Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations

Evan M. Gordon[1], Timothy O. Laumann[1], Babatunde Adeyemo[1], Jeremy F. Huckins[5], William M. Kelley[5]
and Steven E. Petersen[1,2,3,4]

[1]Department of Neurology, [2]Department of Psychology, [3]Department of Radiology, [4]Department of Anatomy and Neurobiology,
Washington University School of Medicine, St. Louis, MO, USA and [5]Department of Psychological and Brain Sciences, Dartmouth
College, Hanover, NH, USA

Address correspondence to Evan M. Gordon. Email: egordon@npg.wustl.edu
E.M.G and T.O.L contributed equally to this work

The cortical surface is organized into a large number of cortical areas; however, these areas have not been comprehensively mapped in the human. Abrupt transitions in resting-state functional connectivity (RSFC) patterns can noninvasively identify locations of putative borders between cortical areas (RSFC-boundary mapping; Cohen et al. 2008). Here we describe a technique for using RSFC-boundary maps to define parcels that represent putative cortical areas. These parcels had highly homogenous RSFC patterns, indicating that they contained one unique RSFC signal; furthermore, the parcels were much more homogenous than a null model matched for parcel size when tested in two separate datasets. Several alternative parcellation schemes were tested this way, and no other parcellation was as homogenous as or had as large a difference compared with its null model. The boundary map-derived parcellation contained parcels that overlapped with architectonic mapping of areas 17, 2, 3, and 4. These parcels had a network structure similar to the known network structure of the brain, and their connectivity patterns were reliable across individual subjects. These observations suggest that RSFC-boundary map-derived parcels provide information about the location and extent of human cortical areas. A parcellation generated using this method is available at http://www.nil.wustl.edu/labs/petersen/Resources.html.

Keywords: cortical areas, functional connectivity, parcellation, resting state

Introduction

The cortical surface of the brain is organized into a large number of interacting cortical areas (Sejnowski and Churchland 1989). Accurate identification of these cortical areas is a major goal of modern systems neuroscience, as it would provide substantial benefits to many areas of neuroscientific investigation. For example, identification and functional characterization of visual areas in the macaque has provided a detailed hierarchical wiring diagram of the primate visual system that has greatly aided our understanding of visual processing (Felleman and Van Essen 1991). Identifying human cortical areas would be a critical first step toward the same sort of comprehensive characterization of information flow within the brain's various processing systems. Additionally, identification of cortical areas would greatly improve investigations of brain function using graph theory (Bullmore and Sporns 2009), because such areas could serve as rationally defined, neurobiologically based network "nodes" (Wig et al. 2011; Power et al. 2013). Finally, identified areas can serve as a priori regions of interest (ROIs) for the analysis of functional neuroimaging data. Averaging data within predefined areas would improve signal-to-noise (SNR) and reduce multiple comparison problems in statistical testing.

Identification of distinct cortical areas is based on observing dissociations in one or more critical underlying brain properties, including functional responses, topography, architectonics, and connectivity (Felleman and Van Essen 1991; Carmichael and Price 1994, 1996). In the macaque, decades of research using these modalities have provided a reasonable first-order approximation of a complete cortical areal parcellation (Lewis and Van Essen 2000; Paxinos et al. 2000; Saleem et al. 2007; Van Essen et al. 2012; Markov et al. 2014). Although a limited number of similar areal dissociations have been identified in humans (e.g., Brodmann 1909; Öngür et al. 2003; Schleicher et al. 2005), the measurement of these brain properties often relies either on invasive neural recordings or on postmortem examinations of brain tissue, both of which are difficult to obtain for large expanses of cortex in humans. As such, definitions of cortical areas in humans have lagged behind those in other primates.

Advances in functional neuroimaging techniques offer the potential for noninvasive in vivo recording of brain activity. In principle, cortical areas may be dissociated by their differential responses to specific task conditions (Petersen et al. 1988). However, application of this approach to the cortex broadly has been challenging, as most tasks recruit large networks of coactivated areas. This lack of specificity makes it difficult to identify fine dissociations between adjacent and functionally related areas using a necessarily limited task set.

Recently, a functional magnetic resonance imaging (fMRI) technique called resting-state functional connectivity (RSFC) has emerged that may provide one modality for noninvasive parcellation of human cortex. RSFC relies on the observation that, in the absence of any task, spatially distant regions of cortex exhibit highly correlated patterns of blood oxygenation level-dependent (BOLD) activity (Biswal et al. 1995) that are both spatially structured (Beckmann et al. 2005; Power et al. 2011; Yeo et al. 2011) and relatively reliable across individuals (Damoiseaux et al. 2006; Shehzad et al. 2009). While the precise significance of RSFC is uncertain, accumulating evidence suggests that regions exhibiting RSFC correlations are also functionally coactive during tasks (Fox and Raichle 2007; Smith et al. 2009; Biswal et al. 2010). In this view, these correlations observed during the resting state at least partly reflect the statistical history of regional coactivation (Dosenbach et al. 2007). RSFC correlations also appear to be at least partly constrained by structural connections, though regions with no direct structural connections can also be functionally connected, likely via indirect pathways (Vincent et al. 2007; Honey et al. 2009). Taken together, this evidence suggests that RSFC measurements reflect some combination of both a region's function, in a manner not limited to any one task, and its direct and indirect connectivity.

RSFC data may be used to perform areal parcellation via a recently proposed approach known as boundary mapping (Cohen et al. 2008; Wig, Laumann, Petersen 2014). The boundary mapping approach relies on the observation that RSFC patterns can abruptly change from one cortical location to a proximate location, mirroring the abrupt changes in function or connections that form the basis of cortical area discrimination in nonhuman primates (Felleman and Van Essen 1991); these locations of abrupt change may thus represent boundaries between cortical areas. The boundary mapping technique has previously been used to identify transition zones in limited sections of cortex, including left lateral parietal cortex (Nelson, Cohen, et al. 2010; Barnes et al. 2012) and parts of frontal cortex (Cohen et al. 2008; Nelson, Dosenbach, et al. 2010; Hirose et al. 2012, 2013), as well as in the whole brain (Wig, Laumann, Cohen, et al. 2014; Wig, Laumann, Petersen 2014). Boundaries identified in this way have been shown to (1) Separate regions with functionally discrete task activation time courses (Nelson, Cohen, et al. 2010); (2) match functional activation patterns; (3) correspond well with system-level divisions, but also further subdivide those systems; and (4) match architectonically defined areal borders between V1 and V2 (Wig, Laumann, Petersen 2014). In summary, boundaries identified using this technique are reasonable candidates for borders between cortical areas. However, no previous work has either used these boundaries to identify cortical areas, or evaluated the resulting cortical areas. Here we present a method for identifying and evaluating putative cortical areas from group-average RSFC-boundary maps.

A parcellation that accurately represents cortical areas of the brain should have, among others, several properties. First, each parcel should generally be homogenous, in that it should have a similar functional connectivity pattern at all points within the parcel (Craddock et al. 2012; Shen et al. 2013). Second, a parcellation that accurately represents cortical areas should contain parcels that overlap known human cortical areas that have been well described with cytoarchitectonics (e.g., Fischl et al. 2008). Third, a parcellation that accurately represents cortical areas should have a large scale network structure that is consistent with the known network structure of the brain (Power et al. 2011; Wig et al. 2011; Yeo et al. 2011). Finally, parcels that accurately represent cortical areas in group-average data should serve as reasonable a priori ROIs in individual subjects. While the known interindividual variability in areal extent (e.g., Amunts et al. 2000) means that cortical area locations in individual subjects are unlikely to precisely match parcels identified from group-average data, these group-average parcels should still represent the central tendency of the group. Thus, for any given parcel, the functional connectivity patterns across subjects should reflect that level of reliability.

We note that some of these criteria—particularly parcel homogeneity and overlap with architectonics—are likely to fail for a minority of cortical areas. For example, some cortical areas are topographically organized (e.g., somatotopy in somatomotor cortex), such that subregions within the area have different functional responses (Rao et al. 1995), including different RSFC responses (Long et al. 2014). These functional dissociations would likely either reduce the observed RSFC homogeneity of a parcel representing the area, or result in the delineation of sub-areal parcels within a single cortical area. These are unavoidable limitations of any RSFC-based technique.

In this study, we constructed a set of parcels derived from a group-average RSFC-boundary map that represent putative cortical areas. We assessed the homogeneity of these parcels, and we compared those homogeneities against an appropriate null model. We additionally assessed the homogeneity of these boundary map-derived parcels using an independent dataset, collected on a different scanner model at a different institution. Furthermore, we compared the homogeneity of the boundary map-derived parcellation with the homogeneities of several other alternative parcellations (Brodmann 1909; Tzourio-Mazoyer et al. 2002), including other candidate approaches for performing whole-brain areal partitioning using RSFC data (Power et al. 2011; Yeo et al. 2011; Craddock et al. 2012; Shen et al. 2013). Each of these sets of parcel homogeneities was also compared with a tailored null model, all within the independent dataset. We also identified boundary map-derived parcels that overlapped with several known human architectonic areas. We further identified the network structure of the boundary map-derived parcellation and compared this structure with the network structure identified using all gray matter points in the brain. Finally, we assessed the level of intersubject reliability of subject-level RSFC patterns from these boundary map-derived parcels.

Methods

For a graphical summary of the methods, see Figure 1.

We acquired two independent datasets: Dataset 1, which we used to create an RSFC-boundary map and generate parcels; and Dataset 2, which we used to compare the boundary map-derived parcellation against other putative areal parcellations.

