing, a face versus shape selection scenario was used as the “faces_shapes” contrast. For reward & decision making, gambling punishment and reward tasks were employed and contrasted (“pun-

ish_reward”). To identify regions of language processing, interleaved math and story tasks are compared in the “math_story” contrast. The relational processing study contrasts two tasks: the first requiring the subject to identify geometrical relationships among objects, and the second involving matching of objects. The two responses are compared in the “match_rel” contrast. In the study of social cognition, the theory of mind (TOM) task involved blocks of moving geometric shapes wherein some moved with intention (TOM) and were compared to random movement (random). These are compared in the “random-TOM” contrast. To sample regions involved in working memory, a 2-back versus 0-back contrast was used (“0bk_2bk”). A number of other contrasts were performed to identify functionally specialized regions, where a certain task is compared to the average of other tasks in that category. In the motor category, regions specific to a particular body part were identified using a contrast of the response of that part to the average response of all other body parts. The contrasts used below involve movement of the tongue (“t_avg”), left hand (“lh_avg”), right hand (“rh_avg”), left foot (“lf_avg”) and right foot (“rf_avg”). Similarly, perception of different categories of images was contrasted with the average response over the other categories as follows: faces (“face_avg”), place processing (“place_avg”), tools (“tool_avg”) and body perception (“body_avg”).

Atlases for initialization

For the *in vivo* studies described below, we choose two atlases, Yeo et al., (2011) and Glasser et al., (2016) for initialization. Yeo's atlas uses resting-fMRI data from a group of 1000 subjects. These data were resampled onto their respective cortical surfaces and co-registered to a common atlas space. Clustering was used to identify a common set of 17 networks, each consisting of multiple contiguous regions on the cortical atlas.¹ We attached a distinct label to each contiguous region, merging those with less than 30 vertices with their neighbors, to produce a final set of 58 cortical parcels. Glasser's atlas is a more detailed group-based parcellation constructed from both functional and anatomical MRI data from 210 subjects in the HCP database. A machine learning approach was taken to identify a set of 180 cortical parcels per hemisphere.² We reduced this number to a total of 338 regions by again merging those with less than 30 vertices with their neighbors.

Evaluation and metrics

The GPIP parcellation algorithm described above was applied to the 40 HCP subjects' resting fMRI (LR_1 sessions) initialized with both the Yeo and Glasser atlases. The resulting parcellations, customized to each subject, were then compared to the initial common group atlas parcellation with which the algorithm was initialized. To quantify the degree of refinement required in each subject, we computed the fraction of vertices whose label is changed from initialization to convergence.

For each subject, we used the resting fMRI LR_2 sessions to quantify homogeneity and connectivity variability based on the parcellations obtained from LR_1. As in Craddock et al. (2012), homogeneity of a parcel is defined as the average temporal correlation coefficient between all pairs of vertices in that parcel. A larger homogeneity value indicates greater similarity in time series within that region and is an indication of increased functional homogeneity. As an overall measure, we compute the average homogeneity values over all parcels for each subject and then average these over all subjects. This score is computed with respect to the initial atlas and final parcellation, and as a function of iteration number as an indication of convergence behavior of the algorithm.

¹ Available from http://surfer.nmr.mgh.harvard.edu/fswiki/CorticalParcellation_Yeo2011.

² Available from <https://balsa.wustl.edu/>.

We also studied the variability in connectivity across the group, again using the LR_2 data. We first computed the Pearson correlation between each pair of regions using the time series averaged over each region. For each pair of regions we computed the variance of this correlation across the 40 subjects. We then computed the mean of these variances across all region pairs together with its standard error. This variability measure was computed as a function of iteration to indicate convergence behavior as well as allowing a comparison of the initial group atlas and the final individualized parcellations. Lower values of variability will reflect a higher degree of consistency across subjects in the strength of the correlations between parcels. This in turn would indicate improved similarity across the group in network connectivity between the regions making up the parcellated cortices. Conversely, larger variability would suggest inconsistent parcellation across the group. This could be caused by poor accuracy in identification of region boundaries. Alternatively, high variability could be a consequence of mislabeling (*i.e.* one or more regions in the atlas end up assigned to different functional areas in different subjects).

Finally, we compared resting fMRI parcellations against task-fMRI contrast activation maps for each subject. To compute task variability for a certain task pair, for each subject we computed the variance in each parcel of that activation map's z-scores as provided by the HCP (Barch et al., 2013; Woolrich et al., 2001). The task variability was then averaged over all regions and all subjects to produce a single task variability metric. We computed task variability for the initial atlas and final parcellation separately for all 15 task contrasts described above. We then applied a Wilcoxon rank-sum test to determine if there was a significant reduction in task variability. Lower task variability would suggest higher functional homogeneity in parcels in terms of a specific task-fMRI contrast and implies better alignment between parcel boundaries and task activation regions.

Perturbation analysis

We explored the robustness of our algorithm in the case where the perturbed initialization contains regions that do not overlap spatially with the locations of these regions in the original atlas. To do this we choose a parcel in the posterior inferior supramarginal gyrus in Yeo's atlas. We then perturbed the atlas so that this region was shifted in the posterior direction so that it has zero overlap with its position in the original Yeo atlas. The surrounding parcels were correspondingly modified so that the spatial topography of the parcels remained unchanged. We then applied the parcellation algorithm to the same 40 subjects resting fMRI LR_1 data sets initialized with this perturbed Yeo atlas. To determine the impact on the resulting parcellation, for each subject we computed the Dice coefficient for the right posterior inferior supramarginal gyrus parcel between results obtained using the two versions of Yeo's atlas. We then computed the mean and standard error of the Dice coefficient across all subjects as a function of iteration.

Simulation studies

We also explored the performance of our parcellation approach using simulated data. Synthetic data was generated to simulate resting fMRI functional connectivity across a group of subjects with variation in both their cortical parcellations and concentration matrices. We start with a randomly generated sparse symmetric positive definite matrix \tilde{C} representing the mean partial correlation matrix for the population. Individual concentration matrices were obtained by varying each off-diagonal concentration matrix strength (non-zero elements of \tilde{C}) by sampling from a normal distribution with mean 1 and standard deviation 0.5. Results were tested for positive definiteness and the sampling repeated in cases where this condition was not met. The resulting matrices $C = \{C_1, C_2, \dots, C_N\}$ were used to define the ground-truth connectivity for each of $N=24$ subjects. We used the correspond-

ing correlation matrices to generate synthetic Gaussian data representing resting BOLD interactions between cortical parcels. These were then mapped to a 100×100 lattice gridded surface and independent Gaussian spatially and temporally white noise added separately to each vertex. Gaussian smoothing was performed both temporally and spatially. The cortical parcellations for each subject, except subject 1, were based on random perturbations of the boundaries of an initial common parcellation that has 25 square parcels each of size 20×20 as shown in Fig. 8. Individual boundaries were generated by random application of a sequential single-vertex morphological dilation to 12 parcels randomly selected out of 25. Repeating this process 6 times produces a shift of at least 6 vertices between the initial atlas and ground truth parcellations for each subject. Subject 1 was specially designed so that the location of one of the parcel (the red parcel indicated by a black arrow in Fig. 8) does not overlap with its location in the initial atlas. This allows us to examine the ability of GPIIP to correctly identify individual parcels that have no overlap with their initial location in the atlas. The average error rate (percentage of mislabeled vertices) between these ground truth parcels and the regular atlas with which the algorithm was initialized was 27.7%.

