Published in final edited form as:

J Neurosci Methods. 2017 February 01; 277: 1–20. doi:10.1016/j.jneumeth.2016.11.014.

sGraSP: a Graph-based Method for the Derivation of Subject-specific Functional Parcellations of the Brain

N. Honnorat^a, T. D. Satterthwaite^{a,b}, R. E. Gur^b, R. C. Gur^b, and C. Davatzikos^a

^aCenter for Biomedical Image Computing and Analytics, Department of Radiology, University of Pennsylvania, Philadelphia PA 19104, USA

^bBrain and Behavior Laboratory, Department of Psychiatry, University of Pennsylvania, Philadelphia PA 19104, USA

Abstract

Background—Resting-state fMRI (rs-fMRI) has emerged as a prominent tool for the study of functional connectivity. The identification of the regions associated with the different brain functions has received significant interest. However, most of the studies conducted so far have focused on the definition of a common set of regions, valid for an entire population. The variation of the functional regions within a population has rarely been accounted for.

New Method—In this paper, we propose sGraSP, a graph-based approach for the derivation of subject-specific functional parcellations. Our method generates first a common parcellation for an entire population, which is then adapted to each subject individually.

Results—Several cortical parcellations were generated for 859 children being part of the Philadelphia Neurodevelopmental Cohort. The stability of the parcellations generated by sGraSP was tested by mixing population and subject rs-fMRI signals, to generate subject-specific parcels increasingly closer to the population parcellation. We also checked if the parcels generated by our method were better capturing a development trend underlying our data than the original parcels, defined for the entire population.

Comparison with Existing Methods—We compared sGraSP with a simpler and faster approach based on a Voronoi tessellation, by measuring their ability to produce functionally coherent parcels adapted to the subject data.

Conclusions—Our parcellations outperformed the Voronoi tessellations. The parcels generated by sGraSP vary consistently with respect to signal mixing, the results are highly reproducible and

Corresponding authors: N. Honnorat and C. Davatzikos; Phone: (215) 349 8646; Fax: (215) 614 0266; 3600 Market Street, Suite 380, Philadelphia, PA 19104.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Conflict of interest

All authors declare no financial interests or potential conflicts of interest.

All other study authors report no competing interests.

the neurodevelopmental trend is better captured with the subject-specific parcellation, under all the signal mixing conditions.

Keywords

rs-fMRI; parcellation; tessellation

1. Introduction

Functional MRI is a prominent tool for studying brain connectivity. This imaging modality measures the oxygen consumption almost in real time and simultaneously in the whole brain with a fairly good resolution. Resting-state fMRI (rs-fMRI) focuses on the description of the brain activation patterns that are observed when the subjects are asked to not focus on a particular task (Fox and Raichle, 2007). It was shown that the networks activated during rest span the whole brain and correspond quite well to the networks that were described using task fMRI (Smith et al., 2009). For this reason, rs-fMRI is nowadays considered as a major tool for studying brain organization, in both health (Yeo et al., 2011; Power et al., 2011) and disease (Leistedt et al., 2009; Baker et al., 2014; Lynall et al., 2010; Varoquaux et al., 2010a).

Several approaches have been proposed for analyzing how the complex tasks performed by the brain and observed via rs-fMRI decompose into simpler tasks, processed by smaller brain regions interacting dynamically. Independent Component Analysis (ICA) was for instance used for decomposing the brain into subnetworks presenting independent activation dynamics (Calhoun et al., 2001; Smith et al., 2012). Dictionary learning methods were proposed for decomposing brain connectivity into sparse brain networks (Varoquaux et al., 2011). Following the segregation and integration principles (Tononi et al., 1994) a large number of functional parcellation methods partitioning the brain into connected non-overlapping regions were developed (Langs et al., 2015; Baumgartner et al., 1997), in particular methods based on normalized cuts (Thirion et al., 2006; van den Heuvel et al., 2008; Shen et al., 2013; Craddock et al., 2012), based on K-means clustering (Kim et al., 2010; Bellec et al., 2010; Mezer et al., 2009), based on mixture models (Lashkari et al., 2010), on statistical models (Baldassano et al., 2015), on region growing (Blumensath et al., 2013), on hierarchical clustering (Salvador et al., 2005; Cordes et al., 2002) and on Markov Random Fields (Honnorat et al., 2015; Parisot et al., 2016). In order to study the effects of aging and disease on the brain function, the ICA components, dictionary loadings and connected regions are commonly considered as the nodes of a high-level connectivity graph. A graphical theoretical analysis is then performed for extracting significant group differences (Bullmore and Sporns, 2009; Rubinov and Sporns, 2010) or for further grouping strongly connected nodes into communities (Olhede and Wolfe, 2014). Because community detection algorithms are not all able to handle full brain resolution and do not take the spatial proximity of the nodes into account when grouping them, these algorithms are often used for connecting brain parcels into networks. There is a strong synergy between community detection and parcellation methods: community detection become more effective when applied after brain parcellation, by virtue of starting with parcels that are optimally

partitioned to be functionally homogeneous. Networks of parcels are also more interpretable than networks of voxels, particularly in the presence of noise.

When applied to large populations of several hundred of subjects, these methods successfully identify stable networks (Yeo et al., 2011) and prominent foci of activity (Power et al., 2011) that are now often used as templates. However, the shape and the location of the functional units can vary quite a lot within a population. This large inter-subject functional variability cannot be handled correctly when a unique segmentation is defined for the entire population. This concern prevents the definition of highly reliable functional biomarkers. For this reason, interest has emerged recently for the definition of subject-specific parcels, activity foci and networks (Beckmann et al., 2009; Filippini et al., 2009; Schultz et al., 2014; Langs et al., 2015) with the aim of developing diagnosis/prognosis frameworks exploiting individual functional differences (Finn et al., 2015). The study of subject-specific variations in regionally-coherent functional units can also provide biomarkers of various diseases associated with loss of functional coherence.

The derivation of subject-specific networks and parcels is still an active topic of research (Parisot et al., 2015; Arslan et al., 2015; Parisot et al., 2016; Schultz et al., 2014; Wang et al., 2015). A relatively simple approach consists in assigning the voxels of the individual scans to the population parcels directly, according to the shape of their correlation maps (Wang et al., 2015). In this recent work, a population parcellation is first registered to each brain and an average time series is computed for each network/parcel. The parcels/networks are then iteratively deformed by reassigning their voxels to the network/parcel with the most correlated signal and updating the average time series. Most of the other approaches in the literature impose a form of smoothness when deriving subject-specific components. For instance, some methods estimate the individual variations at the very time of the atlas construction (Varoquaux et al., 2011; Parisot et al., 2015; Arslan et al., 2015; Parisot et al., 2016). The subject-specific parcellations and dictionary loadings are then obtained as a co-product of the population parcellation, and they benefit from similar smoothness properties. Dual regression was proposed for deriving subject-specific ICA components by adapting a set of population ICA components to individual data (Beckmann et al., 2009; Filippini et al., 2009). Template based rotation (TBR) performs a matching in the opposite direction: during TBR, the individual rs-fMRI scans are matched to the template component by performing a procrustean rotation of their spatial PCA components (Schultz et al., 2014). For both approaches, the smoothness of the templates induces smooth individual ICA components.

Most of these methods either impose a strong similarity between the individual parcellations and the group parcellation by assuming a spatial similarity (Beckmann et al., 2009; Filippini et al., 2009), or rely on a kind of functional matching that constrains the subject parcellations (Parisot et al., 2015; Arslan et al., 2015; Parisot et al., 2016). On the contrary, we propose to weakly constrain the individual parcellation, in order to capture the relevant subject variability as much as possible. We impose the location of only one node per parcel and let the parcellation adapt freely to the subject data while ensuring that the parcels remain connected. We demonstrate that, when such constraint is necessary, the similarity between subject and population parcellations can be enforced by mixing their time series. This

strategy allows us to explore a large range of deformations that the population parcellations can undergo for fitting subject data.

Our framework extends the population parcellation method GraSP (Honnorat et al., 2015). We propose to derive subject-specific parcellations from the parcellations generated by GraSP for the entire population. Our method, which will be referred as subject-specific GraSP (sGraSP), is seed-based: we assume that the functional centers of the subject-specific parcels are located at the same place as the functional center determined by GraSP for the entire population (Honnorat et al., 2015). This property ensures that the generated parcels are comparable across the population. We compare sGraSP with two simpler approaches, generating subject-specific parcellation by tessellating the brain according to a data-dependent geodesic distance. In addition to these distance maps, sGraSP exploits connectivity maps measuring the connectivity between the seeds and the entire brain directly. The partition is then solved as a discrete optimization problem combining these two sources of information. The fact that sGraSP outperforms the tessellation approaches for a large sample of the participants (n=100) recruited as part of the Philadelphia Neurodevelopmental Cohort (PNC) (Satterthwaite et al., 2014) demonstrates that this additional complexity of sGraSP is useful.

The remainder of this paper is organized as follows. In section 2, we present sGraSP and the tessellation approaches for deriving subject-specific parcellations. The experimental results are presented in section 3. Discussions conclude the paper.

2. Methods

2.1. Model Assumption and Related Work

Inter-subject variability is one of the most challenging issues for the extraction of functional biomarkers. The study of the cytoarchitecture of the brain (Zilles and Amunts, 2015) has confirmed that most brain areas defined by Brodmann (Brodmann, 1909) and Economo and Koskinas (Economo and Koskinas, 1925) significantly vary in location and shape from subject to subject (Amunts et al., 1999; Bürgel et al., 2006; Scheperjans et al., 2008; Henssen et al., 2016). This variability impacts brain function (Mueller et al., 2013).

Several approaches can be adopted in response to this issue. Let us assume that a template/population parcellation is available, such as the seven networks of Yeo et al. (Yeo et al., 2011) or a data-driven brain segmentation provided by a parcellation method (Honnorat et al., 2015; Parisot et al., 2015; Arslan et al., 2015; Parisot et al., 2016). The simplest approach, adopted so far in most of the studies, considers that the inter-subject variability is captured by the variation of the features extracted from the population parcellation. According to this approach, adapting the parcellations to the subjects for deriving subject-specific regions of interest is not necessary/useful.

In an attempt of decoupling the variation of the shape of the regions from the variation of their content (their functional homogeneity, their correlations with other regions, etc.) the second approach still assumes that a common set of regions can be described for all the subjects of a population and that these regions are anchored in a fixed location. However, the

regions are assumed to vary in shape and extension, which implies that their geometric center and area can change across the population. The population parcellation needs to be adapted to each subject individually, to capture the variability of each brain.

A larger degree of adaptation can be obtained by relaxing the anchoring constraint. This third approach can be implemented with a non-rigid registration, by matching patches of data extracted from the individual scans with patches extracted from the population (Wang et al., 2015), by refining ICA components for each subject scan independently as in dual regression (Beckmann et al., 2009; Filippini et al., 2009) or by refining the parcellations iteratively, which allows subject parcellations to diverge freely and independently from the population template (Wang et al., 2015).

Finally, a complete adaptation could be achieved by letting the number of parcels adapt to each brain as well. This last approach, however, complicates the definition of a common set of biomarkers, which is necessary for comparing subjects and for deriving population statistics. To our knowledge, it was therefore never adopted broadly. This approach will not be further considered in this paper.