Dataset 1

Subjects

Data were collected from 120 healthy young adult subjects during relaxed eyes–open fixation (60 females, mean age = 25 years, age range = 19–32 years). All subjects were native speakers of English and right-handed. Subjects were recruited from the Washington University community and were screened with a self-report questionnaire to ensure that they had no current or previous history of neurological or psychiatric diagnosis. Informed consent was obtained from all subjects. The study was approved by the Washington University School of Medicine Human Studies Committee and Institutional Review Board.

Data Acquisition

Structural and functional MRI data were obtained with a Siemens MAGNETOM Trio Tim 3.0-T Scanner (Erlangen, Germany) and a Siemens 12-channel Head Matrix Coil. A T_1-weighted sagittal magnetization-prepared rapid acquisition gradient-echo (MP-RAGE) structural image was obtained [time echo (TE) = 3.08 ms, time repetition, TR(partition) = 2.4 s, time to inversion (TI) = 1000 ms, flip angle = 8°, 176 slices with 1 × 1 × 1 mm voxels; Mugler and Brookeman 1990]. An auto align pulse sequence protocol provided in the Siemens software was used to align the acquisition slices of the functional scans parallel to the anterior commissure–posterior commissure plane of the MP-RAGE and centered on the brain. This plane is parallel to the slices in the Talairach atlas (Talairach and Tournoux 1988).

During functional MRI data acquisition, subjects were instructed to relax while fixating on a black crosshair that was presented against a white background. Functional imaging was performed using a BOLD contrast-sensitive gradient-echo echo-planar imaging (EPI) sequence (TE = 27 ms, flip angle = 90°, in-plane resolution = 4 × 4 mm). Whole-brain EPI volumes (MR frames) of 32 contiguous, 4-mm-thick axial slices were obtained every 2.5 s. A T_2-weighted turbo spin-echo
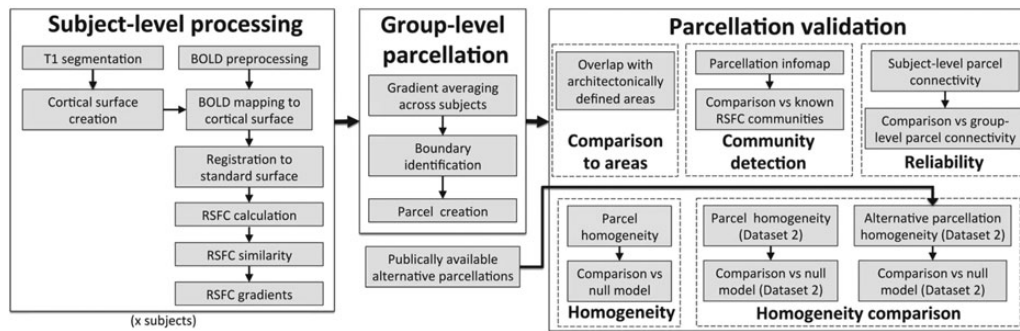
Figure 1. Visual outline of analysis methods.

structural image (TE = 84 ms, TR = 6.8 s, 32 slices with 1 × 1 × 4 mm voxels) in the same anatomical planes as the BOLD images was also obtained to improve alignment to an atlas. The number of volumes collected from subjects ranged from 184 to 729 (mean = 336 frames, 14.0 min).

Dataset 2

Subjects

Data were collected from 108 healthy young adult subjects during relaxed eyes–open fixation (69 females, mean age = 21 years, age range = 18–33 years). Subjects were recruited from the Dartmouth College community and were screened with a self-report questionnaire to ensure that they had no neurological problems, were not using psychoactive medications, and had normal or corrected-to-normal vision. Participants were given course credit or monetary compensation in exchange for their participation and were provided informed consent in accordance with the guidelines set by the Committee for the Protection of Human Subjects at Dartmouth College. These subjects were selected as the subjects with minimal in-scan head motion from a larger cohort of 746 subjects.

Data Acquisition

Structural and functional MRI data were obtained with a Philips Achieva 3.0-T scanner and a 32-channel phased array coil. A T_1-weighted sagittal MP-RAGE structural image was obtained (TE = 4.6 ms, TR = 9.9 ms, flip angle = 8°, 160 slices with 1 × 1 × 1 mm voxels).

During functional MRI data acquisition, subjects were instructed to relax while fixating on a white crosshair that was presented against a black background. Functional imaging was performed using a BOLD contrast-sensitive gradient-echo echo-planar sequence (TE = 35 ms, flip angle = 90°, in-plane resolution = 3 × 3 mm, sense factor = 2). Whole-brain EPI volumes (MR frames) of 36 3.5-mm-thick axial slices were obtained every 2.5 s with 0.5 mm skip between slices. Two 5 min runs (240 volumes total) were collected from each subject.

Further analysis of both datasets was identical, except where noted.

Preprocessing

Functional images were first processed to reduce artifacts (Miezin et al. 2000). These steps included (1) Correction of odd versus even slice intensity differences attributable to interleaved acquisition without gaps, (2) correction for head movement within and across runs, and (3) across-run intensity normalization to a whole-brain mode value of 1000. Atlas transformation of the functional data was computed for each individual using the MP-RAGE scan. Each run was then resampled to an isotropic 3-mm atlas space (Talairach and Tournoux 1988), combining movement correction and atlas transformation in a single cubic spline interpolation (Lancaster et al. 1995; Snyder 1996). All subsequent operations were performed on the atlas-transformed volumetric time series.

Functional Connectivity Processing

Additional preprocessing steps to reduce spurious variance unlikely to reflect neuronal activity were executed as recommended in Power et al. (2014). RSFC preprocessing was performed in two iterations. In the

first iteration, the processing steps were (1) demeaning and detrending, (2), multiple regression including: whole-brain, ventricular and white matter signals, and motion regressors derived by Volterra expansion (Friston et al. 1996), and (3) a band-pass filter (0.009 Hz < f < 0.08 Hz).

Following the initial RSFC preprocessing iteration, temporal masks were created to flag motion-contaminated frames. Motion-contaminated volumes were identified by frame-by-frame displacement (FD, described in Power et al. 2012). Volumes with FD > 0.2 mm (Dataset 1)/FD > 0.25 mm (Dataset 2; different thresholds were used based on observations of different motion "noise floors" in the two datasets, following Power et al. 2012), as well as uncensored segments of data lasting fewer than 5 contiguous volumes, were flagged for removal. In Dataset 1, these masks censored 16 ± 14% (range: 0.7–66%) of the data across subjects; on average, subjects retained 279 ± 107 volumes (range: 151–719). In Dataset 2, these masks censored 8 ± 2% (range: 4–12%) of the data across subjects; on average, subjects retained 221 ± 5 volumes (range: 212–230).

The data were then reprocessed in a second iteration, incorporating the temporal masks described above. This reprocessing was identical to the initial processing stream, but ignored censored data. Finally, the data were interpolated across censored frames using least squares spectral estimation (Power et al. 2014) of the values at censored frames, so that continuous data can be passed through (4) a band-pass filter (0.009 Hz < f < 0.08 Hz) without contaminating frames near high motion frames (Power et al. 2012; Carp 2013). It should be noted that, even following this processing, censored frames are still ignored during the final correlation calculations between time courses.

Surface Processing and CIFTI Creation

Surface generation and sampling of functional data to anatomical surfaces followed a similar procedure as described in Glasser et al. (2013). First, following volumetric registration, anatomical surfaces were generated from each subject's MP-RAGE image using FreeSurfer's default recon-all processing pipeline (version 5.0). This pipeline included brain extraction, segmentation, generation of white matter and pial surfaces, inflation of the surfaces to a sphere, and surface shape-based spherical registration of the subject's "native" surface to the fsaverage surface (Dale and Sereno 1993; Dale et al. 1999; Fischl et al. 1999; Ségonne et al. 2004, 2005). The fsaverage-registered left and right hemisphere surfaces were then brought into register with each other (Van Essen et al. 2012) and resampled to a resolution of 164 000 vertices using Caret tools (Van Essen et al. 2001). Finally, each subject's surface was downsampled to a 32 492 vertex surface (fs_LR 32k), which allowed for analysis in a computationally tractable space while still oversampling the underlying resolution of BOLD data used in subsequent analyses. The above procedure results in a surface space that allows for quantitative analysis across subjects. A script for this procedure is available on the Van Essen Lab website (Freesurfer_to_fs_LR Pipeline, http://brainvis.wustl.edu/wiki/index.php/Caret:Operations/Freesurfer_to_fs_LR).

Surface processing of the BOLD data proceeded through the following steps. First, the BOLD volumes are sampled to each subject's individual "native" midthickness surface (generated as the average of the

white and pial surfaces) using the ribbon-constrained sampling procedure available in Connectome Workbench 0.84 that samples data from voxels within the gray matter ribbon (i.e., between the white and pial surfaces; Glasser and Van Essen 2011). Voxels with a time series coefficient of variation 0.5 SDs higher than the mean coefficient of variation of nearby voxels (within a 5-mm sigma Gaussian neighborhood) are excluded from the volume to surface sampling, as described in Glasser et al. (2013). Once sampled to the "native" surface, time courses were deformed and resampled from the individual's "native" surface to the 32k fs_LR surface. Finally, the time courses were smoothed along the 32k fs_LR surface using a Gaussian smoothing kernel ($\sigma = 2.55$).

These surfaces are then combined with volumetric subcortical and cerebellar data into the Connectivity Informatics Technology Initiative (CIFTI) format using Connectome Workbench (Glasser et al. 2013), creating full brain time courses that exclude nongray matter tissue. Subcortical (including accumbens, amygdala, caudate, hippocampus, pallidum, putamen, and thalamus) and cerebellar voxels were selected based on a mask generated by finding the modal assignment of voxels by Freesurfer segmentation across all subjects. Volumetric data were smoothed within this mask with a Gaussian kernel ($\sigma = 2.55$) before being combined with the surface data.