Results

In vivo data

Fig. 2 shows parcellation results for three subjects, selected from a group of 40. Results for all 40 subjects can be found in Figs. S1 and S2 (Supplementary material). For visual comparison, we show all results mapped to the common cortical surface to which each subject is registered. In this common space, the initial parcellations are identical for all subjects. The algorithm then adapts the parcel boundaries using each individual's resting fMRI to produce the results shown in rows 2–4. The update procedure guarantees that the spatial topology of the original atlas parcels is retained, producing a matched set of regions across subjects that differ only in the locations of their boundaries. The black arrow in Fig. 2 (row 3, column 3) points to a parcel (green patch) which has no overlap with its initial position in the Glasser atlas. This observation indicates the ability of GPIIP to move individual parcels to the point where they no longer overlap with their initial location, a point we return to below in the Discussion. For each subject, the fraction of vertices whose labels changed was computed and resulting statistics over the $N=40$ subjects (mean, standard deviation, minimum and maximum) are shown in Table 1 for both the Yeo and Glasser atlas.

Fig. 3 shows changes in homogeneity computed from the LR_2 resting fMRI data using parcellations obtained with LR_1. Homogeneity is averaged across all parcels as a function of iteration and then averaged across all 40 subjects. In both cases there are significant improvements in homogeneity from initialization to the final iteration. Convergence with the Glasser atlas is slower, presumably because of the considerably larger number of regions.

Fig. 4 shows the variability in connectivity, computed using the LR_2 data, as a function of iteration number for the Yeo and Glasser atlases. In both cases there is a near monotonic reduction in variability with iteration with significant reductions from initialization to the final result.

Examples comparing resting fMRI parcellation with task localizers are shown in Figs. 5 and 6. Quantitative results for all task pairs are included in Tables 2 and 3. Fig. 5 shows result for the faces *vs.* shapes visual task pair (Barch et al., 2013) for three individuals. Results for all 40 subjects are included in Figs. S3 and S4 in the Supplemental material. For both initializations, Fig. 5 shows apparent improvement from the initial atlas to the individualized results in the sense that regions of higher z-score appear better constrained between parcel boundaries after application of GPIIP, particularly in the fusiform face area, para-hippocampal gyrus and visual cortex. Fig. 6 is an example using the TOM *vs.* random (theory of mind *vs.* random motion) social




Fig. 2. Examples of individualized resting fMRI parcellation, medial and lateral view of left hemisphere. Top row: initialization with Yeo (left) and Glasser (right) atlases. Rows 2–4: Final parcellations for three of 40 subjects. Each color represents a separate parcel, with colors matched across subjects. The black arrow in row 3, column 3 points to a parcel (green patch) which has no overlap with its initial position in the Glasser atlas.

Table 1
Statistics over n=40 subjects of the fraction of vertices that were relabeled per subject.

	Mean	Standard deviation	Minimum	Maximum
Yeo's Atlas	0.3555	0.0156	0.3286	0.3874
Glasser's Atlas	0.4574	0.0101	0.4364	0.4839

cognition task contrast initialized with the Glasser atlas for two subjects and projected to the medial surface. Results for all 40 subjects are included in Fig. S5 in the Supplemental material. With the high density of the parcellation, the results are perhaps difficult to interpret visually, however bounded parcels in medial prefrontal and parietal regions appear more homogeneous.

To quantify the consistency between resting fMRI parcellations and task-fMRI for all task pairs we computed task variability as the variance of each task contrast's z-score within parcels as described above. In Table 2 we tabulate average task variability with respect to the initial

Yeo atlas and after application of GPIP. We also performed a Wilcoxon rank-sum test for differences in task variability. Equivalent results for initialization with Glasser's atlas are shown in Table 3. Note that in all cases the difference between initialization and GPIP results are positive, indicating reduced variability with respect to the parcellations for all task pairs for both the Yeo and Glasser atlases. The p-values listed in the final column indicate significant reductions ($p < .05$, Bonferroni corrected) for all but the face_avg working memory task with the Yeo atlas.

Perturbation analysis

Fig. 7 shows behavior in the cases where the Yeo atlas initialization is modified so that a parcel in the posterior inferior supramarginal gyrus is moved posteriorly so that it no longer overlaps with its original position (red patch, top row of Fig. 7). Using this as an initialization we again applied the GPIP algorithm to the LR_1 data and compared results with those obtained with the original Yeo atlas. We show the final parcellations




Fig. 3. Homogeneity of resting fMRI parcellations as a function of iteration. Shown is mean homogeneity across all regions and subjects (plus/minus the standard error). (a) Initialization using the Yeo atlas; (b) initialization using the Glasser atlas.




Fig. 4. Variability of connectivity (mean plus/minus the standard error) as a function of iteration number computed as the variance of correlation across the population between each pair of regions, averaged over all pairs. (a) Initialization using the Yeo atlas; (b) initialization using the Glasser atlas.

for two subjects using the Yeo and perturbed Yeo atlases in Fig. 7. We also plot the Dice coefficient between the posterior inferior supramarginal gyrus parcels for the two atlases, averaged over all 40 subjects, as a function of iteration. The final average Dice coefficient is 0.65. While this does not indicate an exact overlap, this result does indicate that the algorithm is able to produce large changes in individual parcels when required, and, importantly, that the perturbed parcel is not misaligned to a different adjacent functional region. This robustness stems from the fact that both initializations share the same topography or spatial connectivity between parcels. Movement of the boundaries of the better aligned parcels will affect those of their neighbors, tending to correct initial functional misregistration of individual parcels.

Simulated data

We investigated the convergence properties of the GPIP method using the simulated dataset described above. For each ‘subject’ we computed the error rate in labeling (fraction of vertices that are mislabeled) as a function of iteration by comparing with known ground truth. Fig. 8(a) shows examples of the initial and final parcellations for four subjects selected from the group of 24. Final parcellations for all synthetic subjects can be found in Fig. S6 (Supplemental material). The group averaged error rate at each iteration is plotted in Fig. 8(b). Since the error rate is zero at the 20th iteration, the images at convergence also show the degree of variability in spatial structure of the ground truth images. The special case of subject 1 also converged to the ground truth, indicating that GPIP can correctly identify parcels even in cases where the initial and true (final) parcels do not overlap. These results also illustrate convergence for data that are consistent with the model described above. Further, the results indicate the ability of the algorithm to adapt individual parcellations to match ground truth that is substantially (27.7% on average) different from the initialization.

Discussion

The GPIP algorithm described above takes a common cortical atlas that delineates regions and adapts the boundaries of each region within subject to better fit the individual subject’s regional functional profile from resting fMRI. Through this procedure we are able to produce individualized parcellations but with the property that all subjects have the same labeled parcel topology, facilitating group analysis. Our approach can be viewed as a means of refining a common group atlas. As the results above indicate, this refining process produces consistency, relative to the initial atlas, of regions of functional specialization identified using both independent resting and task fMRI.