In this work, we adopt the second approach and we propose a method that could be easily extended when more flexibility is required. We make the assumption that each brain region has an anchor, a cortical location that we refer as “functional center” and that should be the same for all the subjects, by definition. In the meantime, the shapes/extents of the cortical regions are free to adapt to each individual. We show that our approach can extract more reliable biomarkers than the fixed parcellation of the first approach, in particular when small brain regions are studied. These results suggest that adapting the shape of the parcels might significantly improve the quality of regional functional biomarkers subsequently extracted. In the future, a larger degree of flexibility could be introduced by adapting the location of the functional centers as well, for instance by matching patches surrounding the functional centers (Wang et al., 2015) or by registering subject and population data (Sabuncu et al., 2010; Langs et al., 2010).

2.2. Notations

In this work, we assume that n_s rs-fMRI scans, acquired for n_s different subjects s , have been registered and projected to a common mesh $\mathcal{M} = (V, E)$ containing n_p nodes/vertices. We denote with n_t the number of time points of the scans, and we assume that the $n_p n_t$ -long time series of each subject have been normalized to zero mean and unit Euclidean norm.

Our goal is to produce a segmentation of the mesh adapted to each subject. We also want these individual segmentations to be comparable across individuals, in order to extract a corresponding set of features and networks for all the subjects. As explained in the previous section, we make the assumption that the same parcels can be defined for all the subjects of the population. Under this assumption, we produce first a segmentation for the whole population that is then adapted to all the subjects individually.

The individual Blood Oxygenation Level Dependent (BOLD) time series are first normalized to zero mean and unit Euclidean norm, concatenated and normalized again to

unit Euclidean norm (Kim et al., 2010; Yeo et al., 2011; Lashkari et al., 2010; Honnorat et al., 2015). This standard procedure generates population time series such that the correlation between two brain locations, at the population level, is the average of the correlations between these locations observed at the subject level. As a result, when correlations are directly used for measuring brain connectivity, the population connectivity emerges as a simple average of the subject brain connectivity. For large number of subjects, we observed (not shown here) that this procedure produces very similar correlations as the averaging after Fisher z-transform (Clayton Silver and Dunlap, 1987) preferred by some authors (Shou et al., 2014), in particular because most of the correlations are weak or moderate and therefore close to their Fisher z-transform. For high-resolution rs-fMRI scans and a large number of subjects, the concatenation is a little more computationally efficient.

The graph-based method GraSP is then applied for generating a population-level parcellation (Honnorat et al., 2015). GraSP relies on a Markov Random Field framework. This method has the advantage of selecting a functional center for each parcel: a node of the brain mesh presenting a BOLD time series that can be identified with the activity of the whole parcel. In this paper, we compare two seed-based approaches exploiting this information. First, by modifying GraSP, we obtain a method that takes simultaneously into account (1) the direct connectivity between functional centers and the entire brain and (2) the local brain organization, thanks to the computation and exploitation of a set of geodesic distance maps reflecting the local brain connectivity. This entire approach, a population parcellation using GraSP followed by a refinement for each subject, will be referred as subject-specific GraSP (sGraSP).

Then, noticing that a direct tessellation of these geodesic distances maps produces valid partitions of the brain as well, we propose to focus only on the computation of geodesic distance maps. This second approach is much faster since no optimization is required.

These two approaches, Voronoi tessellations based on the population functional centers and sGraSP parcellations, were compared in detail on a large sample of our database. We present in section 2.3 the two parcellation methods that were used during this work: the GraSP parcellation method (Honnorat et al., 2015) that was used for creating the population-level parcellations and the novel variant of GraSP dedicated to the parcellation of individual rs-fMRI data. The following section presents two alternative methods based on Voronoi tessellations for the segmentation of subject-specific data. They were derived from the first part of GraSP only so that they do not require the use of a Markov Random Field solver. The database used for this work is described in section 2.5. Section 2.6 presents the metric that was used for comparing the subject-specific parcellations with the population level parcellations. Section 2.7 describes the neurodevelopmental trend that was searched in our cohort, in order to demonstrate the advantages of deriving subject-specific parcellations for clinical applications. The last section, section 2.8, explains how a set of additional data was generated by mixing population and individual rs-fMRI signal. These additional test conditions allowed us (1) to investigate the behavior of our optimization framework when the individual scans become increasingly distinct from the population data and (2) to assess the reliability of the neurodevelopmental trend detected with our subject-specific parcellation method.

2.3. Graph-based Parcellation with Shape Priors

In this section, we review the GraSP method (Honnorat et al., 2015) and we describe sGraSP, a variant for the parcellation of individual rs-fMRI scans. In order to produce comparable parcellations, we derive first a population-level parcellation, by segmenting the concatenation of all the subject time series. The subject-specific parcellations are produced by a variant of GraSP that is constrained for producing parcels with the same 'functional center' as the corresponding population-level parcels. For GraSP, the BOLD signal of the 'functional center' is by definition the signal of the parcel 'centered at this node'. Because the parcels are constrained to be *geodesically* star-convex with respect to their functional center (Veksler, 2008; Gulshan et al., 2010), the functional center tends to be selected near the geometrical center of the parcel, but this is not a hard constraint: the shape of the parcel can freely adapt to the rs-fMRI signal and can adopt tortuous shapes if required, depending on the local correlation between BOLD time series. This freedom was illustrated in (Honnorat et al., 2015) for a synthetic dataset. In this work, we take advantage of this property again: by constraining only the parcel centers, we define comparable parcels across the population, while letting the parcel shapes freely fit individual data.

2.3.1. GraSP Model—GraSP is a graph parcellation method based on a discrete Markov Random Field (MRF) framework. In this framework, each graph node p is considered as a candidate parcel functional center and associated with an index. The parcels are identified by these indices. By definition of the functional center, the BOLD signal of a parcel is the signal of the functional center of the parcel. The parcellation is solved as a labeling problem: for each graph node p , a label l_p indicating what parcel should contain p is searched. These labels are determined by minimizing an energy of the form:

$$E(\{l_p\}) = \sum_p V_p(l_p) + \sum_q L_q(\{l_p\}) + \sum_q S_q(\{l_p\}) \quad (1)$$

where the costs $V_p(l_p)$ diminishes when the rs-fMRI signal y_p of the node p and the signal y_{l_p} of the parcel l_p where it lies are highly correlated (Honnorat et al., 2015):

$$V_p(l_p) = 1 - \rho(y_p, y_{l_p}) \quad (2)$$

Where $\rho(y_p, y_{l_p})$ denotes the Pearson correlation between the rs-fMRI signal of the mesh node p and the rs-fMRI signal of the mesh node l_p . The costs $L_q(\{l_p\})$ penalize all the parcels that are not empty by a similar cost K which is the only parameter of the method and indirectly fixes the refinement, the scale of the parcellation (Delong et al., 2012; Ryali et al., 2013):

$$L_q(\{l_p\}) = \begin{cases} K & \text{if } \exists p \text{ such that } l_p = q \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

The last costs, $S_q(\{I_p\})$, are shape priors that are set large enough for acting as hard constraints. These shape priors are designed for enforcing the star convexity of the parcels with respect to their functional center (Veksler, 2008). This property ensures that the parcels are not fragmented and that the number of parcels is actually controlled by the cost K . The fit of the shape prior to the data was improved by using a geodesic distance (Gulshan et al., 2010): the geodesic distance computed from the Pearson distance between neighboring cortical mesh nodes (Honnorat et al., 2015). The use of a geodesic distance makes also star shape priors able to segment very tortuous connected objects containing even loops and holes Honnorat et al. (2015).

2.3.2. Geodesic Star-convex Shape Priors—The construction of the shape priors $S_q(\{I_p\})$ starts with the definition of a distance (more precisely: a pseudometric) between neighboring nodes of the cortical mesh. In this work, we have used the Pearson distance between the BOLD signals of these nodes. If p and q are two neighboring mesh nodes, y_p and y_q their normalized rs-fMRI signal (to zero mean and unit Euclidean norm) and $\langle \cdot, \cdot \rangle$ the standard inner product, the Pearson distance between p and q is defined by:

$$\pi(p, q) = 1 - \langle y_p, y_q \rangle \quad (4)$$

For each parcel functional center s , a shape prior ensures that when a node p is assigned to the parcel of functional center s , the nodes that are part of the shortest path between p and s (according to the Pearson distance) are also assigned to the parcel s . This property, also known as star convexity (with respect to the center s), can be expressed as follows, where $p \rightarrow s$ indicates that p is assigned to the center s and $\mathcal{P}(p, s)$ the shortest path between p and s :

$$\text{star convexity: } \forall p, [p \rightarrow s] \Rightarrow [\forall q \in \mathcal{P}(p, s), q \rightarrow s] \quad (5)$$

Star convexity ensures the connectedness of the parcels because it guarantees that any two nodes belonging to the same parcel can be connected, at least through their parcel functional center. Star shape priors are easy to introduce in a Markov Random Field framework, because they decompose into a set of simple pairwise constraints that can be handled by standard MRF optimizers (Boykov et al., 2001; Kolmogorov and Zabih, 2004; Boykov and Kolmogorov, 2004; Delong et al., 2012; Veksler, 2008; Honnorat et al., 2015). During our experiments, the shortest path distances were computed, similarly to Johnson's algorithm (Johnson, 1977), by using Dijkstra's algorithm repeatedly for each parcel functional center. We reduced the computational burden significantly by using an advanced implementation of Dijkstra's method based on Fibonacci heaps (Fredman and Tarjan, 1987).

So far, GraSP was only used for producing group-level parcellation, by parcellating the concatenation of the time series of several subjects (Honnorat et al., 2015). In the following section, we explain how the parcellation obtained at the population level can be fitted to each subject, in order to produce subject-specific parcellations that are comparable across subjects and can be used for defining subject-specific geometric and connectivity-based features. This subject-specific GraSP variant will be referred as sGraSP.

2.3.3. sGraSP: Constrained Subject-specific Parcellation—Assuming that the brains of a population of subjects have been registered, population-level cortical segmentations are commonly found by applying the parcellation methods to the time series obtained by concatenating, at each brain location, the normalized BOLD signals of all the subjects of the population (Baumgartner et al., 1997; Thirion et al., 2006; Shen et al., 2013; Yeo et al., 2011; Lashkari et al., 2010), the parcellation-level correlation matrices produced in that way being equal to the average of the individual correlation matrices. The parcellations obtained in that way were found to be quite reproducible (Blumensath et al., 2013; Honnorat et al., 2015).

In order to generate subject-specific parcellations, we have modified the original GraSP implementation for imposing the population-level parcel centers, by (1) assigning a very large label cost to the nodes that should not be selected as parcel centers, (2) reducing the label cost of the desired parcel centers to zero and (3) penalizing the assignment of these nodes to other parcels with very large singleton costs. The definitions (2) and (3) were thus modified as follows:

$$V_p(l_p) = \begin{cases} 1 - \langle y_p, y_{l_p} \rangle & \text{if } l_p \in L \\ \infty & \text{otherwise} \end{cases} \quad (6)$$

$$L_p(\{l_p\}) = \begin{cases} 0 & \text{if } q \in L \\ \infty & \text{otherwise} \end{cases} \quad (7)$$

where L is the list of parcel centers inherited from the population parcellation. In addition, restricting the optimization to the desired parcel centers significantly reduced the computational burden.