Boundary Map Generation

RSFC-boundary mapping identifies transitions in resting-state correlations across the cortical surface. The original approach described in Cohen et al. (2008) applied 2D image processing tools to BOLD data sampled from patches on a flattened cortical surface. The current implementation performs all calculations directly on a closed surface topology and applies to the entire cortical surface. The RSFC-Boundary Mapping procedure is implemented using Connectome Workbench and Matlab (Version 7.14, Mathworks, Inc., Sherborn, MA, USA) and follows a similar sequence as described in Wig, Laumann, Petersen (2014) with some notable distinctions that will be highlighted below.

For each subject, the time course of each surface vertex was correlated with that from every other surface vertex and subcortical voxel in the CIFTI space. Each correlation map was transformed using Fisher's r-to-z transformation. For each hemisphere, the subject's RSFC map similarity matrix was created by calculating the pairwise spatial correlations between all vertex's RSFC correlation maps, producing a $32k \times 32k$ matrix. To find positions where RSFC similarity exhibited abrupt changes, the first spatial derivative was computed using the "cifti-gradient" function in Connectome Workbench. This resulted in 32k gradient maps for each hemisphere. These gradient maps were then averaged across subjects. At this point, instead of using nonmaxima suppression to identify boundaries in the gradient maps, as in Wig, Laumann, Petersen (2014), we used the "watershed by flooding" algorithm (Beucher and Lantuejoul 1979), implemented using custom Matlab scripts. This standard image segmentation procedure defines regions in the gradient maps by starting from local minima (vertices with values smaller than of their neighbors that were <3 vertices away) and iteratively growing until reaching locations that could ambiguously be assigned to more than one region. These boundary locations identify putative boundaries in the gradient maps. Finally, the 32k boundary maps from each hemisphere were averaged to indicate the frequency with which a given vertex was identified as a boundary.

Boundary Map Reliability

To determine the reliability of the boundary maps, we calculated the degree of spatial correlation between the boundary maps from the two datasets as an overall measure of reliability. To further determine whether the strongest boundaries in particular were highly reliable, we then thresholded the two boundary maps to retain the top quartile of boundary values (i.e., retaining the cortical vertices most likely to be boundaries) and assessed the overlap of the two thresholded boundaries by calculating Dice's coefficient.

Parcel Creation

Parcels were created from the Dataset 1 boundary map only using custom Matlab scripts. We identified all local minima on the boundary map image as seeds to be used for parcel creation. Parcels were grown from these seeds using the "watershed by flooding" procedure described above, such that parcels were allowed to expand outward from the seed until they either reached a height threshold on the boundary map or met another parcel. This resulted in a large number of parcels tiling the cortical surface (>1000), with one-vertex wide borders (i.e., the watershed zones) separating them. Pairs of parcels were then merged together based on the values of the boundary map in the border vertices between the parcels, which represent the local change in connectivity patterns, and therefore can be considered a measure of the dissimilarity of the parcels. If the median boundary value between two parcels was below a threshold, then the parcels were considered not sufficiently dissimilar and were merged together. We visually examined multiple border thresholds, and the optimal threshold that captured all major divisions in the boundary map image appeared to be at the 60th percentile of the values in the boundary map (see Supplementary Fig. 3 for parcellations resulting from other threshold values). As areas of the cortex with very high boundary map values are likely to be transition zones between parcels rather than parcels themselves, we then eliminated all parcels and portions of parcels in vertices with high boundary map values (defined as the top quartile of values in the boundary map).

This procedure produced an anatomically plausible number of parcels that visually appeared to fit the contours of the boundary map. Parcels in low-SNR areas (defined as regions with mean BOLD signal <750, consisting primarily of orbitofrontal cortex and anterior ventral lateral temporal lobe; see Ojemann et al. 1997; Wig, Laumann, Petersen 2014), which are likely to be noisy and unreliable, were excluded from further analysis. Finally, we eliminated parcels containing fewer than 15 cortical vertices (~30 mm^2) because the effective resolution of the BOLD data (originally $4 \times 4 \times 4$ mm, then upsampled and smoothed on the surface) suggested that accurate identification and evaluation of objects of that small might be dubious.

Parcel Evaluation

The parcel creation procedure outlined above creates parcels based on strong boundaries that indicate large differences in connectivity patterns between adjacent cortical regions. However, a parcel that accurately represents a cortical area should not only be distinct from its neighbors but, in most cases (i.e., nontopographic regions), it should also have a single, consistent connectivity pattern across the parcel—in other words, its connectivity pattern should be homogenous within the parcel. Thus, the degree to which the created parcels are homogenous can serve as a quality metric of the parcellation (Craddock et al. 2012; Shen et al. 2013). We assessed the homogeneity of the parcellation using the following technique: for each parcel, we computed the average whole-brain connectivity pattern of each vertex in the parcel across subjects in Dataset 1. We then entered the connectivity patterns from all vertices in a parcel into a principal components analysis. The homogeneity of the parcel was calculated as the percent of total variance across all vertices' connectivity patterns that can be explained by the first (largest) principal component. A higher homogeneity value indicates that the connectivity patterns of vertices within the parcel can be better described by a single connectivity pattern. We then averaged the homogeneity values across parcels to determine the overall homogeneity of the whole parcellation. Compared with other metrics of parcel homogeneity, this novel metric has the advantage of being highly interpretable: The homogeneity of a parcel represents the percent of variance in the parcel explained by the most common connectivity pattern. Homogeneity analyses conducted with a previously devised homogeneity metric (average z-transformed pairwise correlations between all vertex connectivity patterns within a parcel, from Craddock et al. 2012) yielded very similar results (Supplementary Fig. 11).

However, we note that any metric of parcel homogeneity is likely to be dependent on parcel size, with smaller parcels being intrinsically more homogenous. To illustrate this fact, consider that a large, perfectly homogenous parcel could be divided in half, and both halves would still be perfectly homogenous. Furthermore, a direct comparison of the homogeneities of the large and small parcels would not indicate one scheme as superior to the other, even though the large perfectly

homogenous parcel is much more likely to represent a single cortical area. Even in a purely random parcellation scheme, randomly placed small parcels are more likely to contain a single connectivity pattern than randomly placed large parcels, which will more often span multiple cortical areas. Thus, any homogeneity-based evaluation of a parcellation must be compared with a null model—it should consider not only how homogenous the parcels are, but also whether they are more homogenous than would be expected from randomly placed parcels of the same size and shape. Thus, we assessed the degree to which a parcellation was more homogenous than a null model consisting of many parcellations with randomly placed parcels of the same size, shape, and relative position to each other.

To create such random parcellations, we rotated each hemisphere of the original parcellation a random amount around each of the *x, y,* and *z* axes on the spherical expansion of the 32k fs_LR cortical surface. This procedure randomly relocated each parcel while maintaining the relative positions of parcels to each other. Each parcel was then slightly dilated or contracted to adjust for vertices gained or lost due to the non-uniform vertex density across the surface of the sphere, thus maintaining the same number of vertices within the rotated parcel while approximately maintaining the same shape. Random rotation was repeated 1000 times to generate distributions of average homogeneities calculated from randomly placed versions of each tested parcellation. Notably, in any random rotation, some parcels will inevitably be rotated into the medial wall (where no data exist) or into low-SNR regions (where we believe the homogeneity of data to be particularly low). The homogeneity of a parcel rotated into one of these regions was not calculated; instead, we assigned this parcel the average homogeneity of all random versions of the parcel that were rotated into valid (high-SNR) cortical regions.

The average homogeneity of the original parcellation was compared with the homogeneities of the set of rotated parcellations. We assessed (1) the number of rotated parcellations that had worse average homogeneity than the original parcellation and (2) the difference between the original parcellation homogeneity and the distribution of random homogeneities, calculated as a *Z*-score [(original homogeneity − mean of random homogeneities)/standard deviation of random homogeneities].

Comparison of Parcel Homogeneity Against Alternative Parcellations

We compared the homogeneities of boundary-derived parcels against those of several alternative parcellations, created using a variety of methods (and excluding all parcels in low-SNR regions). These alternative parcellations included "Power ROIs": A set of functional ROIs derived from a combination of meta-analytic and functional connectivity analyses (Power et al. 2011); "Craddock": A parcellation created by the NCUT method (Craddock et al. 2012); "Shen": A parcellation created using a multiclass spectral clustering approach to the NCUT criterion (Shen et al. 2013); "Power communities": A parcellation created using the Infomap community detection technique (Power et al. 2011); "Yeo": A parcellation created using a signal clustering technique (Yeo et al. 2011); "Brodmann": A parcellation created from canonical Brodmann areas (Brodmann 1909); and "AAL": A parcellation created from the Automated Anatomical Labeling atlas (Tzourio-Mazoyer et al. 2002). Each parcellation was sampled to the cortical surface where necessary, and parcels containing <15 cortical vertices outside of low-SNR regions were eliminated from further analysis. For parcellation approaches with multiple solutions (the Craddock and Shen parcellations), we selected the solution with the number of parcels most similar to the boundary map-derived parcellation. We repeated these analyses for all other available Craddock and Shen parcellations with at least 50 parcels; these produced similar results to the chosen parcellation (Supplementary Figs 8 and 9). Table 1 provides additional details for each of these parcellations.

To ensure that the boundary map-derived parcellation created using Dataset 1 was not advantaged by being tested in the same dataset, we tested all parcellations' homogeneity using Dataset 2. For each parcellation scheme, we evaluated homogeneity using Dataset 2, and compared it with the homogeneity of randomly rotated versions of the parcellation.