In the *in vivo* results presented, we initialize with two published

cortical atlases although the approach could be applied to any parcellation defined with respect to coregistered cortical surfaces. Of the two atlases, the first is based solely on resting functional MRI data (Yeo et al., 2011) while the second uses a combination of anatomical and functional data to define a far denser parcellation (Glasser et al., 2016). Results for both atlases indicate improved homogeneity, reduced inter-subject variability of connectivity, and reduced variability of task-related activation within parcels. However, further work will be needed to validate whether resting fMRI data alone can be used to define dense parcellations which originate from integrated multi-modal imaging (e.g. Glasser et al., 2016) especially certain cortical regions that are anatomical driven.

The maps showing individual parcellations mapped on to a common cortical representation in Fig. 2 (and Figs. S1 and S2 in Supplemental material) show substantial deviation across subjects in region boundaries for both atlases. The statistics in Table 1 also indicate the fraction of vertex relabeling required to individualize the atlas boundaries is similar for all subjects. Initializing with the Glasser atlas requires a larger fractional change (mean 0.4574) than does the Yeo atlas (mean 0.3555). Our results also indicate that the GPIP boundaries have improved concordance, relative to the original atlas, with both resting and task related fMRI data. The inter-subject variability in these boundaries, as shown in Fig. 2, S1 and S2, indicates the limited accuracy with which functional regions can be identified by anatomically-driven mapping to a common atlas. This in turn will limit the power of analyses based on a common atlas, which have been shown to have moderate inter-subject confidence in peak regions (e.g. Yeo et al., 2011, Table 5), and low confidence in boundary regions (e.g. Wang et al., 2015; Wig et al., 2014). The approach described here has the potential to improve statistical power of region-based statistical analysis through improved delineation of functional parcels while retaining a common set of regions across subjects.

Functional areas of the brain demonstrate substantial spatial variability across subjects, and the individual localization of these functional regions optimizes estimation of inter-regional connectivity (e.g. Stevens, et al., 2015). However, task-localizers for functional area identification across cortex exceed the constraints of a single imaging study. Comparisons of the resting fMRI based parcellation with a series of functional localizers shows consistent improvements in concordance with task contrasts. Maps of parcel boundaries overlaid on the z-scores in Figs. 5 and 6 illustrate improved consistency of these boundaries with known areas of functional specialization. Note that for the multiple subjects we presented, each individual subject has its own distinct activation region that cannot be well-delineated by a common atlas. In Fig. 5 the GPIP refinement of the Yeo atlas shows clear changes in boundaries so that we are able to better contain the fusiform

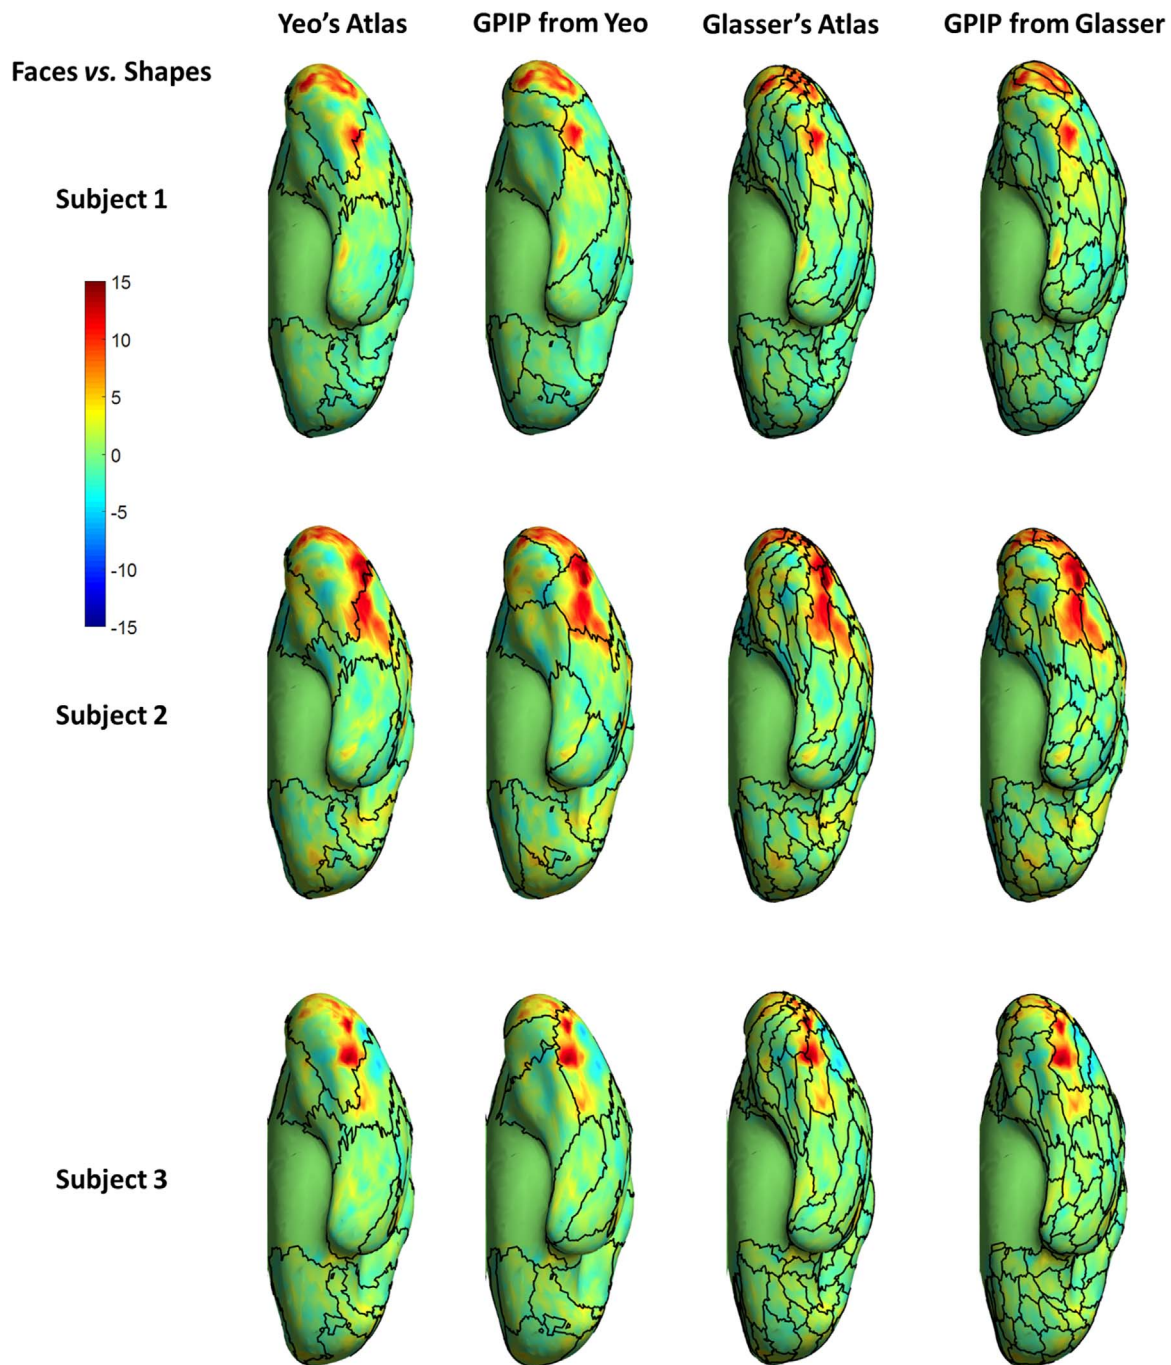


Fig. 5. Comparison of resting fMRI parcellation and activation maps for the faces *vs.* shapes visual task pair, ventral view of right hemisphere. First and third columns: initial parcellations using the Yeo and Glasser atlas, respectively, with boundaries indicated by black curves. These are overlaid with color maps of z-scores computed for the faces *vs.* shapes contrast for three subjects. Second and Fourth columns: final parcellation boundaries initialized from the Yeo and Glasser atlas respectively.