When developing our method, we noticed that, because the subject-level parcel centers are fixed, the geodesic maps produced by GraSP could be exploited directly without any optimization step. We propose two tessellations methods following this approach in the next section. Detailed experiments were carried out for checking if these methods can achieve good performances or if the complete optimization performed by GraSP is necessary. Their results are presented in section 3.

2.4. Tessellations

GraSP shape priors require the computation of the geodesic distance D . Here, we propose to exploit the geodesic distance maps directly, without additional optimization, by assigning the cortical nodes to the closest parcel functional center inherited from the population-level parcellation. The resulting segmentation corresponds to the Voronoi diagram of the parcel centers, also called Voronoi tessellation. During our experiments, two families of geodesic pseudometrics were computed. A pseudometric exhibits all the properties of a metric, except the distinguishability: two different mesh nodes can have a null pseudometric distance.

The first family, that will be denoted by d , is derived from the geodesic Pearson distance used by GraSP. It was defined for all pairs of neighboring nodes p and q of the cortical mesh as follows:

$$d(p, q; c) = (1 - \langle y_p, y_q \rangle)^c = \pi(p, q)^c \quad (8)$$

where $c \geq 1$. A geodesic pseudometric between the parcel centers and all the nodes of the cortical mesh was computed as follows. Similarly to the geodesic Pearson distance used by GraSP, the following distance between a functional center p and any mesh node q was computed by an implementation of Dijkstra's algorithm with Fibonacci heaps (Fredman and Tarjan, 1987):

$$D(p, q; c) = \min_{(x_1, x_2, \dots, x_n) \in \mathcal{P}(p, q)} \sum_{i=1}^{n-1} d(x_i, x_{i+1}; c) \quad (9)$$

where $\mathcal{P}(p, q)$ is the set of paths between p and q in the mesh $\mathcal{M} = (V, E)$:

$$\mathcal{P}(p, q) = \{(x_1, x_2, \dots, x_n) \in \mathbb{P}(V) \mid x_1 = p, x_n = q, (x_i, x_{i+1}) \in E \forall i = 1 \dots n-1\}$$

where $\mathbb{P}(V)$ denotes the power set of the set of mesh vertices V . By definition, the distance D between nodes separated by a path of highly correlated nodes is small. This property guarantees the functional coherence of the Voronoi cells.

The second family that we have investigated, δ , is asymmetric and depends only on the correlation between the local signal and the parcel functional center. δ is derived from a quasimetric defined for each parcel functional center r and for $c \geq 1$:

$$\delta(p, q; r; c) = (1 - \langle y_r, y_q \rangle)^c = \pi(r, q)^c \quad (10)$$

As a quasimetric, δ exhibits all the properties of a distance (positivity, positive definiteness, triangle inequality) except the symmetry. Dijkstra's shortest path algorithm was used for defining a geodesic pseudo(quasi)metric between the parcel functional center r and all the mesh nodes, for all the parcel centers:

$$\Delta(p, q; r; c) = \min_{(x_1, x_2, \dots, x_n) \in \mathcal{P}(p, q)} \sum_{i=1}^{n-1} \delta(x_i, x_{i+1}; r; c) \quad (11)$$

These quasimetrics increase slowly when traveling in the mesh through cortical nodes highly correlated with the parcel center. As a result, when nodes are highly correlated with a functional center and connected with this center, they have a strong likelihood of being assigned to its parcel/cell.

During our experiments, several $c = 1$ were tested as well. For these values, the local $d(p, q; c)$ is not a metric anymore, because the triangle inequality does not hold. However, $\delta(p, q; r; c)$ remains a quasimetric and D remains pseudo metric. Six geodesic distances used in this work are illustrated in the figure (1). The second tessellation is slower than the first one because it requires the computation of a metric between each pair of neighboring nodes for every functional parcel center inherited from the population-level parcellation. It is unclear if this approach is faster than sGraSP proposed in the previous section since it requires more distance computations, but no MRF optimization. By contrast, the first tessellation is necessarily faster than sGraSP, because it performs the same geodesic distance computation without performing any MRF optimization. In general, because all the methods tested here (tessellations and sGraSP) focus on the set of parcel centers defined at the population level, they are much faster than the original GraSP that considers all the cortical nodes as candidate parcel centers.

We compared the energy of the tessellations produced by the different approaches as explained in the following sections. In particular, our goal was to determine if the tessellations can provide good subject-specific segmentations, or if an optimization is necessary for adapting the population-level parcels.

2.5. Data

The methods presented in this paper were validated with 859 subjects between the ages of 8 and 23 years, from the Philadelphia Neurodevelopmental Cohort (Satterthwaite et al., 2014). For each subject, a T1 image was acquired prior to a 6 minute long fMRI scan at TR of 3000 ms (120 time-points). Each fMRI scans was first registered to its related T1 using boundary-based registration (Greve and Fischl, 2009) with distortion correction provided by FSL 5 (Jenkinson et al., 2012). The T1 images were registered to the FreeSurfers fsaverage5 cortical template (10,242 nodes/hemisphere) (Dale et al., 1999). The fMRI subject time series was projected to the fsaverage surface space using FreeSurfers *mri vol2surf* command. Time series data was processed using a validated confound regression procedure that has been optimized to reduce the influence of subject motion, which is of particular concern for studies of development. The first four volumes of the functional time series were removed to allow signal stabilization, leaving 120 volumes for subsequent analysis. Functional images were realigned using MCFLIRT14 and band-pass filtered to retain frequencies between 0.01 and 0.08 Hz. Improved confound regression (Satterthwaite et al., 2013) included nine standard confounding signals (6 motion parameters + global /WM/ CSF) as well as the temporal derivative, quadratic term, and temporal derivative of the quadratic of each (36 parameters total). Prior to confound regression, all motion parameters and confound time courses were band-pass filtered in an identical fashion as the time series data itself in order to prevent mismatch in the frequency domain and allow the confound parameters to best fit the retained signal frequencies. No additional spatial smoothing of the time series was performed.

Given that the left and right hemispheres are segmented separately by FreeSurfer, all the experiments presented here were performed independently for the two hemispheres. This allowed us to replicate all our results.

2.6. Parcellations Comparison

Many different metrics can be computed for comparing parcellations (Pfitzner et al., 2009). In this paper, the similarity between parcellations was measured by the adjusted Rand index, an index that measures the overlap of multiple structures simultaneously (Hubert and Arabie, 1985). This index can be computed very efficiently in the following manner. Let $X = \{x\}$ and $Y = \{y\}$ denote two parcellations, where x and y denote parcels. Let $|x|$ denote the number of nodes of the parcel x . The mismatch matrix ($m_{x,y}$) is first computed, by counting the number of nodes included in the intersection between the parcels of the two parcellations:

$$m_{x,y} = |x \cap y|$$

The adjusted Rand index is then computed from $m_{x,y}$ and the sums $a_x = \sum_y m_{x,y}$ and $b_y = \sum_x m_{x,y}$ as follows:

$$\text{adjusted Rand}(X, Y) = \frac{\sum_{x,y} \binom{m_{x,y}}{2} - \left[\sum_x \binom{a_x}{2} \sum_y \binom{b_y}{2} \right] / \binom{N}{2}}{\frac{1}{2} \left[\sum_x \binom{a_x}{2} + \sum_y \binom{b_y}{2} \right] - \left[\sum_x \binom{a_x}{2} \sum_y \binom{b_y}{2} \right] / \binom{N}{2}}$$

2.7. Neurodevelopment Trend

The functional units of the brain are commonly identified by parcellating the cortex into functionally coherent regions (Blumensath et al., 2013; Honnorat et al., 2015). Once these units have been extracted and their function has been described, it becomes possible to interpret the effects of development, aging and brain diseases. In particular, neurodevelopment is known to modify the strength of short range and long-range connectivity: as the brain matures, its segregation and integration improve. The better segregation is associated with a decrease of the connectivity between adjacent regions, that tend to become more specialized. Improved integration leads on the contrary to a reinforcement of long range connectivity (Fair et al., 2007, 2009). Because our parcellation and tessellation methods group adjacent brain regions together, the functional coherence of the parcels that we build measures directly the strength of local brain connectivity. This local functional coherence can be estimated by measuring the following energy:

$$e = \sum_p \left[1 - \langle y_p, y_{l_p} \rangle \right] \quad (12)$$

where l_p denotes the index of one of the parcel functional centers inherited from the population-level parcellation: the functional center of the parcel that contains the node p after the parcellation of the individual scan. This energy e diminishes when all the parcels are functionally coherent. In practice, because of our specific choice of singleton costs $V_p(l_p)$ in GraSP/sGraSP and because the number of parcels is fixed during the parcellation

of subject data, e is the objective function that is optimized by sGraSP, under the constraint that the parcels remain geodesically star-convex (Honnorat et al., 2015).

In our previous work, we have already observed that the energy increases very significantly with age in our database, by dividing the population into three age groups that were parcellated using GraSP (Honnorat et al., 2015). In this work, we measured this trend with more details. We searched for a linear effect of neurodevelopment on the energy, corrected for subject motion during the scan. As a result, the following linear regression was performed:

$$e \approx \beta_0 + \beta_1 \times \text{age} + \beta_2 \times \text{MRD} \quad (13)$$

where the mean relative displacement (MRD) is a measure of subjects motion computed by the motion correction algorithm used for pre-processing our data (Jenkinson et al., 2012; Satterthwaite et al., 2013). β_1 was considered as the amplitude of the neurodevelopment effects and the p-value associated with β_1 was reported.

2.8. Parcellation Robustness

In order to check that sGraSP produces parcels that are similar to the population parcels when the individual rs-fMRI signal is close to the population brain activity, we replaced the standard Pearson correlation matrix Σ_s used by GraSP by the following mixture: $\overline{\Sigma}_s = (1 - \lambda)\Sigma_s + \lambda\Sigma_p$ where Σ_p denotes the correlation matrix computed for the whole population and the parameter λ was varied between 0 and 1. During these experiments, setting λ to 0 led us to generate individual parcellation solely depending on subject resting state data. We checked that λ values close to 1 were producing functional units very similar to the population-level parcels. This signal mixing was also used for checking if the parcellations and the tessellations were able to produce good results in the case of attenuated inter-subject variability, which is a pre-requisite for capturing the natural inter-subject variability to its full extent.

3. Results

The experiments carried out during this work are presented in this section. We present first the population-level parcellations that were generated. The parcel centers defined by these parcellations were used for all the tessellations and subject-specific parcellations. The performances of the tessellation algorithms are presented in the section 3.2. Section 3.3 presents the results obtained with sGraSP. We demonstrate in particular that it is possible to generate subject-specific parcellations very similar to the population parcellations by mixing individual and population rs-fMRI signals. Small perturbations of the subject signal with population signal always lead to very moderate parcellation stretching. This result also suggests that our parcellation approach is very reliable/reproducible. Section 3.4 presents all the experiments related to neurodevelopment. We emphasize that all the experiments were performed independently for both brain hemispheres during this work. This replication demonstrates that the results presented here are strongly reliable.