Comparison of Parcels with Known Cytoarchitectonic Areas

If the boundary-derived parcellation created above is an accurate representation of the cortical areas in the brain, then it should contain parcels that are similar to known human cytoarchitectonic areas. We visually compared the boundary map-derived parcels to the probabilistic borders of areas 17, 1, 2, 3 (combining 3a and 3b), 4 (combining 4a and 4p), and hOc5 that were mapped to the 32k fs_LR by Van Essen et al. (2012) (publicly available through the SumsDB database, http://sumsdb.wustl.edu:8081/sums/index.jsp) based on cytoarchitectonic mapping by Fischl et al. (2008).

Identification of Parcel Network Structure

If the boundary-derived parcels created above are accurate representations of the cortical areas in the brain, then the network structure of the temporal correlations between these parcels should be highly similar to previously published descriptions of the network structure of the temporal correlations between all gray matter voxels.

Closely following Power et al. (2011), we assessed the network structure of the parcel-wise graph using the Infomap algorithm (Rosvall and Bergstrom 2008). In each subject, we calculated the average time course of each parcel from Dataset 1, and cross-correlated these time courses to form the parcel-wise correlation matrix. These correlation matrices were then Fisher-transformed and averaged across subjects. The resulting average correlation matrix was thresholded at a variety of correlation thresholds calculated to create connection matrices with specific degrees of sparseness (ranging from 1 to 3% of all possible connections surviving the threshold, in steps of 0.1%). Furthermore, connections passing these thresholds were removed if the geodesic distance along the cortex between the centroids of the connected parcels was <20 mm. The resulting connection matrices at each threshold were then evaluated using the Infomap algorithm, which assigned parcels to communities at each correlation threshold based on the maximization of within-community random walks in the connection matrix. Communities with 5 or fewer parcels were eliminated from consideration, and those parcels were considered unassigned.

We then collapsed across Infomap thresholds using a "consensus" procedure, with the goal of incorporating information both from more sparse thresholds, in which smaller networks were likely to emerge, and more dense thresholds, in which more parcels were likely to be successfully assigned. In this procedure, each node was given the community assignment it had at the sparsest possible threshold at which it was successfully assigned. The node assignments were "cleaned up" by removing small communities that were only present at one threshold. This procedure is nearly identical to the method used to collapse previously published voxel-wise community assignments (Power et al. 2011) across thresholds to create a single network map (the "Power communities" map described above). We note that this procedure does not attempt to comprehensively describe all features of the network, and may be especially poor at capturing nonhierarchical network features (which do occur infrequently). Rather, it provides a single, summary view of the brain's networks.

We assessed the overlap between the consensus parcel-wise network communities and the surface-mapped voxel-wise Power consensus communities described above. Overlap was calculated as the number of cortical vertices that had the same community identity in both parcel- and voxel-wise Infomap analyses divided by the total number of vertices that were assigned to a community in both analyses.

Use of Parcels in Individual Subjects

Ideally, the boundary-derived parcellation could be used to interrogate individual subject data. However, applying a group-level parcellation to individual subjects should only be performed if there is reasonable confidence that the parcellation truly does reflect the central tendency of the overall group, such that data in a given parcel from an individual will tend to look like the average data in that parcel across individuals. Thus, to determine whether the parcellation derived from the group boundary map could also be used to investigate individual subjects, we examined how reliably individual subjects' parcel connectivity maps looked like the group-average parcel connectivity maps.

Table 1
Previously published parcellations compared against present boundary map-derived parcellation

Name	Reference	Number of parcels	Notes
Power ROIs	ROI centers from Power et al. (2011), with subcortical ROIs excluded; available from http://www.nil.wustl.edu/labs/petersen/Resources_files/Consensus264.xls	226 (111 L, 115 R)	10 mm radius circles drawn around cortical locations. ROIs based, in part, on the same dataset used to create the boundary map.
Craddock	Craddock et al. (2012); available from http://ccraddock.github.io/cluster_roi/	353 (175 L, 178 R)	Sampled from volume to surface. Used 400-cluster parcellation generated by two-level group clustering using temporal correlation.
Shen	Shen et al. (2013); available from http://www.nitrc.org/frs/?group_id=51	213 (106 L, 107 R)	Sampled from volume to surface. Used 300-cluster parcellation.
Power communities	Power et al. (2011); available from http://sumsdb.wustl.edu:8081/sums/index.jsp	103 (53 L, 50 R)	Sampled from volume to surface. "Consensus" procedure applied to collapse across multiple thresholds. Separate parcel created from each spatially contiguous community cluster.
Yeo	Yeo et al. (2011); available from http://sumsdb.wustl.edu:8081/sums/index.jsp	98 (49 R, 49 L)	Used 17-cluster parcellation. Separate parcel created from each spatially contiguous cluster.
Brodmann	Brodmann (1909); available from http://sumsdb.wustl.edu:8081/sums/index.jsp	77 (39 L, 38 R)	Sampled from volume to surface.
AAL	Tzourio-Mazoyer et al. (2002); available from http://www.gin.cnrs.fr/AAL/aal_for_SPM8.tar.gz	74 (36 L, 38 R)	Sampled from volume to surface.

For each parcel in each Dataset 1 subject, we calculated the whole-brain subject-level connectivity pattern of the parcel by extracting the parcel's mean time course in that subject's data and correlating it against the time courses from every other gray matter point in the brain. We then averaged the Fisher's Z-transformed correlation patterns across subjects. Finally, we calculated the spatial correlation between the group-average Fisher-transformed connectivity pattern for that parcel and all of the Fisher-transformed subject-level connectivity patterns for the same parcel. This analysis produced a subject-group similarity (i.e., spatial correlation value) for each parcel in each subject.

We then explored two dimensions of variability in connectivity patterns: at the subject level and at the parcel level. First, we examined whether some subjects tended to be more or less similar to the group average than others, and whether the degree of similarity was related to the quantity of data remaining for each subject after motion correction. This was done by averaging similarity scores across parcels, for each subject, and then plotting these subject average scores against the number of uncensored time points in each subject's resting-state scan. Second, we examined whether some parcels' subject-specific connectivity patterns tended to be more or less similar to the group average than others. This was done by averaging similarity scores across subjects, for each parcel. Parcels with low average similarity scores can be considered unreliable for use in cross-subject analysis.

Results

Boundary Map Characteristics

Visually, the group boundary map (Fig. 2) appears very similar to our previously published boundary map (Wig, Laumann, Petersen 2014), though close examination indicates that the present boundary map appears cleaner, with sharper boundaries, and lower minimum boundary values. Comparison of histograms of the values in the current boundary map and the previously published boundary map (Supplementary Fig. 1) supports this observation, as the value distribution of the current map is markedly shifted to the left, suggesting a reduction in measurement noise in the map.

Boundary Map Reliability

Boundary maps from the two datasets appeared visually very similar. When thresholded at the top quartile of boundary map values, the boundary maps from the two datasets overlapped closely (Supplementary Fig. 2), with a Dice's coefficient of 0.71.

Parcel Creation

The parcel creation procedure produced 422 cortical parcels (206 in the left hemisphere and 216 in the right hemisphere; Fig. 3). Of these parcels, 356 (178 in each hemisphere) parcels were at least partly ($\geq$15 vertices, ~30 mm^2) outside low-SNR areas (Wig, Laumann, Petersen 2014). The remaining 66 parcels were considered unreliable due to low SNR and were excluded from further analysis.

Parcel Homogeneity

We calculated the homogeneity of each of these parcels within Dataset 1. Homogeneity represents the degree to which the parcel has a uniform connectivity pattern, and is thus a metric of parcel quality. Parcel homogeneities are mapped onto the brain as shown in Figure 3. Mean parcel homogeneity across all parcels was $89.1 \pm 5.8\%$ (maximum 98.4% and minimum 61.2%).

We then compared the mean homogeneity of this parcellation with a null model consisting of mean homogeneities from 1000 matched parcellations randomly rotated on the cortex. We observed that the mean homogeneity of the boundary-derived parcellation was much higher than any of the 1000 randomly rotated null model parcellation homogeneities (Fig. 3); the parcellation was thus significantly more homogenous than random at $P < 0.001$. These null model parcellations had a mean homogeneity of 85.6%, with a standard deviation of 0.29% across parcellations; the boundary-derived parcellation had a homogeneity Z score of 12.07 (i.e., was 12.07 SD away from the mean of the null model parcellations).

We further examined the relationship between parcel homogeneity and parcel size to determine whether the homogeneity measure was dependent on parcel size. The homogeneities of the real parcels (in red) and null model parcels (in gray, medians in black) are plotted against parcel size in Figure 3. We observed a close relationship between homogeneity and parcel size that can be appreciated with the Lowess fit line plotted on

Figure 2. RSFC-boundary map from Dataset 1. Bright colors indicate locations where abrupt transitions in RSFC patterns were reliably found across many cortical vertices, representing putative boundaries between cortical areas. Dim colors represent relatively stable RSFC patterns.

top; this relationship was observed both for the boundary-derived parcels (in red) and for the random-matched parcels (in gray; mean homogeneities of random parcels in black).

In summary, parcels derived from boundary maps are highly homogenous. Overall, this parcellation is also much more homogenous than a null model consisting of randomly replaced versions of the parcellation, suggesting that the present parcels are well placed. We further established that the homogeneity measure has a strong relationship with parcel size, justifying our use of the present null model, which accounts for parcel size and shape.

Comparison of Parcel Homogeneity Against Alternative Parcellations

To demonstrate external validity of the parcellation, we evaluated the homogeneity of the Dataset 1 boundary map-derived parcels using data from Dataset 2. The mean homogeneity across the boundary map-derived parcels was 87.4 ± 6.4%, which was similar to, but slightly lower than, the homogeneities of the parcels derived from and tested in Dataset 1 (89.1 ± 5.8%, as stated above).

As above, we compared the homogeneities of the boundary map-derived parcels with the null model consisting of randomly rotated versions of the parcellations. Once again, the boundary map-derived parcellation tested in Dataset 2 was more homogenous than any randomly rotated parcellation ($P < 0.001$); it had a Z score of 10.91 compared with the distribution of random parcellations.