face area in a single parcel, rather than splitting into two or more regions. Similarly, although Glasser's atlas has a different functional resolution from Yeo's, in Fig. 5 we see that in the refined parcellation, the medial and posterior boundaries of the fusiform area are frequently consistent for both Yeo and Glasser. This provides some evidence that the regional boundary is indeed determined by the individual's resting brain activity and can be found consistently even for different initial atlases. Similarly, as shown in Fig. 6, parcels obtained with GPIP using the Glasser atlas may better functionally align parcels to be consistent with individual differences in activation patterns across subjects in the TOM *vs.* random task contrast (Barch et al., 2013). The quantitative results in Tables 2 and 3 confirm these observations across multiple subjects and every task-pair. GPIP consistently shows significantly

lower task variability than the corresponding initialization.

While we note that boundary refinement did produce improvement (reduction) in within-parcel task related variability with the Glasser atlas, these results should be interpreted with caution. The original Glasser atlas is based on structural as well as functional features in the data, and certain areas (*e.g.* motor cortex and visual cortex) are delineated primarily by anatomical features and may not closely reflect regions of functional specialization. The clearest example of this is in primary motor cortex where hand and face areas are combined in a single parcel in the Glasser atlas while they are readily differentiated in resting fMRI data. As a result, refinement of Glasser with GPIP can result in recruiting neighboring parcels to produce a subdivision between hand and face areas, as can be seen in some of the subjects




Fig. 6. Comparison of resting fMRI parcellation and activation maps for the TOM *vs.* random social cognition task pair, medial view of left hemisphere. Left column: initialization using the Glasser atlas (black curves) overlaid with color maps of z-scores computed for the TOM *vs.* random contrast for three subjects. Right column: final individualized parcellation boundaries, again overlaid with task activation maps.

Table 2

Task variability within each parcel averaged across all parcels and all subjects for each of the 15 task pairs. Shown are task variability computed with respect to the Yeo atlas initialization (column 3) and final result using the GPIP algorithm (column 4). Differences between the two are shown in column 5 and p-values for the Wilcoxon rank-sum test in column 6.

	Contrast	Yeo's	GPIP from Yeo's	Difference (Yeo – GPIP)	p-value (ranksum)
Emotion	<i>faces_shapes</i>	3.269	2.920	0.350	1.01E-10
Gambling	<i>punish_reward</i>	2.255	1.795	0.460	1.66E-89
Language	<i>math_story</i>	12.180	7.290	4.889	1.33E-200
Motor	<i>lf_avg</i>	2.826	2.413	0.414	2.27E-26
	<i>lh_avg</i>	3.240	2.861	0.379	6.72E-10
	<i>rf_avg</i>	2.506	2.115	0.396	1.27E-25
	<i>rh_avg</i>	3.001	2.599	0.402	4.03E-15
	<i>t_avg</i>	7.296	6.393	0.903	6.39E-10
Relational	<i>match_rel</i>	1.952	1.504	0.443	5.03E-83
Social	<i>random_tom</i>	3.159	2.279	0.886	2.63E-92
WM	<i>Obk_2bk</i>	7.203	4.860	2.342	9.57E-90
	<i>face_avg</i>	4.948	4.680	0.268	3.24E-02
	<i>place_avg</i>	4.487	4.056	0.425	1.15E-09
	<i>tool_avg</i>	3.817	3.246	0.571	1.05E-40
	<i>body_avg</i>	4.178	3.654	0.525	1.50E-14

in Fig. S2. More generally, some of the parcels may be adapted to locally fit a functional area that differs from the original anatomical-functional region identified in the Glasser atlas. As a result further cross-validation is required to evaluate which regions are most reliably identified when using this atlas.

We also note that the Glasser atlas exhibits lower task variability

that Yeo's in the initial atlas and after GPIP refinement (Table 2 *vs.* Table 1). This is expected because the smaller parcel sizes in Glasser will have lower variance than those in Yeo for spatially inhomogeneous z-score maps. Common to both atlases is the significant reduction in variability for all task pairs, with the exception of the face *vs.* average working memory task for Yeo's atlas.

The performance of our GPIP algorithm is highly dependent on the initial atlas since this defines the number and arrangement of parcels. Given the range of individual variation, it is important the approach is able to identify equivalent functional regions across subjects even in cases where location of some functional regions are highly perturbed in the initial atlas. To investigate this, we perturbed the Yeo atlas as shown in Fig. 7 and described above. In this case, despite the fact that one parcel in the perturbed atlas was completely shifted so that it had zero overlap with its original location, the GPIP algorithm was able to produce individualized mapping with an average dice coefficient of 0.65 in the shifted parcel relative to results where the original atlas was used. This indicates a robustness of the approach that likely arises from the constraint we impose of a common atlas topology across subjects. In this way, mislocated regions will be pulled back towards their true location by the refinement of the boundaries of their neighboring parcels. The results shown here is an extreme case since the parcel is mislocalized in the atlas so will be incorrect with respect to most of the individuals. In practice, we may instead expect to see a few subjects whose functional areas are substantially shifted, or even have zero overlap, relative to the initial atlas. This will become increasingly common with finer and finer parcellations. To address the question of whether GPIP will be able to correctly identify parcels in this case, we searched the *in vivo* results across the 40 subjects shown in Fig. S2 for

Table 3

Task variability within each parcel averaged across all parcels and all subjects for each of the 15 task pairs. Shown are task variability computed with respect to the Glasser atlas initialization (column 3) and final result using the GPIP algorithm (column 4). Differences between the two are shown in column 5 and p-values for the Wilcoxon rank-sum test in column 6.