3.1. Population Parcellations

In order to test our methods at different scales, the concatenated/population time series were parcellated ten times, for the following parcellation resolution/parameter: $K = 1, 2, \dots, 10$. Figure (3) recalls that increasing K correspond to larger parcels. All the numbers of parcels are shown in the table (2). The coarser parcellation contains 69 parcels, while the finer one contains 2001 parcels. This range is larger than the one adopted during our original experiments (Honnorat et al., 2015), and we think that it covers most of the standard applications. The parcel centers of these twenty parcellations (ten scales, two hemispheres) were used by the parcellation and tessellation methods described above for generating subject-specific segmentations, for different signal mixtures.

3.2. Tessellations Performances

In order to estimate if the tessellation approaches can produce low energy subject-specific parcellations, we segmented a large subset of our database, for both $d = (1 - \langle y_p, y_q \rangle)^c$ and $\delta = (1 - \langle y_p, y_q \rangle)^c$ distances for distance parameters covering a large range of values: $c = 0.1, 0.25, 0.5, 1.0, 2.0, 5.0, 10.0$. One hundred subjects were selected, which lead us to generate 2800 tessellations. A functionally coherent sample was obtained by retaining only the youngest subjects (8 to 11.5 years old). This strategy prevented gender discrepancies during adolescence from introducing variability in our results. The variability was further attenuated by mixing a lot the individual signal with population signal ($\lambda = 0.9$). The tessellations fit to individual data was evaluated by comparing their energy with the 'baseline' energy obtained when the population level parcellation is used for parcellating the subjects, without any individual adaptation.

The experiment was carried out for both hemispheres, using the parcels centers of the population-level parcellations of the scale $K = 5$, a value that produces a medium number of parcels, close to 155 parcels per hemisphere and similar to many fMRI studies (Power et al., 2011). According to the figure (4) the faster tessellation, based on the distance d , generates the worst results. The energy increases when the distance parameter c differs from two and, even at that optimal point, the tessellations energy is higher than the baseline. The results obtained with the δ metric are clearly better. The energy of the tessellation decreases when the parameter c increases. However, the baseline energy, obtained when subject data is segmented using the population-level parcellation, is never reached/improved. Given that only the youngest subjects of our database were selected, and that their differences with the population average was attenuated a lot by signal mixing ($\lambda = 0.9$), the inter-subject variability in this experiment was much smaller than in our entire dataset. The fact that a poor functional coherence was obtained even in such a coherent sample suggests that the tessellations are never producing regions that are highly functionally coherent.

These results can be partially explained by looking at the shape of the Voronoi cells produced by the tessellation. Figure 6 shows that for the distance d , when the distance parameter increases, the cells become anisotropic and so thin that at some point, they seem to be fragmented. When the parameter is reduced, the cells become on the contrary increasingly isotropic. Both extremes lead to parcels that are not functionally coherent. A similar evolution is observed with δ but, as shown in figure 7, even large values of distance

parameter c do not produce narrow cells. This second tessellation scheme is more stable. These observations are in agreement with the results presented in figure 5 presenting the similarity between the population level parcellations and the tessellations derived from their functional centers. We observed in particular that the tessellations with narrow cells are significantly less similar to the population parcellations than the best tessellations, obtained for the metric with large parameter values.

Despite their differences, though, none of the tessellations produced parcels more functionally coherent than the population-level parcels. We conclude that an optimization similar to the one performed by the subject-specific variant of GraSP presented in this paper is necessary. We demonstrate in the next section that contrary to the tessellations, this approach successfully improves the fit of the population-level parcellation.

3.3. Subject-specific Parcellations

The performances and the reliability of sGraSP were measured by parcellating the hundred youngest subjects, for twenty different signal mixtures, corresponding to $\lambda = 0.0, 0.05, 0.1, \dots, 0.95$ and for the parcel resolution corresponding to $K = 5$. The parcellations obtained were compared to the population level parcellations, by measuring their energy and the adjusted Rand index between them. Because the same sample as in the previous experiment was used, a comparison with the tessellation performances presented in the previous section is valid and straightforward. The λ were selected so that the entire range of admissible λ is covered, from pure individual signals ($\lambda = 0$) to slight perturbations of the population parcellation ($\lambda = 0.95$).

Figure 8 illustrates four subject-specific parcellations that were obtained for the same signal mixing. The comparison between the subject-specific parcellation energy and the energy obtained when the population-level parcellation is used for segmenting the subject data is shown in figure 9. These results demonstrate that, contrary to the tessellations, sGraSP is able to significantly and strongly improve the functional coherence of the parcels: a reduction of the parcellation energy by as much as 15% was observed. The amplitude of the subject adaptation diminishes almost linearly with λ because the population-level signal is increasingly mixed with the individual signals and the parcellation. We note also that the amplitude of the adaptation to individual data is quite similar across the population, the variance being low for all the λ values. This result demonstrates again the strong stability and reliability of our parcellation algorithm: if our method was producing random parcellations or not fitting the subject data correctly, the variance would be higher. We also note that the results are perfectly reproducible between the two hemispheres: very similar energy improvements are observed for the two hemispheres for similar signal mixing.

Figure 10 shows that the adjusted Rand also evolves as we hoped: for both hemispheres, the adjusted Rand index tends to one when λ increases. This result indicates that the variability of the subject-specific parcellations diminishes, as an increasing part of the subject-specific signal is replaced by the population average. It indicates that sGraSP is very stable and the parcellations produced are not random. We highlight also that most of the adjusted Rand indices are quite small (below 0.5). These low adjusted Rand index values indicate that the population-level parcels were deformed a lot by our subject-specific sGraSP.

3.4. Neurodevelopment

3.4.1. Parcellation Energy—As explained in section 2.7, a neurodevelopmental effect was searched at all the parcellation scales and for both hemispheres by regressing age from the energy of the subject-specific parcellations while correcting for subject motion in the scanner. The reliability of our findings was assessed by repeating the experiments five times, after five slight mixtures of the individual time series with the population BOLD signal. These experiments were carried on a subset of 107 subjects, one eighth of our database. Contrary to the previous experiments, we ordered the subjects by increasing age before selecting every eighth of them.

The results presented in Figure 11 indicate that parcellation energy is increasing and strongly correlated with age during neurodevelopment. The effect is more significant at small scale, which is coherent with the fact that the energy measured at small scale corresponds to the shortest brain connections, that are the most impacted by aging (Fair et al., 2007, 2009). The effect was stronger and more reliable for the left hemisphere, but both hemispheres exhibit the same trends. The amplitude of the regression coefficient decreases when the part of individual subject rs-fMRI signal is reduced. In addition, the significance of the results was almost identical for all the signal mixtures. These results once again highlight the robustness of our parcellation results with respect to signal mixing/perturbation.

3.4.2. Advantages of Subject-specific Parcellations—In order to demonstrate the benefit of deriving subject-specific parcellations, we compared the trends and p values presented in the previous section with the trend and p-values obtained when the population parcellation is used for segmenting the individual rs-fMRI scans directly, without any adaptation. Figure 12 shows that the increase of parcellation energy with age during neurodevelopment can be detected without adapting the population level parcellation to the individual data. However, the trends observed are slightly smaller (for the left hemisphere) and much less significant (for both hemispheres), because individual variations disturb the estimation of the parcel functional homogeneity. Adapting the parcels to the subject data helps to account for their shape and location variation across the population, and produces an estimation of their coherence that is closer to the reality. These results demonstrate that subject-specific parcellations better capture the local functional changes. The difference is particularly noticeable at the finer scales that are more important in practice because they better reflect the effects of neurodevelopment on short range connectivity.

3.4.3. Intra-parcel Correlation and Parcellation Entropy—In order to quantify the improvements introduced by subject-specific parcellations over the original population parcellation, we fixed the parameters and we measured, for each of the 859 subjects in our database, two criteria that were not optimized directly by GraSP: the intra-parcel average correlations and the entropy of the parcellation. Signal mixing was set at $\lambda = 0.2$ and the scale $K = 3$ was chosen. Intra-parcel average correlation was computed by measuring the correlation between all the pairs of nodes of the parcel, averaging their Fisher z-transform and applying the inverse Fisher z-transform to the result (Clayton Silver and Dunlap, 1987). Parcellation entropy was measured by normalizing the average parcels time series to zero

mean and unit Euclidean norm, quantizing the average times series obtained into a thousand bins and measuring the parcels time series entropy $H(T_c)$ as follows:

$$H(T_c) = \sum_{i=1}^{1000} -p(i;q) \log(p(i;q)) \quad (14)$$

where $p(i; q)$ denotes the fraction of the time points for which the normalized BOLD signal T_c of the parcel c falls in the bin i . The quantity $p(i; q) \log(p(i; q))$ was set to zero when the bin i is never populated, in other words when $p(i; q) = 0$. The parcellation entropy H_p was measured by summing the parcels entropies multiplied by the number of nodes of the parcels:

$$H_p = \sum_{\text{parcel } c} \#(c) H(T_c) \quad (15)$$

where the number of nodes of the parcel c is denoted by $\#(c)$. Subject-specific parcellation entropies were compared with population parcellation entropies, with the entropies measured for ten random permutations of the subject-specific parcellations, with the entropies measured for ten random permutations of the population parcellation and with full brain entropies H_B , measured by summing the entropies of the node time series y_p , $H(y_p)$, over the brain:

$$H_B = \sum_p H(y_p) \quad (16)$$

The results presented in figure 13 indicate that subject-specific parcels exhibit stronger intra-parcel correlations than the corresponding population parcels. For this reason, features derived by averaging the BOLD time series inside subject-specific parcels should be more reliable and less corrupted by noise than their counterparts extracted without adapting the parcellation to the individual resting state data.

These results are confirmed by the parcellation entropies presented in figure 14. Contrary to the entropy of population parcellations, the entropy of the subject-specific parcellations is clearly smaller than the entropy of random parcellations for most of the subjects in our database. This result indicates that subject-specific parcellations are the only one that reduces the complexity of the data by combining correlated time series. Several random and population parcellations are even associated with entropies higher than the full brain entropy, which indicates that the averaging inside their parcels followed by the normalization of the time series has produced signals with higher randomness. This result might be an indication that, because these parcellations were not aligned with the brain structure, the averaging has canceled the reliable signal instead of the noise and reduced the SNR. This situation happens rarely with the subject-specific parcellation.

Full brain entropy doesn't vary significantly with age. We observed on the contrary a very significant increase of parcellation entropy for subject-specific and population parcellation, for both hemispheres, while regressing motion MRD and sex. Sex was not significantly associated with the entropies. We observed a large inter-subject variability, potentially due to remaining noise: a large part of the variance of the entropy was linearly explained by none of the variables (R-squared between 0.25 and 0.3). The model used for this analysis and all the p-values obtained are reported in figure 15. These results suggest that original subject time series contain a large amount of noise masking the effect of age. Parcellating the brain allows to reduce the noise and make the trend emerge. However, residual motion artifacts are also revealed in that way.

3.4.4. Parcels Boundary Variations—For the set of parcellations generated in the last section, we measured also the stability of the parcellation by measuring, for each node in the cortical mesh, the fraction of subjects for which the node was assigned to the parcel where it was assigned the most often. This stability measure was, therefore, equal to one when a node was always assigned to the same cluster, and close to zero when the node was randomly assigned to many different clusters. Figure 16 presents the stability measured for the entire brain. These results indicate a higher stability for the visual cortex and a higher variability in frontal, temporal and insular cortices.