We further evaluated the homogeneity of several alternative parcellations using Dataset 2. Parcel homogeneities from these alternative parcellations can be seen in Figure 4; average homogeneities of each parcellation are listed in Table 2.

We then compared the homogeneities of each alternative parcellation against the homogeneities of a null model consisting of 1000 randomly placed versions of the parcellation. See Table 2 for comparisons to the null model parcellations. The Power ROIs, Yeo parcels, and Brodmann parcels were more homogenous than any of their null model parcellations, while the Shen parcellation and Power communities were significantly better than the set of random parcellations, but not better than all possible random parcellations. The Craddock and AAL parcellations were not significantly more homogenous than their null models.

Finally, we examined homogeneity versus parcel size relationships for the boundary map-derived parcellation, as well as for each alternative parcellation (Supplementary Fig. 4). Relationships between homogeneity and size were observed for each parcellation and the null model of each parcellation, though size–homogeneity relationships appeared weaker for parcellations with less variance in parcel size, as would be expected. When the fit lines of all the parcellations and random parcellations were plotted on the same scale (Supplementary Fig. 5), it became evident that while all parcels exhibited homogeneity decreases as size increased, the boundary map-derived parcels had superior homogeneity even when parcel size was taken into account.

Comparison of Parcels with Known Cytoarchitectonic Areas

We observed a strong visual overlap between the boundary map-derived parcels and several known cytoarchitectonic areas. The left side of Figure 5 illustrates these overlaps on the left hemisphere (see Supplementary Fig. 6 for right hemisphere overlap). The architectonic boundary of area 17 almost perfectly encompassed a single RSFC-defined parcel in both hemispheres. In contrast, area hOc5 also appeared to correspond with a single parcel in the left hemisphere, but that parcel extended significantly beyond the probabilistic border of the area. In the right hemisphere, no parcel corresponded with area hOc5. Area 1 did not correspond with any parcels, falling directly on top of a border between parcels in both hemispheres.

In both hemispheres, cytoarchitectonically defined areas 2, 3, and 4 aligned well with a string of parcels running down the pre- and postcentral gryri. Taken together, these strings of parcels matched areas 3 and 4 almost perfectly, and overlapped most of area 2, failing only to capture a ventral posterior section of the area. Thus, we hypothesized that while the parcels do not conform well to strict anatomical definitions of cortical areas, they may be capturing some unknown functional subdivisions within the areas.

One possible functional subdivision these parcels could be capturing is the known somatotopic divisions within areas 2–4, in which dorsomedial somatomotor cortex receives sensory input and projects motor output to the feet, dorsolateral somatomotor cortex to the hands, and ventrolateral somatomotor cortex to the mouth and tongue. We conducted a post
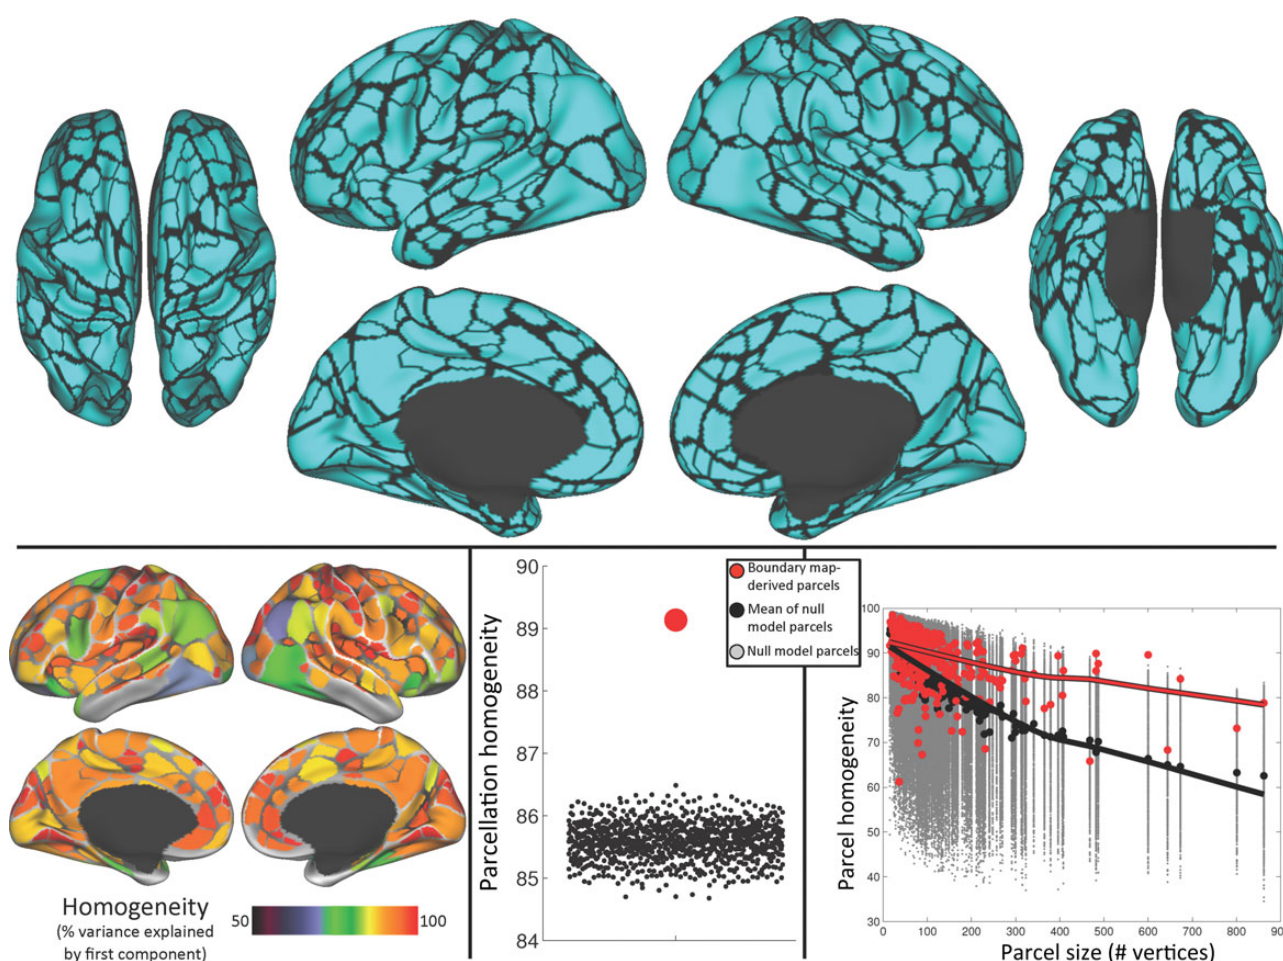
Figure 3. Boundary map-derived parcels are both highly homogenous and more homogenous than a null model. Top: 422 cortical parcels were created from the Dataset 1 boundary map. Bottom left: homogeneity of each parcel, calculated as the percent of the variance in RSFC patterns explained by the parcel's first PCA eigenvariate. Green indicates a parcel is >70% homogenous; red indicates >90% homogenous. Bottom middle: average homogeneity across parcels (red dot) was significantly higher than that across parcels of each null model iteration (black dots). Bottom right: homogeneity of individual real parcels (red dots) was higher than that of null model parcels (gray dots) when plotted against parcel size. Black dots indicate the median homogeneity across iterations for each null model parcel. Lowess fit lines in red and black emphasize the homogeneity–size relationship for the real and null model parcels, respectively.

hoc investigation of this possibility using results from a motor fMRI task collected as part of the Human Connectome Project. This task involved blocks of cued left or right finger tapping, left or right toe squeezing, and tongue movement (see Barch et al. 2013 for details). Preliminary findings from this task conducted in 20 subjects were presented in Barch et al. (2013); the present investigation used results from 219 subjects, analyzed using the same procedures as in Barch et al. (2013). We thresholded these data at a very high statistical threshold (arbitrarily selected to be $Z > 8.0$, though similar results were observed for any threshold between $Z > 5.0$ and $Z > 10.0$) and examined the overlap between the task activations and the various parcels in the pre- and postcentral gyri.

We observed that each task activation cluster very well-matched multiple parcels in the pre- and postcentral gyri (left hemisphere activations shown in Fig. 5, right; see Supplementary Fig. 6 for right hemisphere activations). The correspondence was particularly clear for the hand and tongue activation clusters. The left hemisphere (though not right hemisphere) foot cluster extended anteriorly and posteriorly to the pre- and postcentral gyri. Importantly, the dorsal/ventral borders of each activation cluster conformed very well to some of the parcel

borders that split the putative cytoarchitectonic areas into multiple parcels. This suggests that these borders represent differences in function within a topographically organized area that are not captured by cytoarchitectonics.