	Contrast	Glasser's	GPIP from Glasser's	Difference (Glasser - GPIP)	p-value (ranksum)
Emotion	<i>faces_shapes</i>	2.970	2.220	0.749	3.41E-71
Gambling	<i>punish_reward</i>	1.804	1.632	0.172	8.07E-27
Language	<i>math_story</i>	6.394	5.114	1.280	6.10E-91
Motor	<i>lf_avg</i>	2.340	2.023	0.317	1.67E-44
	<i>lh_avg</i>	2.629	2.387	0.242	9.08E-22
	<i>rf_avg</i>	2.056	1.797	0.258	7.17E-37
	<i>rh_avg</i>	2.444	2.190	0.254	3.75E-26
	<i>t_avg</i>	6.072	5.444	0.628	3.30E-16
Relational	<i>match_rel</i>	1.469	1.275	0.194	1.05E-57
Social	<i>random_tom</i>	2.175	1.841	0.333	1.32E-76
WM	<i>Obk_2bk</i>	4.510	3.799	0.711	4.74E-67
	<i>face_avg</i>	4.804	3.902	0.903	3.83E-42
	<i>place_avg</i>	4.024	3.268	0.756	1.03E-67
	<i>tool_avg</i>	3.300	2.874	0.426	1.33E-39
	<i>body_avg</i>	3.637	3.079	0.557	3.11E-54

cases where GPIP converges with individual parcels that have zero overlap with their initial location. We found three subjects for which this occurred, and show one example in Fig. 2. We also explored this question in the simulation study, and again found that GPIP is able to identify parcels that do not overlap with their initial location (Fig. 8). As noted above, this probably follows from the assumption of a common topology across subjects, so that neighboring parcels tend to pull or push the non-overlapping parcel towards its correct location.

A potential disadvantage of the GPIP approach is that the model cannot accommodate differences in the number of functional regions across subjects or functional connectivity that is topological different from the group atlas, for which there is evidence in the literature (Amunts et al., 1999; Harrison et al., 2015), limiting some potential applications. Since GPIP adapts an initial atlas to individual resting fMRI data, it is unable to identify functional variability that is systemically absent from the group atlas (Laumann et al., 2015; Gordon et al., 2015; Mueller et al., 2013).

Our individualized GPIP functional parcellation method uses an

optimization framework to compute both the parcellation and the connectivity concentration matrix for each subject. This requires solving a large non-convex optimization problem with a mixture of continuous and discrete variables. As a result a globally optimal result is not guaranteed. However, the results of the simulation (Fig. 8) show convergence of the algorithm in the case where the data conforms to the assumed model. The two-steps algorithm, Eq. (8) and Eq. (9), exhibits monotonic reduction of the cost function in Eq. (7), and has the useful property that the function is convex with respect to the concentration matrices that are updated in Eq. (8). For the 32 K Conte-69 (Smith et al., 2013) surface mesh used here we found that after approximately 20 iterations, less than 1% of vertices are relabeled in any further iteration, irrespective of the number of parcels, when a common initialization is used. Although this may vary with quality of data and mesh size, we found for the HCP data that 40 iterations were sufficient to achieve effective convergence.

Results are dependent on the choice of regularization or hyper parameters, λ and β in Eq. (9). For all results in this paper we used

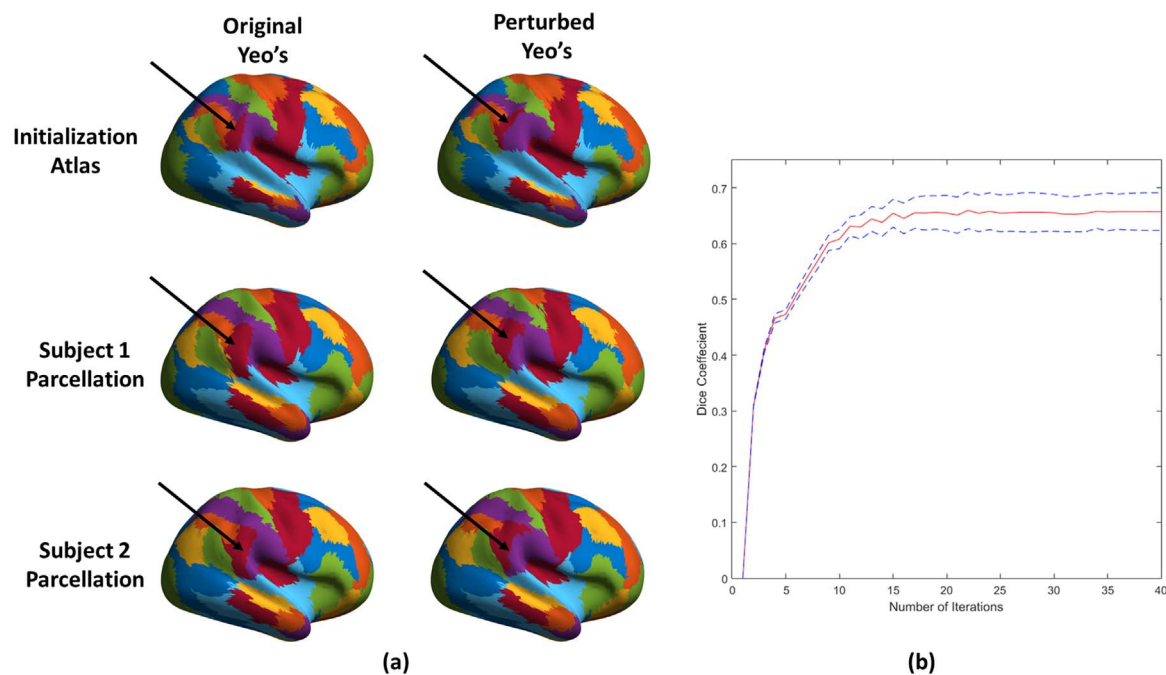


Fig. 7. Illustration of parcellation stability to a perturbed initialization: (a) (top row) initialization with Yeo atlas (left) and Yeo atlas perturbed with respect to red parcel in right posterior inferior supramarginal gyrus (right), as indicated by the black arrow; (rows 2 and 3) individualized parcellation results using the two initializations for two of forty subjects; (b) plot of Dice coefficient, averaged over 40 subjects for the posterior inferior supramarginal gyrus parcel, as a function of iteration.