3.5. Implementation Details

All the parcellations were obtained in less than five minutes without particular efforts for reducing computational time. With a better choice of computational parameters, this computational time could have been reduced by a factor two or three. The software used for this study can be downloaded on request from the following web page: <https://www.cbica.upenn.edu/sbia/software/index.html>

4. Discussion

In this work, we have compared two approaches for the generation of subject-specific segmentations of the cortex: sGraSP, a variant of the graph-based parcellation method GraSP (Honnorat et al., 2015) and Voronoi tessellations. The results of our experiments can be summarized in three points. First, we show that the fastest methods, based on tessellations, do not produce subject-specific parcels that are highly functionally coherent. This suggests that an optimization approach similar to sGraSP is necessary for generating parcellations able to capture local changes of functional coherence. We demonstrate that sGraSP, on the contrary, produces parcellations that are significantly and strongly more adapted to the individual rs-fMRI scans than the population-level average. The parcels generated by sGraSP vary consistently with respect to signal mixing and the results are perfectly reproducible between the two hemispheres. We show finally that a neurodevelopmental trend is captured with much more strength when subject-specific parcels are derived with sGraSP, under several signal mixing conditions. These results demonstrate that our parcellation strategy is strongly reliable/reproducible and that adapting a template population parcellations to the rs-fMRI scans of a clinical population might improve the reliability of the subsequent analyses.

In this section, we would like to discuss six aspects of our work in details.

4.1. Anchoring Constraint

Many subject-specific parcellation approaches in the literature constrain only the number of subject-specific parcels/networks to be the same across a population. In practice, however, the parcels/networks are never completely free to move and deform because a permutation of their locations would completely disturb the subsequent analysis. A strong regularization of the deformations is always included in the model, most of the time implicitly. For all the methods producing the subject-specific segmentations at the same time as a population consensus, the regularization is explicitly imposed (Varoquaux et al., 2011; Parisot et al., 2015; Arslan et al., 2015; Parisot et al., 2016). For dual regression (Beckmann et al., 2009; Filippini et al., 2009) template based rotation (Schultz et al., 2014) and the direct matching (Wang et al., 2015), the subject-specific segmentations are obtained after a single adaptation of the population model, which ensures that the subjects parcellations are close to the population template. A larger amplitude of deformation is in theory possible with the approach adopted by Wang et al. (Wang et al., 2015). However, the small number of networks considered in this work ensured that the average time series were very stable. This stability prevented the optimization from diverging and deforming too much the population segmentation.

With these comparisons in mind, are we overconstraining our model by imposing the functional center of the subject parcels? We think that this is not the case because in our model the location of the functional centers is not constraining the final shape and location of the parcels that much. Indeed, because adjacent cortical locations tend to be highly correlated, all the nodes close to the optimal functional parcel center will produce similar parcels, very close to the one obtained with the best possible functional center. In the ideal case where the nodes inside a parcel would be perfectly correlated with each other and very weakly correlated with the other parcels nodes, all the nodes within a parcel would be interchangeable: any of them could be selected as functional parcel center. As a result, we think that as soon as moderately large parcels are sought and the subject parcels contain the population functional center, our current approach is able to find nearly optimal parcels.

We would also like to emphasize that fixing the functional centers of the subject parcellation is very similar to fixing the average time series of the parcels in the method proposed by Wang et al. (Wang et al., 2015). As a result, updating the functional centers would lead us to perform an optimization very similar to their approach. In this work, similarly to most other methods (Beckmann et al., 2009; Filippini et al., 2009; Schultz et al., 2014) we are updating the subject parcellations only once, but such an iterative optimization would be of course straightforward to implement in the future.

In addition, a functional registration (Sabuncu et al., 2010; Langs et al., 2010) could be easily introduced for improving the location of the subject functional centers prior to the segmentation, as in (Wang et al., 2015). This step would further increase the ability of our method to deform population parcels for fitting them to subject data.

4.2. Signal Mixing

In this work, we have generated datasets requiring different amplitudes of parcellation deformations by replacing the correlations of the subject time series by a convex combination of the correlations of the subject time series and the correlations between the time series of the population, that were obtained by concatenating all the subject time series. This operation was referred as signal mixing.

A parcellation obtained after signal mixing is still “subject-specific” in the sense that it teases out the specificities of an individual rs-fMRI scan. These subject-specificities are attenuated by signal mixing, but not removed. In the extreme case where the population signal is only slightly modified by the adjunction of individual signal we think that, in analogy with physics, the parcellations obtained can be considered as a set of perturbations of the template parcellation. These perturbations reveal the forces that are acting for stretching the template/population parcellation in order to fit the individual scans.

In this work, signal mixing was used as a tool for generating new experimentation settings and attenuating the deformations. The significance of the variation of parcellation energy with age in our database was not impacted by signal mixing. However, it was shown that this operation, also known as correlation shrinkage, removes a part of the noise of the individual rs-fMRI scans, which can sometimes improve the reliability of the parcellation/feature extraction (Ng et al., 2011, 2012, 2013; Shou et al., 2014). This method can also be used for discarding poor quality fMRI scans (Fritsch et al., 2012).

4.3. Optimal Number of Parcels

In this paper, we have explored a wide range of parcellation refinements, ranging from a few dozen parcels and to several hundred. When too many parcels are chosen, their coverage of the cortex is too small for allowing a reliable estimation of the brain activation within them. On the opposite, large parcels aggregate distinct brain functions and their signal appears as a heterogeneous mix of latent functional processes. A “good” number of parcels for our parcellations should, therefore, be neither too small nor too big. Between these two extremes, however, the choice of a precise number is an open question.

Several authors have proposed strategies for fixing this number (Yeo et al., 2011; Blumensath et al., 2013; Honnorat et al., 2015). A popular approach consists in defining a stability criterion for the parcellation and finding the number of parcels that produces the most stable segmentations of the brain. This approach was for instance implemented in (Yeo et al., 2011), where clustering stability was measured by estimating the reproducibility of the segmentation when the vertices to parcellate are randomly separated into two groups and parcellated independently (Lange et al., 2004). However, all the stability criteria rely on a similarity measure between parcellations and many such measures exist (Pfitzner et al., 2009). It was even demonstrated that the most standard similarity measures for clusterings are not equivalent and that no measure satisfies all the desirable properties that we could expect from them (Meil, 2005). As a result, several parcellation stability measures are possible, and they select different ‘optimum’ numbers of parcels because they are sensitive to different properties of the parcellations.

Moreover, the scale of a useful parcellation will strongly depend on the clinical study. A parcellation containing a single parcel for the cerebellum, for instance, will be of little use for neurologists interested in studying the alteration of the cerebellum across a population of diseased patients. Instead of arbitrary selecting a criterion and therefore indirectly selecting the number of parcels, we prefer delivering an algorithm that can be tuned for parcellating the brain at different scales, so that the users can choose the scale which is the most relevant for their clinical study.

Varying the parameter K penalizing the number of parcels offers also the possibility to produce a spectrum of parcellations. Similarly to Fourier decompositions and scale-space for images, studying the properties of the parcellations as a function of K could potentially provide a lot of insight about the brain in normal and diseased conditions. It becomes in particular possible to study how the fragmentation of a brain area into smaller functional units is impacted by a disease.

4.4. Quality of the Empirical Correlation Matrices

Recent studies have shown that the standard six minutes rs-fMRI scans are slightly too short for a perfect estimation of brain connectivity from empirical correlations between BOLD time series (Birn et al., 2013; Gonzalez-Castillo et al., 2014). Some studies conclude that rs-fMRI scans should last at least twenty-five minutes (Anderson et al., 2011). However, these results are difficult to extrapolate because these minimal scan durations depend on scanners, acquisition protocols and subject motion. They depend also strongly on the type of neuroscience study conducted. When a study concludes for instance that ten minutes long scans are long enough for distinguishing healthy adult brains from each other, such scans might be sufficient for distinguishing healthy adults from diseased patients but too short for distinguishing children twin brains. Moreover, a precise estimation of brain connectivity is not necessary for all the studies and clinical applications. For instance, the largest differences between young children and elder adults brains should be noticed very quickly. In addition, the optimal scan length depends on the scale chosen for describing the connectivity. Even if ten minutes long scans were too short for a precise estimation of a dense high-resolution brain connectivity, they might still provide very good estimations of downsampled connectivity matrices.

The protocol implemented for the present study is based on previous investigations with similar scanners and populations, as explained in (Satterthwaite et al., 2014). We estimated the quality of our empirical correlation matrices by comparing them with the matrices produced by two better covariance estimators, based on linear shrinkage.

Linearly shrinking empirical covariance matrices is a simple and efficient way of improving their conditioning and quality (Ledoit and Wolf, 2003; Chen et al., 2010). This approach was first proposed in the seminal work of Ledoit and Wolf (Ledoit and Wolf, 2003). The estimator proposed in this paper was improved, under the assumption that the signal is Gaussian, by Rao-Blackwellization by Chen et al. (Chen et al., 2010), who described also the Oracle Approximating Shrinkage (OAS), an estimator that performs even better for small sample sizes. For Gaussian distributions, another estimator was proposed soon after by Fisher and Sun (Fisher and Sun, 2011) (FSS). An extension to elliptical distributions was

obtained by plugging the OAS estimator in Tyler's robust covariance estimator (Tyler, 1987; Chen et al., 2011). This approach paved the way for estimators that are robust to outliers and led to numerous improvements such as (Pascal et al., 2014; Ollila and Tyler, 2014; Couillet and McKay, 2014).

Under the assumption that the rs-fMRI data are Gaussian (Chen et al., 2010), the estimators proposed for correlation matrices replace an empirical correlation matrix Σ by an estimation $\bar{\Sigma}$ of the form $\bar{\Sigma} = (1 - \lambda) \Sigma + \lambda I$, where the positive parameter λ is smaller than 1 and depends on Σ , on the duration of the time series and on the assumptions of the method. When the empirical correlation matrix is not reliable, a large parameter λ is computed and the methods take the conservative decision of strongly shrinking the extra diagonal correlations towards zero. In the extreme case, a λ parameter equal to 1 indicates that the identity matrix is the best estimation of the correlation matrix, the extra diagonal non-zero values in Σ are not sufficient for ensuring that the data are actually correlated.

In order to estimate the quality of our empirical correlation matrices used in this work, we computed the shrinkage intensity λ prescribed by the OAS method and by the FSS method, for all the scans in our database. Figure 17 indicates that the median intensity was close to 0.15, which means that a shrinkage would not have reduced the empirical correlations by more than fifteen percent. The agreement between the two shrinkage methods is also very strong. The difference between the shrinkage intensity prescribed by the two approaches was smaller than 1.910^{-5} . These results indicate that the empirical matrices that we have used were of good quality, despite the restricted length of the time series (120 timepoints) compared to the size of the correlation matrices (10.242 times 10.242 voxels).

4.5. Parcels Shape Variations and Noise

A convergent set of observations points the presence of a very significant neurodevelopmental trend in our dataset (Honnorat et al., 2015). This trend, corresponding to a decrease of local functional coherence in the cortex with neurodevelopment between childhood (age of 8) and adulthood (age of 23), is compatible with a hypothesis of increasing segregation of brain function with neurodevelopment, suggested for instance by (Fair et al., 2007, 2009). Because young children tend to move more than adults during the scans, the origin of this trend is still debated. The observed trend could indeed be induced by a motion residual not handled properly by our motion correction algorithm (Satterthwaite et al., 2013) or could result from a side effect of the motion correction.