Parcel Network Structure

We conducted community detection in the parcel-wise graph across many density thresholds (see Supplementary Fig. 7 for results from all thresholds), and we collapsed across thresholds using a consensus procedure. There was a considerable visual overlap between the cross-threshold consensus parcel-wise communities (Fig. 6, top) and the Power communities (Fig. 6, middle). Every community found in the Power communities was also observed in the parcel-wise communities except for one in the anterior medial temporal lobe. These included all of the classic large scale RSFC networks/systems that have been consistently identified using multiple techniques (community detection, Power et al. 2011; independent components analysis, Beckmann et al. 2005; Smith et al. 2009; signal clustering, Yeo et al. 2011), such as Visual (dark blue in Fig. 6), Dorsal somatomotor (light blue), Ventral somatomotor (orange), Auditory
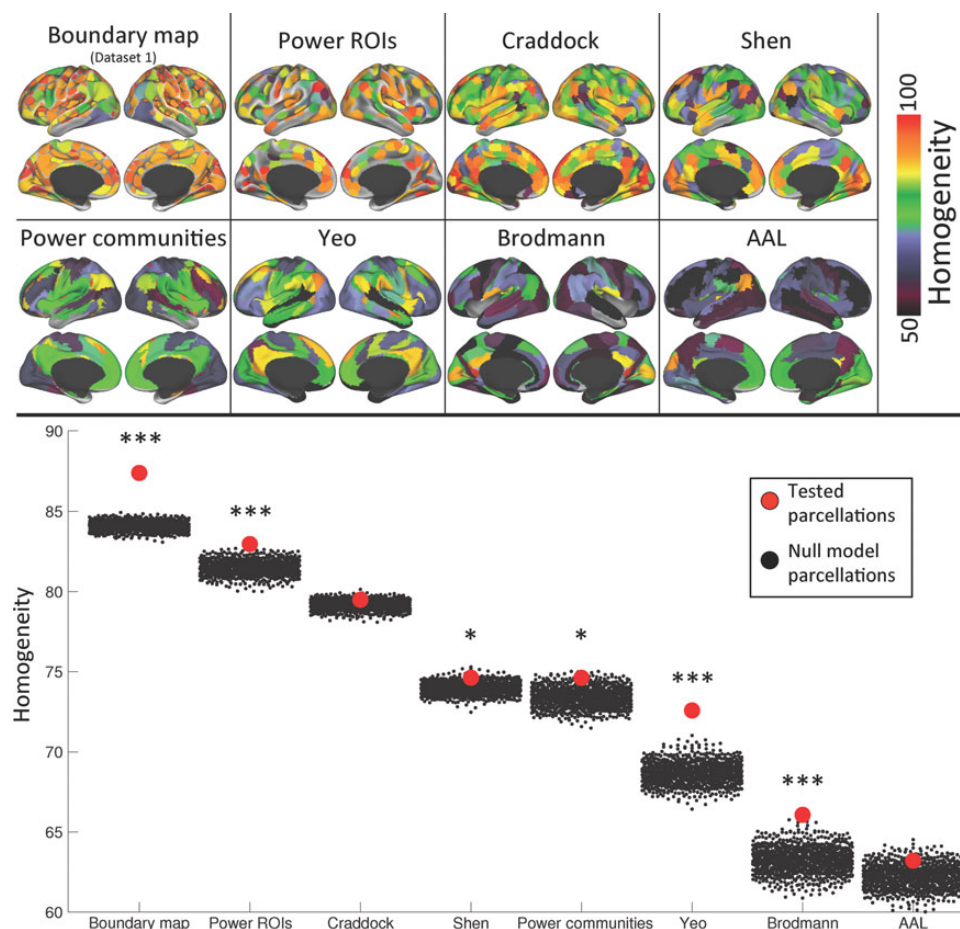
Figure 4. When tested in an independent dataset, the boundary map-derived parcellation is more homogenous than any other parcellation, and does better relative to its null model than any other parcellation. Top: parcel homogeneities of each competing parcellation when tested in Dataset 2. Bottom: average homogeneity across parcels of each parcellation (red dots) compared with the average homogeneity across parcels of each of 1000 null model iterations (black dots), which vary in homogeneity because of differing parcel sizes. ***indicates the parcellation was more homogenous than all of its 1000 null model iterations (i.e., $P < 0.001$); *indicates the parcellation was more homogenous than at least 950 of its null model iterations ($P < 0.05$).

Table 2

Average homogeneity and comparison of homogeneity against a null model for each parcellation

Parcellation	Mean ± SD homogeneity (tested in Dataset 2)	Number of null model iterations with worse homogeneity	Z-score relative to null model
Boundary map-derived	87.4 ± 6.4%	1000 ($P < 0.001$)	10.91
Power ROIs	83.0 ± 9.7%	1000 ($P < 0.001$)	3.29
Craddock	79.5 ± 10.4%	871 ($P = 0.129$)	1.12
Shen	74.6 ± 11.6%	960 ($P = 0.040$)	1.71
Power communities	74.6 ± 15.7%	976 ($P = 0.024$)	1.97
Yeo	72.6 ± 14.9%	1000 ($P < 0.001$)	5.70
Brodmann	66.1 ± 14.0%	1000 ($P < 0.001$)	3.51
AAL	63.2 ± 11.7%	920 ($P = 0.080$)	1.38

(light purple), Default (red), Frontoparietal (yellow), Dorsal attention (green), Cingulo-opercular (purple), Ventral attention (teal), and Salience (black). They also included a number of less well-known systems that have been identified only in more recent investigations (Power et al. 2011; Yeo et al. 2011), such as (1) A superior temporal sulcus-centered community (pink in Fig. 6); (2) a community in the anterior and posterior lateral frontal cortex, ventral inferior parietal lobule, and dorsomedial prefrontal cortex (tan); (3) a community in retrosplenial and

ventral temporal cortex (white); and (4) a community in posterior cingulate and ventral posterior precuneus (medium blue). Meanwhile, only one community emerged in the parcels was not found in the Power communities: A community in the marginal sulcus and frontal eye fields (colored magenta in Fig. 6, top).

Overall, the overlap between the two methods was 71.2%. Multiple parcels with 100% overlap were observed in medial prefrontal, parietal, and occipital cortices, anterior and posterior insulae, and pre- and postcentral gyri. In contrast, parcels with poor overlap between the two methods were observed in lateral occipital and retrosplenial cortex, marginal sulcus, and frontal eye fields (Fig. 6, bottom).

Use of Parcels in Individual Subjects

We examined how similar the group-average parcel connectivity patterns were to the connectivity patterns seeded from the same parcel in each individual. Across all subjects and parcels, the average Fisher-transformed spatial correlation ($Z(r)$) between subject and group connectivity patterns was $0.57 ± 0.15$. However, we observed that the average $Z(r)$ across parcels was not uniform across subjects, ranging from 0.34 to 0.69. We tested whether this variability was related to how much data had been collected on a subject. We observed a

Figure 5. Boundary map-derived parcels match known cortical areas and functional activation patterns. Left and middle: a variety of cytoarchitectonically defined cortical areas (Fischl et al. 2008) were matched by boundary map-derived parcels. Area 17 overlapped very well with one parcel, whereas area hOc5 overlapped moderately well with another parcel. Areas 2, 3, and 4 overlapped with several adjacent parcels. Right: parcel divisions within cytoarchitectonic areas 2, 3, and 4 corresponded with divisions between activation clusters from motor movements of the right foot, right hand, and tongue (Barch et al. 2013).


Figure 6. The network structure of the boundary map-derived parcellation closely corresponds with the previously described network structure of the brain. Top: communities identified with the Infomap community detection procedure using the boundary map-derived parcels as network nodes. See the text for names of each colored community. Middle: the network structure of the brain calculated using every voxel as a network node (Power et al. 2011). Bottom: spatial overlap of the parcel- and voxel-wise community assignments.

nonlinear relationship between the number of time points analyzed and the average subject-group $Z(r)$ across parcels (Fig. 7, left). The average $Z(r)$ for 84 subjects with <300 uncensored time points (12.5 min) ranged from 0.35 to 0.64, with 48 subjects having a $Z(r)$ <0.55, but the average $Z(r)$ for 36 subjects with >300 uncensored time points ranged from 0.56 to 0.68.

We also observed that subject-group $Z(r)$s were not uniform across different parcels, ranging from 0.32 to 0.73 (Fig. 7, top right). Specifically, parcels in medial occipital cortex, lateral and medial parietal cortex, insular cortex, medial prefrontal cortex, and pre/postcentral gyrus tended to have a $Z(r)$ around 0.6 or above, with parietal Default mode parcels (posterior cingulate/precuneus and angular gyrus) having the highest $Z(r)$,

around 0.7. When analysis was restricted to the 36 subjects with >300 time points, $Z(r)$ values increased in 355 of 356 parcels; however, the spatial pattern of $Z(r)$ across parcels did not change (Fig. 7, bottom right). This suggests that including subjects with insufficient data reduces the reliability of parcel connectivity estimates globally rather than in specific parcels.

Discussion

In this study, we described a method for building discrete parcels from RSFC-boundary maps. We also described a homogeneity-based metric to evaluate the quality of the parcellation, and we demonstrated that the boundary map-derived
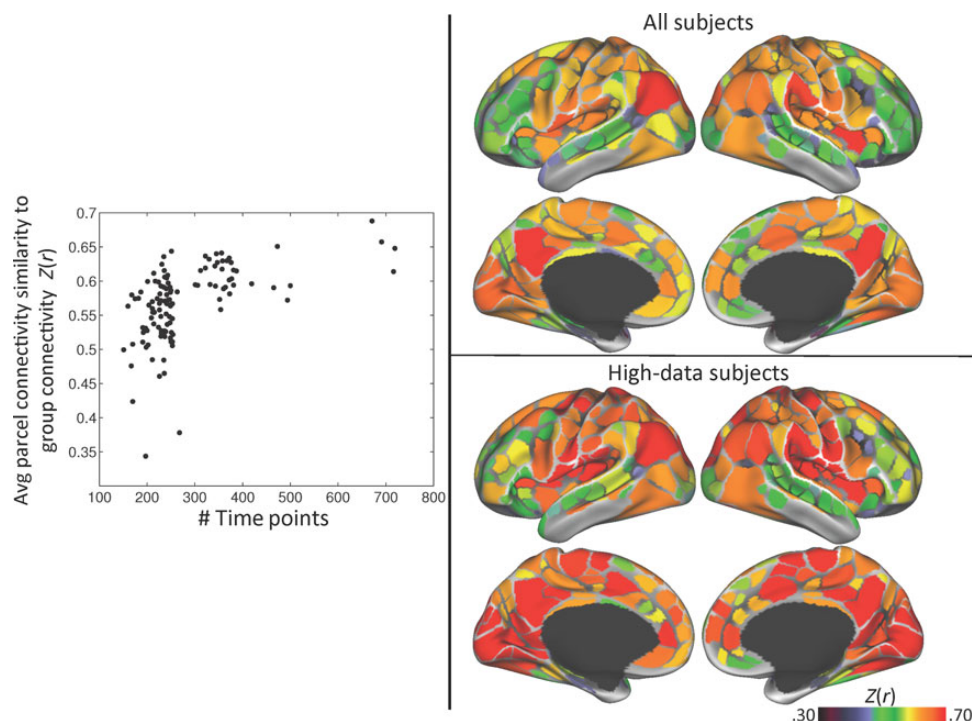
Figure 7. Group-average parcel connectivity is similar to subject-level connectivity, but this similarity varies across parcels and subjects. Left: the average Fisher-transformed correlation between group- and subject-level parcel connectivity patterns for each subject, plotted against the number of time points in each subject's resting-state data. Top right: the average group–subject correlation for each parcel, averaged across all subjects. Bottom right: the average group–subject correlation for each parcel, averaged across subjects with >300 time points (12.5 min) of resting-state data.