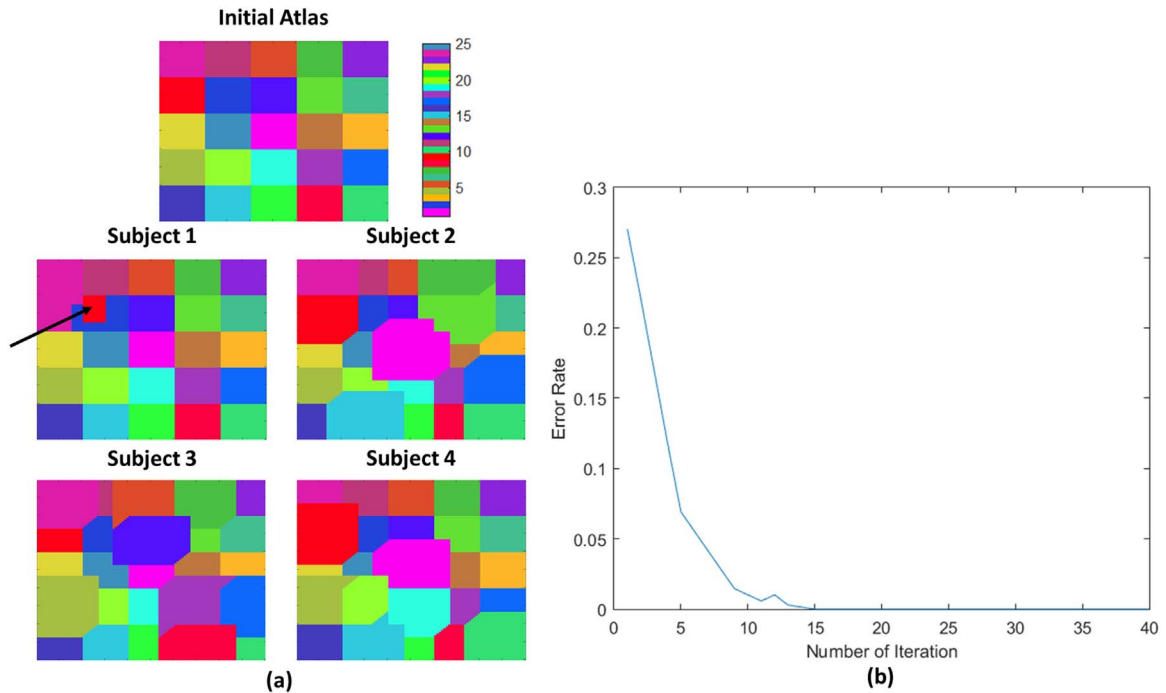


Fig. 8. Simulation results: (a) Initial atlas and examples of the final parcellations for four ‘subjects’. Each color represents one of the synthetic parcels which are matched across subjects. The red parcel for subject 1, indicated by a black arrow, has zero overlap with the common initialization. (b) The error rate (fraction of mislabeled vertices) in the parcellations, averaged over all 24 subjects, as a function of iteration number.

$\lambda = 10$ with $\beta = 0.1$ for GPIP from Ye’s and $\lambda = 10$ with $\beta = 1$ for GPIP from Glasser’s, chosen by visual evaluation of results for several *in vivo* datasets. Further optimization of parameter selection by, for example, modifying parameters to maximize homogeneity by resting-fMRI data cross session validation could further improve performance.

In this paper we present all results for data sampled on the cortical surface. The approach generalizes to volumetric and grayordinate parcellation (Glasser et al., 2013) where minor modifications of the MRF neighborhood system could be used to perform volumetric parcellation of cortical and/or subcortical regions.

Conclusion

We have described GPIP, an algorithm that takes an initial cortical atlas representing regions of functional specialization and refines the boundaries of these regions for each subject using that individual’s resting fMRI. The approach combines a Markov random field prior, which represents individual parcellations of the cerebral cortex, with a multivariate Gaussian model that uses a group-sparsity prior to model similarity across the population of the connectivity between these parcels. GPIP applied to *in vivo* data

shows distinct individual differences in functional boundaries across subjects with respect to a common anatomical alignment. Validation with a second independent fMRI data set indicates improved functional homogeneity within parcels and reduced variation across subjects of the connectivity between pairs of parcels. These boundaries also appear qualitatively and quantitatively consistent with regions of functional specialization identified using functional task localizers. We assume an invertible mapping exists that maps each subject to the common initial atlas. As parcellations become increasingly fine, we would expect this assumption to break down and in those cases, the method described here would not be able to correctly identify all functional subdivisions. However, for studies involving larger functional parcels or networks we would expect a reasonable degree of consistency across groups. It is at this scale that we envision the proposed approach to be a useful addition, allowing group analysis in which common regions of interest are defined across the population while their boundaries are refined to match the individual’s regions of functional specialization. In these cases, GPIP can be used to individualize functional parcellations of cerebral cortex as a precursor to region-based analysis of either resting or task-related fMRI studies.

Appendix A

Optimization by graph partitioning

The parcel labels for each subject are determined as part of the two-steps procedure summarized in Eqs. (8) and (9). Here we describe the computational approach to solving the optimization subproblem in (9).

We first map the optimization problem to an equivalent graph-cut problem. For each subject, we define a graph $G = (Y, E)$, with nodes Y and edges E . In the node set Y , there are two types of nodes: (a) the surface vertex nodes $\{b, b \in B = [b_1 \dots b_V]\}$, where B represents the set of all vertices on the cortex as circles in Fig. A1(a), and (b) the set of parcel nodes $\{p, p \in P = [p_1 \dots p_M]\}$, one node per parcel as diamonds in Fig. A1(b). The parcel nodes are used to assign a label to each vertex so that after partitioning each vertex is connected to a single parcel node. Initially each vertex is connected to every parcel node with an edge strength proportional to the log-likelihood for the data at that vertex given that it belongs to each of the parcels in turn. The edge strength between vertex nodes is determined by the MRF model in Eq. (6) and attaches a weight β between the vertex and each of its neighbors.

We could proceed directly to partition this graph by computing the edge strengths as in Fig. A1(c) and using a graph partitioning method similar

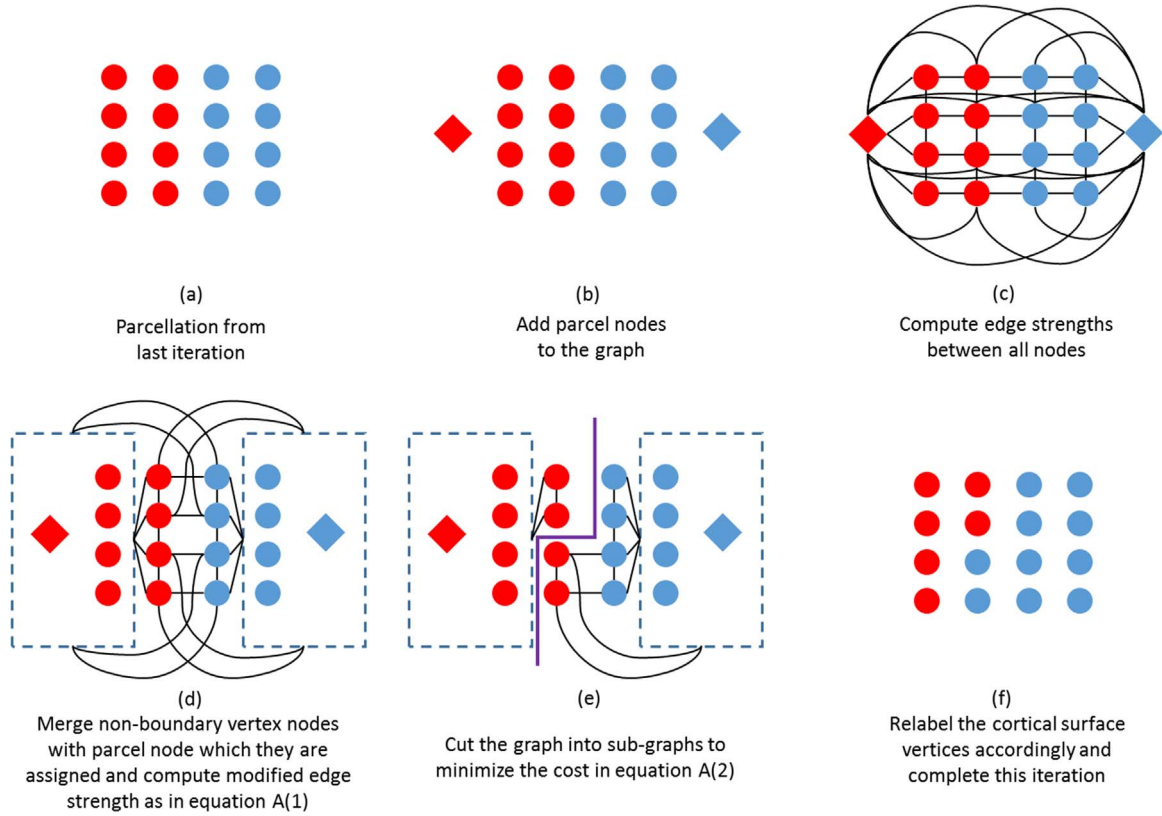


Fig. A1. Illustration of the graph formation and partition approach. The red circles in (a)–(e) represent cortical vertices that have been assigned to the red diamond parcel in the previous iteration, similarly for the blue circles and diamond. At that last step (f) a new cut is made, changing the assignment of parcel labels across the set of cortical vertices. See text for detailed description.

to (Boykov et al., 2001), but the cost would be prohibitive for a problem of this size. We therefore use a modified approach to reduce the size of the graph.