Whatever the source, however, we know that a trend is present. During our experiments, we confirmed that the parcellations generated by sGraSP were better capturing this trend than the population parcellation. We observed also that, in accordance with the fact that the trend impacts local correlations more than long-range correlations, the effect was better detected with fine scale parcellations, where the advantage of sGraSP was also stronger. We think that these results indicate that the biomarkers extracted from sGraSP parcellation are more reliable than the biomarkers extracted with a fixed population parcellation. Indeed, if sGraSP parcellations were capturing random deformations of the parcels, not related to a global effect in the population, the significance of the trend detection would have been reduced.

The observed increase of significance suggests that the global effect is better captured and that subject-specific parcellations are beneficial.

Noise probably induces a part of the deformations of the population parcellation when it is fit to individual scans. Unfortunately, this part is difficult to estimate because, in the absence of ground truth, estimating the signal-to-noise ratio of the individual rs-fMRI scans is impossible. Many groups have recently adopted a replication strategy for addressing this concern (Wang et al., 2015). However, these studies require large datasets containing several scans per subject, which means that replication cannot be implemented for all the studies. Furthermore, replication cannot avoid a circular argument: the scans of a same subject might still be altered by a large amount of noise, that cannot be canceled out by combining the different scanning sessions of the subject. A significant amount of experiments would be necessary for addressing these concerns. In the meantime, synthetic experiments constitute the most convenient way to check the stability of our algorithms with respect to controlled amounts of noise. This approach was adopted here. By perturbing our population rs-fMRI signal with small amounts of individual signal, we have altered the population signal in a realistic way. In the future, we plan to carry out replication studies as well.

4.6. Model Extensions

The Pearson distances based on temporal correlations between rs-fMRI time series, used in this paper for computing geodesic distances, could be easily replaced by most of the other connectivity measures proposed so far (Smith et al., 2011). In this work, temporal correlations were preferred because Pearson distances are the most standard measure of functional connectivity, they are the most robust (Smith et al., 2011) and they are the easiest to compute, since they do not require Fourier or wavelet transforms like coherence measures (Chang and Glover, 2010) or a regularized matrix inversion, such as partial correlations (Varoquaux et al., 2010b) and partial coherence measures (Fiecas and von Sachs, 2014; Fiecas and Ombao, 2011). Also to date, and despite impressive efforts, the fastest methods require several dozens of minutes per scan for computing whole brain sparse partial correlation matrices (Hsieh et al., 2014, 2013).

However, sGraSP could be easily extended for exploiting these alternative connectivity measures. The parcellation generated in such a way would capture slightly different brain structures. For instance, a parcel derived from Granger causality would describe a region of the brain where the activity is triggered by the parcel center. These centers would appear as functional entry points of their parcels, instead of representing the entire parcel activity as when Pearson distances are used for the parcellation. Similarly, parcels generated from partial correlations or partial coherence would only group brain locations with strong direct connectivity. In fine, the choice of connectivity measure for the parcellation will depend on the subsequent analysis.

5. Conclusion

In this paper, we propose a novel approach for the functional parcellation of individual resting-state fMRI scans. The individual parcellations are derived from parcellations obtained for the entire population, by applying a variant of a recent graph-based parcellation

method. Using a large neurodevelopment study, we show that our method outperforms several faster alternatives, based on Voronoi tessellations. Replications experiments demonstrate that the parcellation generated by our method are stable with respect to rs-fMRI signal perturbations and that the subject- specific parcellations generated by our approach capture the neurodevelopmental trends underlying our data with much more significance. In the future, we will investigate advanced methods for the functional matching of subject and population signals. We will also explore novel strategies for mixing the population and individual resting state scans, in order to control the variability of the subject-specific parcellations.

Acknowledgments

This work was supported by the grant NIH R01AG14971 and by RC2 grants from the National Institute of Mental Health MH089983 and MH089924 and P50MH096891. Support for developing statistical analyses (TDS) was provided by a seed grant from the Center for Biomedical Computing and Image Analysis (CBICA) at Penn. Additional support was provided by K23MH098130 to TDS and the Dowshen Program for Neuroscience. Raquel E. Gur served on Otsuka Advisory Board.

References

- Amunts K, Schleicher A, Bürgel U, Mohlberg H, Uylings H, Zilles K. Broca's region revisited: Cytoarchitecture and intersubject variability. *The Journal of Comparative Neurology*. 1999; 412:319–341. [PubMed: 10441759]
- Anderson JS, Ferguson MA, Lopez - Larson M, Yurgelun - Todd D. Reproducibility of single-subject functional connectivity measurements. *American Journal of Neuroradiology*. 2011; 32(3):548–555. [PubMed: 21273356]
- Arslan S, Parisot S, Rueckert D. Joint spectral decomposition for the parcellation of the human cerebral cortex using resting-state fmri. *Information Processing in Medical Imaging*. 2015:85–97. [PubMed: 26221668]
- Baker J, Holmes A, Masters G, Yeo B, Krienen F, Buckner R, Öngür D. Disruption of cortical association networks in schizophrenia and psychotic bipolar disorder. *JAMA Psychiatry*. 2014; 71:109–118. [PubMed: 24306091]
- Baldassano C, Beck D, Fei-Fei L. Parcellating connectivity in spatial maps. *Peer J*. 2015
- Baumgartner R, Scarth G, Teichtmeister C, Somorjai R, Moser E. Fuzzy clustering of gradient-echo functional mri in the human visual cortex. part i: reproducibility. *Journal of magnetic resonance imaging (JMIR)*. 1997; 7(6):1094–1108. [PubMed: 9400854]
- Beckmann C, Mackay C, Filippini N, Smith S. Group comparison of resting-state fmri data using multi-subject ica and dual regression. *OHBM*. 2009
- Bellec P, Rosa-Neto P, Lyttelton OC, Benali H, Evans AC. Multi-level bootstrap analysis of stable clusters in resting-state fm. *NeuroImage*. 2010; 51:1126–1139. [PubMed: 20226257]
- Birn R, Molloy E, Patriat R, Parker T, Meier T, Kirk G, Nair V, Meyerand M, Prabhakaran V. The effect of scan length on the reliability of resting-state fmri connectivity estimates. *Neuroimage*. 2013; 83:550–558. [PubMed: 23747458]
- Blumensath T, Jbabdi S, Glasser MF, Van Essen DC, Ugurbil K, Behrens TE, Smith SM. Spatially constrained hierarchical parcellation of the brain with resting-state fmri. *NeuroImage*. 2013; 76:313–324. [PubMed: 23523803]
- Boykov Y, Kolmogorov V. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE PAMI*. 2004 Sep; 26(9):1124–1137.
- Boykov Y, Veksler O, Zabih R. Fast approximate energy minimization via graph cuts. *IEEE PAMI*. 2001; 23(11):1222–1239.
- Brodmann, K. Vergleichende Lokalisationslehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. Leipzig: Barth JA.; 1909.

- Bullmore E, Sporns O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*. 2009 Mar;10:186–198. [PubMed: 19190637]
- Bürgel U, Amunts K, Hoemke L, Mohlberg H, Gilsbach J, Zilles K. White matter fiber tracts of the human brain: Three-dimensional mapping at microscopic resolution, topography and intersubject variability. *NeuroImage*. 2006; 29:1092–1105. [PubMed: 16236527]
- Calhoun V, Adali, Pearlson G, Pekar J. A method for making group inferences from functional mri data using independent component analysis. *Human Brain Mapping*. 2001; 14:140–151. [PubMed: 11559959]
- Chang C, Glover G. Time-frequency dynamics of resting-state brain connectivity measured with fmri. *NeuroImage*. 2010; 50:81–98. [PubMed: 20006716]
- Chen Y, Wiesel A, Eldar Y, Hero A. Shrinkage algorithms for mmse covariance estimation. *IEEE Transactions on Signal Processing*. 2010; 58(10):5016–5029.
- Chen Y, Wiesel A, Hero A. Robust shrinkage estimation of high-dimensional covariance matrices. *IEEE Transactions on Signal Processing*. 2011; 59(9):4097–4107.
- Clayton Silver N, Dunlap W. Averaging correlation coefficients: Should fisher's z transformation be used? *Journal of Applied Psychology*. 1987; 72(1):146–148.
- Cordes D, Haughton V, Carew JD, Arfanakis K, Maravilla K. Hierarchical clustering to measure connectivity in fmri resting-state data. *Magnetic Resonance Imaging*. 2002; 20:305–317. [PubMed: 12165349]
- Couillet R, McKay M. Large dimensional analysis and optimization of robust shrinkage covariance matrix estimators. *Journal of Multivariate Analysis*. 2014; 131:99–120.
- Craddock R, James G, Holtzheimer PI, Hu X, Mayberg H. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Human Brain Mapping*. 2012 Aug; 33(8):1914–1928. [PubMed: 21769991]
- Dale A, Fischl B, Sereno M. Cortical surface-based analysis. i. segmentation and surface reconstruction. *Neuroimage*. 1999; 9:179–194. [PubMed: 9931268]
- Delong A, Osokin A, Isack H, Boykov Y. Fast approximate energy minimization with label costs. *International Journal of Computer Vision*. 2012; 96:1–27.
- Economo, C., Koskinas, G. *Die Cytoarchitektonik der Hirnrinde des erwachsenen Menschen*. Berlin: Springer; 1925.
- Fair DA, Cohen AL, Power JD, Dosenbach NUF, Church JA, Miezin FM, Schlaggar BL, Petersen SE. Functional brain networks develop from a “local to distributed” organization. *PLoS Computational Biology*. 2009; 5(5)
- Fair DA, Dosenbach NUF, Church JA, Cohen AL, Brahmbhatt S, Miezin FM, Barch DM, Raichle ME, Petersen SE, Schlaggar BL. Development of distinct control networks through segregation and integration. *Proceedings of the National Academy of Sciences (PNAS)*. 2007; 104(33):13507–13512.
- Fiecas M, Ombao H. The generalized shrinkage estimator for the analysis of functional connectivity of brain signals. *The Annals of Applied Statistics*. 2011; 5:1102–1125.
- Fiecas M, von Sachs R. Data-driven shrinkage of the spectral density matrix of a high-dimensional time series. *Electronic Journal of Statistics*. 2014; 8:2975–3003.
- Filippini N, MacIntosh B, Hough M, Goodwin G, Frisoni G, Smith S, Matthews P, Beckmann C, Mackay C. Distinct patterns of brain activity in young carriers of the apoe-4 allele. *PNAS*. 2009; 106(17):7209–7214. [PubMed: 19357304]
- Finn E, Shen X, Scheinost D, Rosenberg M, Huang J, Chun M, Papademetris X, Constable R. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature Neuroscience*. 2015; 18:1664–1671. [PubMed: 26457551]
- Fisher T, Sun X. Improved stein-type shrinkage estimators for the high-dimensional multivariate normal covariance matrix. *Computational Statistics and Data Analysis*. 2011; 55:1909–1918.
- Fox M, Raichle ME. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*. 2007; 8:700–711. [PubMed: 17704812]
- Fredman M, Tarjan R. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the Association for Computing Machinery*. 1987; 34(3):596–615.