parcels were highly homogenous. We found that the parcellation was significantly more homogenous than size- and shape-matched random parcellations in two independent datasets. We also found that the boundary map-derived parcellation had higher overall homogeneity and performed better relative to random parcellations than a number of alternative parcellations. We additionally observed a high degree of overlap between the boundary map-derived parcels and several known cytoarchitectonic areas, with subdivisions within the cytoarchitectonic areas corresponding to functional differences. We further examined the network structure of the boundary map-derived parcels, and we found that it closely matched the previously described voxel-wise structure of the brain. Finally, we observed that boundary-derived parcel connectivity patterns were mostly reliable across individual subjects.

There are good a priori reasons to believe that RSFC-boundary maps have real utility for areal parcellation of human cortex. First, RSFC-based techniques are noninvasive and can be applied to any subject population that does not exhibit severe movement during scanning. Second, RSFC is believed to represent some combination of direct and indirect structural connectivity (Vincent et al. 2007; Honey et al. 2009) and a statistical history of functional coactivations (e.g., Dosenbach et al. 2007); as such, it reflects some combination of a region's function and connectivity, which are two of the major measures proposed to dissociate cortical areas (Felleman and Van Essen 1991). Third, RSFC-boundary maps, in particular, have been shown to not only identify where RSFC patterns change, but also to correspond with task activation patterns and to known areal borders based on architectonic divisions (Wig, Laumann, Petersen 2014). This cross-modality validation indicates that strong RSFC boundaries are very likely to index cortical area divisions in many cases.

Boundary Map-Based Parcellation Generates Parcels That Conform to Cytoarchitectonic Areas

We observed that the boundary map-derived parcellation contained parcels that had very strong overlap with the known extent of area 17, as defined by Fischl et al. (2008) and mapped to the cortical surface by Van Essen et al. (2012). Other known cortical areas, such as somatomotor areas 2, 3, and 4, were overlapped by a combination of several parcels. These observations—that parcel borders conform to cytoarchitectonically based estimates of human cortical areas—lend substantial face validity to the parcellation.

However, the fact that somatomotor areas were subdivided into multiple parcels suggests that the present parcellation does not faithfully replicate all architectonic areas, but may instead overparcellate some areas. We predicted that overparcellation would be most likely to occur in topographically organized architectonic areas, such as somatomotor cortex, that are known to have subregions with dissociable functional responses (Rao et al. 1995), including dissociable RSFC responses (Long et al. 2014). Indeed, overparcellation based on function is the most likely explanation for the subdivisions within somatomotor areas, as we observed that at least some of those subdivisions were functionally relevant, conforming to the boundaries between different functional activation patterns resulting from motor movements of different body parts. The present boundary map-derived parcellation should thus be considered a functional parcellation; as such, it provides complementary information about brain organization that cannot be observed via anatomy.

In contrast, area hOc5 (also known as the middle temporal+ complex) was only moderately well matched by a too-large parcel in the left hemisphere, and did not match any parcel in the right hemisphere. This failure to parcellate the area may be

related to the known individual variability in hOc5 (Malikovic et al. 2007), which is greater than that of any other area investigated here (Van Essen et al. 2012). Inconsistent locations of cortical areas across subjects would reduce the likelihood that the boundary mapping procedure can successfully identify the area's border.

In total, the boundary map-derived parcellation consisted of 422 discrete parcels. This number of parcels falls above the range of 150–200 human cortical areas per hemisphere estimated by Van Essen et al. (2012). It is possible that, like the somatomotor cortex, various other architectonic cortical areas may be functionally subdivided by the present parcellation, resulting in an inflated number of parcels.

Boundary Map-Based Parcellation Generates Parcels That Are Functionally Homogenous

Overall, the boundary-derived parcels had highly homogenous RSFC patterns, with an average parcel homogeneity of almost 90%. This high degree of homogeneity in RSFC patterns indicates that most parcels represented regions of uniform BOLD signal, which is an expected characteristic of most cortical areas.

Only a few parcels had low homogeneity (Fig. 3). Some of these parcels—e.g., in medial and anterior inferior temporal lobe, and in inferior insula—were near low-SNR areas, and may have had somewhat degraded signal; thus, low homogeneity is not surprising in these parcels. Other parcels—in right angular gyrus, right occipital cortex, bilateral occipitotemporal cortex, and left frontal eye fields—more likely represent local failures of the RSFC-boundary mapping procedure, in which a true border between cortical areas was not successfully delineated.

Homogeneity-Based Parcellation Evaluation Must Account for Parcel Size and Shape

The boundary map-derived parcels were not only highly homogenous, they were also much more homogenous than a null model consisting of 1000 identical parcellations that were randomly rotated into a new position on the cortical surface. The use of a null model is necessary for true evaluation of a parcellation, as the homogeneity measure of a given parcel is strongly dependent on the parcel's size (Fig. 3 and Supplementary Figs 4 and 5). A similar effect was reported by Craddock et al. (2012), who found that the homogeneity of both clustering-derived and random parcels varied strongly as a function of the number of clusters specified (which will vary inversely with parcel size). By examining homogeneities of individual parcels, we show that this effect is specifically driven by parcel size; this can be appreciated by examination of the parcel size versus homogeneity plots of the randomly rotated parcels (gray points in Fig. 3; black points represent the mean homogeneity across rotations). As discussed in the Methods section, this effect likely arises because small randomly placed parcels are more likely to fall within large homogenous regions such as the medial posterior parietal cortex, whereas large randomly placed parcels are more likely to sprawl across multiple cortical areas. The effect of parcel size is also likely constrained by the smoothness of the data, which is affected by averaging across variable subjects, the application of geodesic Gaussian smoothing during data processing, and the intrinsic spatial autocorrelation of the BOLD signal. If these explanations are correct, then a parcel's homogeneity

will depend not only on its size, but also on the regularity of the parcel's shape, as an elongated parcel is more likely to sprawl across multiple cortical areas and extend beyond the intrinsic smoothness of the data than a circular parcel with the same surface area. This means that any appropriate null model of homogeneity must account both for a parcel's size and its shape. Of previously published RSFC-based parcellation approaches, only Craddock et al. (2012) compared their parcellation to a null model; however, that null model was simply the same number of randomly generated parcels. That null model thus maintains the average parcel size, but it does not attempt to match these sizes on a parcel-to-parcel basis or to maintain the shape of parcels, as the present null model does.

Boundary Map-Derived Parcellation Performs Better Than Alternative Parcellations

We tested the homogeneity of the boundary map-derived parcellation using a second dataset, such that the parcel creation procedure was completely independent of the data in which it was tested. We found that the parcellation was still highly homogenous, and still much more homogenous than its null model, suggesting that these boundary map-derived parcels represent a robust central tendency of the population and can be applied to other datasets, even ones collected with different sequences on different scanners. Further, the boundary map-derived parcellation was both more homogenous than any other putative areal-level parcellation tested, and was more homogenous compared with its null model than any other parcellation tested. This suggests that it better represents functionally homogenous cortical areas than any of the other parcellations.

Parcellations derived from network detection approaches [the clustering-based approach proposed by Yeo et al. (2011) and the community detection procedure described in Power et al. (2011)] performed reasonably well when compared with their null models (particularly the Yeo parcellation), suggesting that these parcellations contain substantial information about the structure in the data. However, the raw homogeneities of the parcels in this parcellation were only moderate. This likely indicates that these approaches, which are designed to identify large scale brain systems or networks, do not parcellate the brain finely enough to represent subsystem-level distinctions between adjacent regions. Such distinctions, as demonstrated by Wig, Laumann, Petersen (2014), likely reflect areal divisions in the brain, as they indicate where multiple regions with similar but discrete connectivity patterns interact within larger systems. The fact that such divisions are not reflected in the Yeo and Power parcellations indicates that those parcellations are closer to system-level divisions of the brain than true parcellations of cortical areas.

Parcellations based on the NCUT criterion (Craddock et al. 2012; Shen et al. 2013) were moderately homogenous; however, the Shen parcellation was only marginally more homogenous than its null model, whereas the Craddock parcellation was not more homogenous than its null model. This poor performance on a homogeneity-based measure is surprising, given that clustering techniques such as these are designed to group similar signals together, which in theory should produce homogenous parcels. Blumensath et al. (2013) recently argued that parcels produced using the NCUT criterion described in Craddock et al. (2012) are dependent primarily on the specified cluster number rather than on the

underlying data, as highly reproducible NCUT parcels could be produced using random data. The present results are a further demonstration that NCUT-derived parcels do not represent the underlying data structure well.