We first prune the vertex nodes to the subset B^k that represents only those vertices on the cortex that lie at the boundary of one of the functional parcels at the previous iteration, k . In other words, all elements of B^k have at least one adjacent vertex with a different label (columns two and three in Fig. A1(a)). The vertices in B that are not in B^k (the first and fourth columns in Fig. A1(a)), will not be changed in the current iteration but since some are mutual neighbors of the boundary vertices, we must still account for their effect on their boundary neighbors, as defined in Eq. (6), in the reduced graph. Since they are already assigned to a parcel label, we can merge these vertices into the node corresponding to the parcel to which they are currently assigned and modify the edge strength from that parcel node to each of the boundary vertex nodes accordingly, Fig. A1(d). Since the modified graph contains only the parcel and boundary vertex nodes it can be as much as two orders of magnitude smaller than the original graph with corresponding computational savings. Provided the edge strengths are modified appropriately, partitioning this modified graph, Fig. A1(e), provides a locally optimal solution, relative to the original graph, by restricting the movement of parcel boundaries to at most one vertex per iteration, Fig. A1(f). A sufficiently large number of iterations can produce an arbitrarily large movement of each boundary.

We define the edge strengths between nodes in the boundary graph as follows (see explanation below):

$$\begin{aligned}
 E_i(a, b) &= \begin{cases} \beta & b \in \mathcal{N}_a, b \in B^k \\ 0 & b \notin \mathcal{N}_a, b \in B^k \end{cases} \\
 E_i(p, q) &= 0, p, q \in P, p \neq q \\
 E_i(a, p) &= \frac{1}{2} \text{Trace} \left(\frac{1}{m_i(p)} \vec{d}_i(a) \vec{l}_p^T C_i^{k+1} Z_i^k \right) + \beta \sum_{b \in \mathcal{N}_a, b \in B^k} (\delta[p - L_i^k(b)]) a \in B^k, p \in P
 \end{aligned} \tag{A1}$$

We use a normalized multi-way minimum cut algorithm to partition the graph into M subgraphs (Yu and Shi, 2003), which as shown by (Boykov et al., 2001; Von Luxburg, 2007), is equivalent to computing a partitioning for the Potts MRF model. In our case, cutting this graph into M subgraphs using minimum cuts, Fig. A1(e), minimizes the cost function (9) over the set of boundary nodes B^k . The restriction to boundary nodes limits the movement of each vertex on the parcel boundaries by at most a single vertex at each iteration, Fig. A1(f). The reduced graph size greatly reduces the per iteration cost. Repeatedly updating and cutting the graph using multiple (k) iterations allows the boundaries to move by up to k vertices and minimizes cost with respect to the entire cortex.

Definition of edge strengths in the graph

The optimization problem in (9) has the form:

$$L^{k+1} = \underset{\{L_i: i=1..N\}}{\operatorname{argmin}} \left\{ \sum_{i=1}^N \left(\frac{1}{2} \operatorname{Trace}(Z_i^T C_i^{k+1} Z_i) \right) + \beta \sum_{a=1}^V \sum_{b \in N_a} (1 - \delta[L_i(a) - L_i(b)]) \right\}. \quad (\text{A2})$$

The edge strengths in (A1) map the cost function on the right to an equivalent graph cut problem. To be consistent with the Potts model, the definition of the edge strength between vertex nodes in (A1) attaches a penalty of β each time a neighboring pair of vertices differ in their parcel labels as a result of a cut. Since we remove vertices from the graph that did not lie on a parcel boundary in the previous iteration, we need to account for the additional cost of moving boundary vertices from the current parcel to a different parcel. This is achieved by the second term in edge strength $E_i(a, p)$, between boundary vertices and parcel nodes, which counts the number of non-boundary vertices that are adjacent to a that currently have the same parcel label.

The first term in $E_i(a, p)$ then computes the contribution of that link to the first term in (A2), the likelihood associated with assigning vertex a to parcel p . The likelihood is computed from the average time series averaged over the vertices in each of the parcels:

$$Z_i = \sum_{a=1}^V \frac{1}{m_i(L_i(a))} \vec{l}_{L_i(a)} \vec{d}_i^T(a), \quad (\text{A3})$$

where $\vec{l}_p = [0..010..0]^T$ is a length M index vector in which the p th element is 1 and 0 otherwise, $m_i(p)$ is the number of vertices with parcel label p and $\vec{d}_i(a)$ is the length T vector representing vertex a 's time series. In order to compute the likelihood as a sum of contributions, one per $E_i(a, p)$ edge, we use the average from the previous iteration, Z_i^k , on the right side of the quadratic form based to give the following approximation to the first term in (A2):

$$\frac{1}{2} \operatorname{Trace}(Z_i^T C_i^{k+1} Z_i) \approx \frac{1}{2} \operatorname{Trace}(Z_i^T C_i^{k+1} Z_i^k) = \frac{1}{2} \operatorname{Trace} \left(\sum_{a=1}^V \frac{1}{m_i(L_i(a))} \vec{l}_{L_i(a)} \vec{d}_i^T(a) C_i^{k+1} Z_i^k \right). \quad (\text{A4})$$

The contribution of the vertex a to this expression given that it is associated with parcel p is then

$$\frac{1}{2} \operatorname{Trace} \left(\frac{1}{m_i(p)} \vec{l}_p \vec{d}_i^T(a) C_i^{k+1} Z_i^k \right), \quad (\text{A5})$$

which gives us the first term in the edge strength $E_i(a, p)$ in (A1).

Appendix B. Supplementary material

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.neuroimage.2017.04.054.