- Fritsch V, Varoquaux G, Thyreau B, Poline J, Thirion B. Detecting outliers in high-dimensional neuroimaging datasets with robust covariance estimator. *Medical Image Analysis*. 2012; 16:1359–1370. [PubMed: 22728304]
- Gonzalez-Castillo J, Handwerker DA, Robinson ME, Hoy CW, Buchanan LC, Saad ZS, Bandettini PA. The spatial structure of resting state connectivity stability on the scale of minutes. *Frontiers in Neuroscience*. 2014; 8(138)
- Greve DN, Fischl B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage*. 2009; 48(1):63–72. [PubMed: 19573611]
- Gulshan, V., Rother, C., Criminisi, A., Blake, A., Zisserman, A. Geodesic star convexity for interactive image segmentation; *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*; 2010. p. 3129–3136.
- Henssen A, Zilles K, Palomero-Gallagher N, Schleicher A, Mohlberg H, Gerboga F, Eickhoff S, Bludau S, Amunts K. Cytoarchitecture and probability maps of the human medial orbitofrontal cortex. *Cortex*. 2016; 75:87–112. [PubMed: 26735709]
- Honnorat N, Eavani H, Satterthwaite T, Gur R, Gur R, Davatzikos C. Grasp: Geodesic graph-based segmentation with shape priors for the functional parcellation of the cortex. *NeuroImage*. 2015; 106106:207–221. 207–211. [PubMed: 25462796]
- Hsieh C-J, Sustik M, Dhillon I, Ravikumar P. Quic: Quadratic approximation for sparse inverse covariance estimation. *Journal of Machine Learning Research*. 2014; 15:2911–2947.
- Hsieh C-J, Sustik MA, Dhillon IS, Ravikumar PK, Poldrack R. Big and quic: Sparse inverse covariance estimation for a million variables. *Advances in Neural Information Processing Systems*. 2013:3165–3173.
- Hubert L, Arabie P. Comparing partitions. *Journal of Classification*. 1985; 2:193–218.
- Jenkinson M, Beckmann CF, Behrens TE, Woolrich MW, Smith SM. FSL. *NeuroImage*. 2012; 62(2): 782–790. [PubMed: 21979382]
- Johnson DB. Efficient algorithms for shortest paths in sparse networks. *Journal of the ACM*. 1977; 24(1):1–13.
- Kim J-H, Lee J-M, Jo H, Kim S, Lee J, Kim S, Seo S, Cox R, Na D, Kim S, Saad Z. Defining functional sma and pre-sma subregions in human mfc using resting state fmri: functional connectivity-based parcellation method. *Neuroimage*. 2010; 49:2375–2386. [PubMed: 19837176]
- Kolmogorov V, Zabih R. What energy functions can be minimized via graph cuts? *IEEE PAMI*. 2004 Feb; 26(2):147–159.
- Lange T, Roth V, Braun M, Buhmann J. Stability-based validation of clustering solutions. *Neural Computation*. 2004; 16:1299–1323. [PubMed: 15130251]
- Langs G, Golland P, Tie Y, Rigolo L, Golby AJ. Functional geometry alignment and localization of brain areas. *Neural Information Processing Systems*. 2010:1225–1233.
- Langs G, Wang D, Golland P, Mueller S, Pan R, Sabuncu M, Sun W, Li K, Liu H. Identifying shared brain networks in individuals by decoupling functional and anatomical variability. *Cerebral Cortex*. 2015 Sep.
- Lashkari D, Vul E, Kanwisher N, Golland P. Discovering structure in the space of fMRI selectivity profiles. *Neuroimage*. 2010; 50(3):1085–1098. [PubMed: 20053382]
- Ledoit O, Wolf M. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*. 2003; 10(5):603–621.
- Leistedt S, Coumans N, Dumont M, Lanquart J, Stam C, Linkowski P. Altered sleep brain functional connectivity in acutely depressed patients. *Human Brain Mapping*. 2009; 30:2207–2219. [PubMed: 18937282]
- Lynall M, Bassett D, Kerwin R, McKenna P, Kitzbichler M, Muller U, Bullmore E. Functional connectivity and brain networks in schizophrenia. *The Journal of Neuroscience*. 2010; 30:9477–9487. [PubMed: 20631176]
- Meil, M. Comparing clusterings: an axiomatic view; *Proceedings of the 22nd international conference on Machine learning*; 2005. p. 577–584.
- Mezer A, Yovel Y, Pasternak O, Gorfine T, Assaf Y. Cluster analysis of resting-state fmri time series. *Neuroimage*. 2009; 45:1117–1125. [PubMed: 19146962]

- Mueller S, Wang D, Fox M, Yeo B, Sepulcre J, Sabuncu M, Shafee R, Lu J, Liu H. Individual variability in functional connectivity architecture of the human brain. *Neuron*. 2013; 77:586–595. [PubMed: 23395382]
- Ng B, Abugharbieh R, Varoquaux G, Poline J, Thirion B. Connectivity-informed fmri activation detection. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2011:285–292. [PubMed: 21995040]
- Ng B, Varoquaux G, Poline J, Thirion B. A novel sparse graphical approach for multimodal brain connectivity inference. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2012
- Ng B, Varoquaux G, Poline J, Thirion B. A novel sparse group gaussian graphical model for functional connectivity estimation. *Information Processing in Medical Imaging (IPMI)*. 2013
- Olhede S, Wolfe P. Network histograms and universality of blockmodel approximation. *Proceedings of the National Academy of Sciences (PNAS)*. 2014; 111(41):14722–14727.
- Ollila E, Tyler T. Regularized M-estimators of scatter matrix. *IEEE Transactions on Signal Processing*. 2014; 62(22):6059–6070.
- Parisot S, Arslan S, Passerat-Palmbach J, Wells III W, Rueckert D. Tractography-driven groupwise multi-scale parcellation of the cortex. *Information Processing in Medical Imaging*. 2015:600–612. [PubMed: 26221706]
- Parisot S, Arslan S, Passerat-Palmbach J, Wells W III, Rueckert D. Group-wise parcellation of the cortex through multi-scale spectral clustering. *NeuroImage*. 2016; 136:68–83. [PubMed: 27192437]
- Pascal F, Chitour Y, Quek Y. Generalized robust shrinkage estimator and its application to stap detection problem. *IEEE Transactions on Signal Processing*. 2014; 62(21):5640–5656.
- Pfutzner D, Leibbrandt R, Powers D. Characterization and evaluation of similarity measures for pairs of clusterings. *Knowledge and Information Systems*. 2009; 19:361–394.
- Power J, Cohen A, Nelson S, Wig G, Barnes K, Church J, Vogel A, Laumann T, Miezin F, Schlaggar B, Petersen S. Functional network organization of the human brain. *Neuron*. 2011; 72:665–678. [PubMed: 22099467]
- Rubinov M, Sporns O. Complex network measures of brain connectivity: Uses and interpretations. *Neuroimage*. 2010; 52(3):1059–1069. [PubMed: 19819337]
- Ryali S, Chen T, Supekar K, Menon V. A parcellation scheme based on von mises-fisher distributions and markov random fields for segmenting brain regions using resting-state fMRI. *NeuroImage*. 2013; 65(0):83–96. [PubMed: 23041530]
- Sabuncu M, Singer B, Conroy B, Bryan R, Ramadge P, Haxby J. Function-based intersubject alignment of human cortical anatomy. *Cereb Cortex*. 2010; 20(1):130–140. [PubMed: 19420007]
- Salvador R, Suckling J, Coleman M, Pickard J, Menon D, Bullmore E. Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral Cortex*. 2005; 15:1332–1342. [PubMed: 15635061]
- Satterthwaite T, Elliott M, Gerraty R, Ruparel K, Loughhead J, Calkins M, Eickhoff S, Hakonarson H, Gur R, Gur R, Wolf D. An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *Neuroimage*. 2013; 64:240–256. [PubMed: 22926292]
- Satterthwaite T, Elliott M, Ruparel K, Loughhead J, Prabhakaran K, Calkins M, Hopson R, Jackson C, Keefe J, Riley M, Mentch F, Sleiman P, Verma R, Davatzikos C, Hakonarson H, Gur R, Gur R. Neuroimaging of the Philadelphia neurodevelopmental cohort. *Neuroimage*. 2014; 86:544–553. [PubMed: 23921101]
- Scheperjans F, Eickhoff S, Hömke L, Mohlberg H, Hermann K, Amunts K, Zilles K. Probabilistic maps, morphometry, and variability of cytoarchitectonic areas in the human superior parietal cortex. *Cerebral Cortex*. 2008; 18:2141–2157. [PubMed: 18245042]
- Schultz A, Chhatwal J, Huijbers W, Hedden T, van Dijk K, McLaren D, Ward A, Wigman S, Sperling R. Template based rotation: a method for functional connectivity analysis with a priori templates. *Neuroimage*. 2014; 102:620–636. [PubMed: 25150630]

- Shen X, Tokoglu F, Papademetris X, Constable R. Groupwise whole-brain parcellation from resting-state fmri data for network node identification. *NeuroImage*. 2013; 82(15):403–415. [PubMed: 23747961]
- Shou H, Eloyan A, Nebel M, Mejia A, Pekar J, Mostofsky S, Caffo B, Lindquist M, Crainiceanu C. Shrinkage prediction of seed-voxel brain connectivity using resting state fmri. *Neuroimage*. 2014; 15:938–944.
- Smith S, Fox P, Miller K, Glahn D, Fox P, Mackay C, Filippini N, Watkins K, Toro R, Laird A, Beckmann C. Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*. 2009; 106:13040–13045. [PubMed: 19620724]
- Smith SM, Miller KL, Moeller S, Xu J, Auerbach EJ, Woolrich MW, Beckmann CF, Jenkinson M, Andersson J, Glasser MF, Van Essen DC, Feinberg DA, Yacoub ES, Ugurbil K. Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*. 2012; 109(8):3131–3136. [PubMed: 22323591]
- Smith SM, Miller KL, Salimi-Khorshidi G, Webster M, Beckmann CF, Nichols TE, Ramsey JD, Woolrich MW. Network modelling methods for fmri. *NeuroImage*. 2011; 54:875–891. [PubMed: 20817103]
- Thirion B, Flandin G, Pinel P, Roche A, Ciuciu P, Poline J. Dealing with the shortcomings of spatial normalization: Multi-subject parcellation of fmri datasets. *Human Brain Mapping*. 2006; 27(8): 678–693. [PubMed: 16281292]
- Tononi G, Sporns O, Edelman G. A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*. 1994; 91:5033–5037. [PubMed: 8197179]
- Tyler D. A distribution-free m-estimator of multivariate scatter. *The Annals of Statistics*. 1987
- van den Heuvel M, Mandl R, Hulshoff Pol H. Normalized cut group clustering of resting-state fmri data. *PLoS ONE*. 2008; 3(4):e2001. [PubMed: 18431486]
- Varoquaux G, Baronnet F, Kleinschmidt A, Fillard P, Thirion B. Detection of brain functional-connectivity difference in post-stroke patients using group-level covariance modeling. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2010a
- Varoquaux G, Gramfort A, Pedregosa F, Michel V, Thirion B. Multi-subject dictionary learning to segment an atlas of brain spontaneous activity. *Information Processing in Medical Imaging (IPMI)*. 2011
- Varoquaux G, Gramfort A, Poline J, Thirion B. Brain covariance selection: Better individual functional connectivity models using population prior. *Advances in Neural Information Processing Systems*. 2010b
- Veksler, O. Star shape prior for graph-cut image segmentation; *IEEE European Conference on Computer Vision (ECCV)*; 2008. p. 454-467.
- Wang D, Buckner R, Fox M, Holt D, Holmes A, Stoecklein S, Langs G, Pan R, Qian T, Li K, Baker J, Stufflebeam S, Wang K, Wang X, Hong B, Liu H. Parcellating cortical functional networks in individuals. *Nature Neuroscience*. 2015; 18:1853–1860. [PubMed: 26551545]
- Yeo B, Krienen F, Sepulcre J, Sabuncu R, Lashkari D, Hollinshead M, Roffman J, Smoller J, Zöllei L, Polimeni J, Fischl B, Liu H, Buckner R. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*. 2011; 106:1125–1165. [PubMed: 21653723]
- Zilles K, Amunts K. Architectonic mapping of the human brain beyond Brodmann. *Neuron*. 2015; 88:1086–1107. [PubMed: 26687219]