References

- Amunts, K., Schleicher, A., Bürgel, U., Mohlberg, H., Uylings, H., Zilles, K., 1999. Broca's region revisited: cytoarchitecture and intersubject variability. *J. Comp. Neurol.* 412 (2), 319–341.
- Arslan, S., Parisot, S., Rueckert, D., 2015, June. Joint spectral decomposition for the parcellation of the human cerebral cortex using resting-state fMRI. In: *Information Processing in Medical Imaging*. Springer International Publishing, pp. 85–97.
- Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L., Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., 2013. Function in the human connectome: task-fMRI and individual differences in behavior. *Neuroimage* 80, 169–189.
- Beckmann, C.F., Mackay, C.E., Filippini, N., Smith, S.M., 2009. Group comparison of resting-state fMRI data using multi-subject ICA and dual regression. *Neuroimage* 47 (Suppl 1), S148.
- Besag, J., 1986. On the statistical analysis of dirty pictures. *J. R. Stat. Soc. Ser. B (Methodol.)*, 259–302.
- Biswal, B., Zerrin Yetkin, F., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* 34 (4), 537–541.
- Blumensath, T., Jbabdi, S., Glasser, M.F., Van Essen, D.C., Ugurbil, K., Behrens, T.E., Smith, S.M., 2013. Spatially constrained hierarchical parcellation of the brain with resting-state fMRI. *Neuroimage* 76, 313–324.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* 3 (1), 1–122.
- Boykov, Y., Veksler, O., Zabih, R., 2001. Fast approximate energy minimization via graph cuts. *IEEE Trans. Pattern Anal. Mach. Intell.* 23 (11), 1222–1239.
- Brodman, K., 1909. Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. Barth.
- Cotter, S.F., Rao, B.D., Engan, K., Kreutz-Delgado, K., 2005. Sparse solutions to linear inverse problems with multiple measurement vectors. *IEEE Trans. Signal Process.* 53 (7), 2477–2488.
- Craddock, R.C., James, G.A., Holtzheimer, P.E., Hu, X.P., Mayberg, H.S., 2012. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Human Brain Mapp.* 33 (8), 1914–1928.
- Eickhoff, S.B., Bzdok, D., Laird, A.R., Roski, C., Caspers, S., Zilles, K., Fox, P.T., 2011. Co-activation patterns distinguish cortical modules, their connectivity and functional differentiation. *Neuroimage* 57 (3), 938–949.
- Eickhoff, S.B., Thirion, B., Varoquaux, G., Bzdok, D., 2015. Connectivity-based parcellation: critique and implications. *Human Brain Mapp.* 36 (12), 4771–4792.
- Fischl, B., 2012. FreeSurfer. *Neuroimage* 62 (2), 774–781.
- Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* 6, 721–741.
- Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178.
- 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., 2013. The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* 80, 105–124.
- Gordon, E.M., Laumann, T.O., Adeyemo, B., Petersen, S.E., 2015. Individual variability of the system-level organization of the human brain. *Cereb. Cortex*, bhv239.
- Harrison, S.J., Woolrich, M.W., Robinson, E.C., Glasser, M.F., Beckmann, C.F., Jenkinson, M., Smith, S.M., 2015. Large-scale probabilistic functional modes from resting state fMRI. *Neuroimage* 109, 217–231.
- Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P., 2009. Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci. USA* 106 (6), 2035–2040.
- Kim, J.H., Lee, J.M., Jo, H.J., Kim, S.H., Lee, J.H., Kim, S.T., Seo, S.W., Cox, R.W., Na, D.L., Kim, S.L., Saad, Z.S., 2010. Defining functional SMA and pre-SMA subregions in human MFC using resting state fMRI: functional connectivity-based parcellation method. *Neuroimage* 49 (3), 2375–2386.
- Laumann, T.O., Gordon, E.M., Adeyemo, B., Snyder, A.Z., Joo, S.J., Chen, M.Y., Gilmore, A.W., McDermott, K.B., Nelson, S.M., Dosenbach, N.U., Schlaggar, B.L., 2015. Functional system and areal organization of a highly sampled individual human brain. *Neuron* 87 (3), 657–670.
- Mueller, S., Wang, D., Fox, M.D., Yeo, B.T., Sepulcre, J., Sabuncu, M.R., Shafee, R., Lu, J., Liu, H., 2013. Individual variability in functional connectivity architecture of the human brain. *Neuron* 77 (3), 586–595.
- Ng, B., Varoquaux, G., Poline, J.B., Thirion, B., 2012. A novel sparse graphical approach for multimodal brain connectivity inference. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI*. Springer Berlin Heidelberg, pp. 707–714.
- Penfield, W., Rasmussen, T., 1950. *The Cerebral Cortex of Man; a Clinical Study of Localization of Function*.
- Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., Shulman, G.L., 2001. A default mode of brain function. *Proc. Natl. Acad. Sci. USA* 98 (2), 676–682.
- Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* 52 (3), 1059–1069.
- Shattuck, D.W., Leahy, R.M., 2002. BrainSuite: an automated cortical surface

- identification tool. *Med. image Anal.* 6 (2), 129–142.
- Shen, X., Papademetris, X., Constable, R.T., 2010. Graph-theory based parcellation of functional subunits in the brain from resting-state fMRI data. *Neuroimage* 50 (3), 1027–1035.
- Shou, H., Eloyan, A., Nebel, M.B., Mejia, A., Pekar, J.J., Mostofsky, S., Caffo, B., Lindquist, M.A., Crainiceanu, C.M., 2014. Shrinkage prediction of seed-voxel brain connectivity using resting state fMRI. *NeuroImage* 102, 938–944.
- Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W., 2011. Network modelling methods for FMRI. *Neuroimage* 54 (2), 875–891.
- Smith, S.M., Vidaurre, D., Beckmann, C.F., Glasser, M.F., Jenkinson, M., Miller, K.L., Nichols, T.E., Robinson, E.C., Salimi-Khorshidi, G., Woolrich, M.W., Barch, D.M., 2013. Functional connectomics from resting-state fMRI. *Trends Cogn. Sci.* 17 (12), 666–682.
- Sporns, O., Tononi, G., Kötter, R., 2005. The human connectome: a structural description of the human brain. *PLoS Comput. Biol.* 1 (4), (p.e42).
- Stevens, W.D., Tessler, M.H., Peng, C.S., Martin, A., 2015. Functional connectivity constrains the category-related organization of human ventral occipitotemporal cortex. *Hum. Brain Mapp.* 36 (6), 2187–2206.
- Van Essen, D.C., Glasser, M.F., Dierker, D.L., Harwell, J., Coalson, T., 2012. Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cereb. Cortex* 22 (10), 2241–2262.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium, 2013. The WU-Minn human connectome project: an overview. *Neuroimage* 80, 62–79.
- Von Luxburg, U., 2007. A tutorial on spectral clustering. *Stat. Comput.* 17 (4), 395–416.
- von Economo, C.F., Koskinas, G.N., 1925. Die cytoarchitektur der hirnrinde des erwachsenen menschen. J. Springer.
- 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., 2015. Parcellating cortical functional networks in individuals. *Nat. Neurosci.*
- Whittaker, J., 2009. *Graphical Models in Applied Multivariate Statistics*. Wiley Publishing.
- Wig, G.S., Laumann, T.O., Petersen, S.E., 2014. An approach for parcellating human cortical areas using resting-state correlations. *Neuroimage* 93, 276–291.
- Woolrich, M.W., Ripley, B.D., Brady, M., Smith, S.M., 2001. Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage* 14 (6), 1370–1386.
- Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., Fischl, B., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* 106 (3), 1125–1165.
- Yu, S.X., Shi, J., 2003, October. Multiclass spectral clustering. In: *Proceedings of the Ninth IEEE International Conference on Computer Vision*. IEEE, pp. 313–319.
- Zilles, K., Amunts, K., 2010. Centenary of Brodmann's map—conception and fate. *Nat. Rev. Neurosci.* 11 (2), 139–145.