Highlights

- we describe sGraSP, a novel subject-specific functional parcellation method
- sGraSP is tested using a large neurodevelopmental cohort (859 scans total)
- additional tests are generated by mixing subject and population signal
- sGraSP outperforms simpler approaches based on Voronoi tessellations
- neurodevelopmental trends are robustly captured by sGraSP for all the signal mixing

Figure 1. Illustration of the geodesic distance maps used for the tessellations

Projection on the pial surface of the left hemisphere of the two discussed geodesic distances, shown for three possible parameter settings. The distance maps were computed from the same parcel functional center located in the intraparietal sulcus (IPS). This functional center was inherited from the population-level parcellation of the left hemisphere obtained for $K = 5$. We present (1) the geodesic distance d , for parameters $c = 0.1, 1.0, 10.0$ and (2) the geodesic quasimetric δ for parameters $c = 0.1, 1.0, 10.0$.

K	1	2	3	4	5	6	7	8	9	10
left hemisphere	2001	842	408	221	152	118	100	87	76	69
right hemisphere	1988	848	411	229	161	114	101	86	76	70

Figure 2.
Number of parcels in the population-level parcellations, for different parameters K .

Figure 3.

(1) Parcellation of the medial surface of the left hemisphere, for $K = 1, 3, 6, 9$ (2001, 408, 118, 76 parcels). (2) Parcellation of the temporal surface of the left hemisphere, for the same parameters.

Figure 4. The tessellations schemes do not produce coherent subject-specific parcels

Energy of the tessellation obtained for the hundred youngest subjects, when the **geodesic distance d is used (in red)** and when the **geodesic quasimetric δ is used (in blue)**, for the following values of the parameter of these distances: 0.1, 0.25, 0.5, 1.0, 2.0, 5.0, 10.0. The energy obtained when the population level parcellation is used for segmenting the subject-specific data is presented in black. This experiment was carried out at the resolution $K=5$, with a mixture of signals corresponding to $\lambda = 0.9$ (1) for the left hemisphere and (2) for the right hemisphere (2800 tessellations total). None of the tessellation improved the baseline energy. As a result, all the paired t-test between the tessellation energies and the baseline energies concluded for a significant positive difference of energy w.r.t. the baseline, instead of an energy decrease (p -values $< 10^{-150}$).

(1)

(2)

Figure 5. The best tessellations are the most similar to the population parcellations

Adjusted Rand Index between the population level parcellation and the tessellations, for all the metric parameters tested in this study and (1) the metric D (2) the metric Δ , separately for the two hemispheres (Left hemisphere in black, right hemisphere in green). These results indicate that the lowest/best tessellation energies in figure 4 correspond to the tessellation that are the most similar to the population parcellations.

Figure 6. Variation of the shape of the parcels produced by the tessellations using the metric d : from overly isotropic to overly tortuous

Tessellation of the left hemisphere of the youngest subject of the database (using the parcel centers inherited from the parcellation corresponding to $K = 5$ that contains 152 parcels), based on the geodesic distance d , for distance parameters 0.1, 0.5, 2.0, 10.0. The tessellations are shown (1) on the lateral and (2) on the medial inflated surface. For small parameter values, the parcels are overly isotropic. For the largest parameters, some parcels are so tortuous that they appear to be fragmented at first glance (see e.g. pink and blue parcels).

Figure 7. Variation of the shape of the parcels produced by the tessellations using the metric δ : moderate decrease of the isotropy

Tessellation of the left hemisphere of the youngest subject of the database (using the parcel centers inherited from the parcellation corresponding to $K = 5$), based on the geodesic quasimetric δ , for distance parameters 0.1, 0.5, 2.0, 10.0. The tessellations are shown (1) on the lateral and (2) on the medial inflated surface. For small parameters, the parcels are extremely isotropic. When the distance parameter is increased, a variety of shapes similar to the one produced by GraSP emerges.

Figure 8.
Subject-specific parcellations of the left hemisphere projected on (1) the lateral fsaverage5 inflated surface (2) the medial fsaverage5 inflated surface.

(1)

(2)

Figure 9. Large and significant fit to the individual data observed for our subject-specific parcellation

For the hundred youngest subjects, improvement/reduction of the parcellation energy when using the subject-specific parcellation with respect to the population-level parcellation (in percent of the second one), at scale $K = 5$, for an increasing mixture of individual and population rs-fMRI signal $\lambda = 0, 0.05, \dots, 0.95$ and for (1) the left hemisphere and (2) the right hemisphere (4000 parcellations total). The parcellation produced by our algorithm fit the subject data much better than the population-level parcellation (up to 20% energy improvement/decrease). All the one sample t test performed concluded that the

improvement/reduction of the energy is very significant ($p\text{-values} < 10^{-90}$). The amplitude of the fit decreases almost linearly with λ (coefficient of determination $R\text{-squared} \approx 0.96$ for both hemispheres). These observations demonstrate the stability/reliability of our method for the entire scale of λ . Please note that, contrary to figure 4, this figure presents a decrease of parcellation/tessellation cost, depending on the signal mixing.

(1)

(2)

Figure 10. Significant deformation of the parcels during their adaptation to individual data, for most of the signal mixtures

For the hundred youngest subjects, adjusted Rand index measuring the similarity between the subject-specific parcellations and the population-level parcellation, at scale $K = 5$, for an increasing mixture of individual and population rs-fMRI signal $\lambda = 0, 0.05, \dots, 0.95$ and for (1) the left hemisphere and (2) the right hemisphere (4000 parcellations total). The adjusted Rand indices observed are very small, which indicates that the shape of the parcels is modified a lot by our parcellation scheme. These results suggest also that inter-subject differences are probably large. The large adjusted Rand indices observed when λ tends to one indicate that this variability is mainly data driven: when the population signal are

slightly perturbed, by introducing 5% of subject signal, the parcellation produced by our method are always very close to the population level (adjusted Rand indices ≈ 0.85).

(1)

(2)

Figure 12. The increase of parcellation energy during neurodevelopment is more difficult to detect without adapting the population parcellation to the subject data

When the population-level parcellation is used for parcellating the subjects, (1) p-values corresponding to the variation of parcellation energy with age, for six different signal mixtures $\lambda = 0$ (no data shrinkage) and $\lambda = 0.05, 0.1, 0.15, 0.2, 0.25$, for both hemispheres ('l' and 'r') and for increasing parcellation scales: $K = 1 \dots 10$ (12840 parcellations total). (2) associated regression coefficient. Although significant, the impact of neurodevelopment is less obvious than using subject-specific parcellations, in particular at small scales ($K = 3$). The coefficients, whereas still positive, are slightly smaller than previously. These results indicate that neurodevelopmental trends are better extracted by subject-specific parcellations.

Figure 13. Stronger intra-parcel average correlations with subject-specific parcellations

For each parcel, median of the intra-parcel average correlation over the entire dataset (859 subjects) for (1) the left hemisphere and (2) the right hemisphere separately. For each parcel, median absolute deviation of the intra-parcel average correlation over the entire dataset (859 subjects) for (3) left and (4) the right hemisphere. These results indicate that the intra-parcel average correlation is stronger for the subject-specific parcellations and more variable, which suggest that the subject-specific parcellations are better capturing individual brain function.

Figure 14. Subject-specific parcellation entropies better exploit redundant information

Full-brain, population and subject-specific parcellation entropies for the entire dataset (859 subjects) as a function of age for (1) the left hemisphere and (2) the right hemisphere separately. Comparison with the random parcellation entropies for (3) left and (4) right hemispheres. The lower subject-specific parcellation entropies indicate that signal averaging inside subject-specific parcels better combines redundant signals for reducing noise. Population parcellations are only slightly better adapted to the signal than random parcellations : their entropy is lower than random parcellations entropies nearly half of the time (49.3% of the time for left hemisphere, 46.5% for right hemisphere), comparable 17.4% and 19.6% of the time and higher for 33.3% and 33.9% of the subjects. By comparison, subject-specific parcellations entropies are lower than random parcellations 95.3% and 96.1% of the time, comparable 1.3% and 1.2% of the time and worse/higher 3.4% and 2.7% of the time.

Model: $entropy \approx \alpha + \beta \times age + \gamma \times MRD + \delta \times sex$

parcellation	hemisphere	α pvalue	β pvalue	γ pvalue	δ pvalue	R-squared
none	left	0	0.33	0.14	0.95	0.004
avg. random	left	0	0.06	1.4×10^{-69}	9.8×10^{-6}	0.338
population	left	0	7.6×10^{-21}	1.1×10^{-39}	0.18	0.291
subject	left	0	3.6×10^{-20}	1.5×10^{-41}	0.28	0.295
none	right	0	0.98	2.9×10^{-4}	0.18	0.018
avg. random	right	0	0.3	3.4×10^{-64}	3.8×10^{-6}	0.315
population	right	0	1.6×10^{-18}	2.5×10^{-35}	0.30	0.264
subject	right	0	1.3×10^{-17}	1.9×10^{-35}	0.61	0.259

Figure 15. A linear trend of functional entropy with respect to age emerges when parcellating the brain

This table presents the p values associated with the regression of full brain, population, subject or random parcellations entropies from age, motion, and sex, for the two hemispheres separately. For random parcellations, we used the average of the twenty random parcellations entropies computed for each subject. These results indicate that an age trend is only present for population and subject-specific parcellations. For both, motion residuals are explanatory variables, but the age trend is very significant and sex differences have no impact. Average random parcellation entropies are not related to age but appear to reflect motion artifacts and sex differences very significantly. Very low R-squared indicate that full brain entropy regressions failed.

Figure 16. Reduced parcellation stability in the temporal, frontal and insular cortices

Parcellation stability reveals the variability of parcels boundaries. The parcels located in the inferior frontal cortex, medial prefrontal cortex, temporal lobe and insular cortex present slightly higher variability. The stability of the visual cortex is, on the contrary, higher than the average. The stability was windowed between 0.1 and 0.5 for improving the visualization.

Figure 17.
Shrinkage intensity prescribed by the OAS estimator for all the scans in our database and both hemispheres.