The Brodmann parcellation (Brodmann 1909) had low homogeneity, but was more homogenous than any of its null model parcellations. This suggests that, like the Yeo and Power parcellations, this parcellation does successfully represent structure in the data, but is too underparcellated to represent true cortical areas. This perspective agrees with modern attempts to anatomically parcellate human cortex, which frequently observe more fine-grained architectonic divisions than those reported by Brodmann (e.g., Morris et al. 2000, retrosplenial cortex; Öngür et al. 2003, orbitofrontal cortex; Morosan et al. 2005, superior temporal gyrus; Caspers et al. 2006, inferior parietal cortex; Scheperjans et al. 2008, superior parietal cortex; Kujovic et al. 2013, extrastriate visual cortex).

The AAL parcellation (Tzourio-Mazoyer et al. 2002) had the lowest homogeneity of all parcellations and was not better than its null model. Indeed, there was no expectation that the AAL parcellation would represent the structure of RSFC data, as previous work has indicated that AAL regions are worse than RSFC-based parcellation schemes at representing cortical areas (Craddock et al. 2012; Blumensath et al. 2013; Shen et al. 2013).

The Power ROIs (from Power et al. 2011) had both high homogeneity and were significantly better than all null model parcellations. These ROIs, which were derived partly from an earlier, less precise version of the present boundary mapping procedures, have been used in the field for a variety of purposes, including investigation of motion-related artifacts (Power et al. 2012), functional connectivity dynamics (Glerean et al. 2012), task control processes (Cole et al. 2013), and deficits related to neuropsychological disorders such as Autism (Rudie et al. 2013), attention deficit hyperactivity disorder (Eloyan et al. 2012), and schizophrenia and bipolar disorders (Argyelan et al. 2014). The present results suggest that these ROIs are reasonable estimates of cortical area centers, though not of full cortical areas, as they do not attempt to define the boundaries of areas.

One other RSFC-based whole-brain areal parcellation scheme has recently been proposed (Blumensath et al. 2013), but we were not able to compare this scheme against the present boundary map-derived parcellation, as it was never applied to group-average data. Blumensath et al. reported that subject-level parcels could be created using a region growing approach constrained by hierarchical clustering. Furthermore, they reported that, compared with parcels derived using the NCUT technique (Craddock et al. 2012), these parcels were more reliable, better represented RSFC pattern transitions, and better aligned with task activation patterns. However, it is unclear if this method could produce reasonable group-average parcels.

Parcel-Based Network Structure Corresponds with Voxel-wise Network Structure

We used a community detection procedure (Infomap; Rosvall and Bergstrom 2008) to identify the network structure of boundary map-derived parcels, and we compared it with the previously described network structure of the brain defined using every voxel in the brain as a node (the "Power communities" described above; Power et al. 2011). Every community found in the Power communities was also observed in the parcel-wise communities except for one in the anterior medial temporal lobe. These included a number of large, highly replicated communities such as the Default, Frontoparietal, and Cingulo-Opercular communities. They also included smaller communities, such as a retrosplenial/temporal community, a cingulate–precuneus community, and a superior temporal lobe community, which have been identified only recently using advanced network analysis techniques (Power et al. 2011; Yeo et al. 2011). The observation that parcel-based communities replicate both large, easily detected RSFC systems and small, subtle RSFC systems indicates that the present parcellation captures the overall network structure of the brain in considerable detail. The fact that this detailed structure is represented without the need for voxel-level granularity suggests that the present parcellation is appropriate for use in certain network analyses, such as graph theory analysis, which benefit from a limited number of rational, neurobiologically based nodes in order to be interpretable (Wig et al. 2011; Power et al. 2013).

One additional community was observed in the parcel-wise analysis that has not been observed in previous work: A community in the marginal sulcus and frontal eye fields (magenta in Fig. 6). These areas were incorporated into the Cingulo-opercular and Dorsal attention systems, respectively, in the Power communities. We are not aware of any work demonstrating that these regions operate as a coherent unit; in contrast, it is well established that the frontal eye fields are a central node of the Dorsal attention system (Corbetta and Shulman 2002). Furthermore, we observed that this community only emerged at relatively sparse thresholds; at more dense thresholds, it was split and incorporated into Cingulo-opercular and Dorsal attention communities, as in the Power voxel-wise communities (Supplementary Fig. 7). We thus speculate that this newly observed community may represent an overseparation of existing communities rather than a real brain system.

Most Group-Defined Parcels Reliably Represent Individual Subject Connectivity, Especially for High-Data Subjects

An important goal of this work is to create parcels representing cortical areas that can be interrogated in individual subject data. Conducting fMRI analysis in a parcel-wise fashion is an ideal form of data reduction (Wig et al. 2011), as it involves analyzing several hundred relatively independent, homogenous parcel-averaged signals rather than 65 000+ noisy, nonindependent voxel signals. In principle, applying these parcels to subject-level task analysis would thus not only decrease the need for multiple comparison correction, but would greatly increase the power of the analysis, as averaging a homogenous signal across a parcel would reduce noise levels. We examined whether the boundary map-derived parcels could be used for individual subject analysis. We found that, on average, subject connectivity maps had high spatial correlations to group-level maps, suggesting that, in general, extracting and averaging subject-level data from a group-average parcel is a valid approach.

However, we also observed that this degree of similarity was not uniform across subjects and parcels. For a given subject,

connectivity similarity with the group was observed to be strongly and nonlinearly related to the amount of data the subject retained after motion censoring: Subjects with greater than 12.5 min of data had high average similarity to the group, whereas those with <12.5 min of data were variable in how similar they were to the group. This finding emphasizes the need to acquire large amounts of data for reliable RSFC estimates, which has been well characterized by Anderson et al. (2011), who similarly demonstrated nonlinear effects of scanning time on RSFC reliability. Specifically, they found that reliability increased as 1/sqrt(scanning time). A similar relationship may be present in the current data, though we found that fitting this curve to the scanning time/group similarity relationship explained only about 33% of the variance in group similarity, so we hesitate to draw any strong conclusions about the nature of this effect.

A number of parcels were observed to have high homogeneity, indicating that the parcel was well formed in the group, but nevertheless had relatively low subject-group similarity. The most likely explanation for this phenomenon is interindividual variability in functional connectivity. Indeed, the locations of the most variable parcels—in lateral prefrontal cortex and lateral temporal–occipital cortex (green and purple in Fig. 3)—correspond to regions previously reported to have particularly high intersubject variability in RSFC patterns (Mueller et al. 2013). While most boundary map-derived parcels are appropriate for subject-level data analysis, these few parcels may be too variable for such a purpose. Ideally, issues of intersubject variability could be avoided by creating single-subject parcels from subject-level boundary mapping. In theory, such subject-level parcels could then be matched to each other for averaging or comparison across subjects; this procedure would constitute an areal-level registration. Blumensath et al. (2013) previously demonstrated that wholebrain parcellations can be created at the individual subject level, though in that work no attempt was made to match parcels to each other across subjects, which would be needed for true parcel-level cross-subject analysis. Future work may explore the feasibility and utility of subject-level parcel matching.

Limitations

While this parcellation scheme is homogenous, replicates the network structure of the brain well, and has similar connectivity patterns across individuals, it may not yet constitute a truly reliable whole-brain parcellation. Most parcels are highly homogenous, but some (e.g., in lateral occipital cortex) appear to be inaccurate and/or underparcellated. Other regions may be somewhat overparcellated. For example, while the parcellation describes some subdivisions in somatomotor cortex that correspond with functional activation patterns, other subdivisions have no known functional relevance, and they divide the motor and somatosensory strips into an arguably implausible number of parcels. It is likely that more accurate parcellations addressing these issues may be generated in the future as higher resolution datasets with more per-subject time points (such as the Human Connectome Project) become available.

It should also be noted that the present approach results in a purely functional parcellation that, while containing substantial information about the location and extent of anatomical cortical areas, nevertheless does not perfectly converge with a true anatomical areal parcellation. Indeed, the topological functional organization of some cortical areas makes it unlikely that specific anatomical area boundaries could ever be derived from purely functional measures like RSFC. In their classic parcellation of macaque visual cortex, Felleman and Van Essen (1991) remark that ideally, each cortical area should be uniquely identifiable using any of several modalities (connectivity, architectonics, topographic organization, functional responses, or behavioral consequences of lesions). In practice, they found that not every area could be identified using every method; often only one or two of these methods dissociated a specific area. This suggests that comprehensive categorization of all cortical areas in the human cortex will require further data from additional modalities.

Conclusions

Here we demonstrate that parcels created from RSFC-boundary maps overlap with known architectonic areas and have highly homogenous connectivity patterns. We also demonstrate that these parcels are far more homogenous than a null model in two independent datasets, indicating that the parcellation not only captures the structure of the data, but that it generalizes across different subject pools, scanners, and scanning sequences. Furthermore, no other parcellation tested was as homogenous or had as large a homogeneity difference compared with its null model. The proposed parcellation scheme thus appears to better represent functional divisions within the human brain than any other RSFC-based parcellation scheme yet published. A modified version of this parcellation created by combining both datasets (Supplementary Fig. 10) is publicly available at http://www.nil.wustl.edu/labs/petersen/Resources.html.

Supplementary Material

Supplementary material can be found at: http://www.cercor.oxford journals.org/

Notes

References

Amunts K, Malikovic A, Mohlberg H, Schormann T, Zilles K. 2000. Brodmann's areas 17 and 18 brought into stereotaxic space—where and how variable? NeuroImage. 11:66–84.

Anderson JS, Ferguson MA, Lopez-Larson M, Yurgelun-Todd D. 2011. Reproducibility of single-subject functional connectivity measurements. AJNR Am J Neuroradiol. 32:548–555.

Argyelan M, Ikuta T, DeRosse P, Braga RJ, Burdick KE, John M, Kingsley PB, Malhotra AK, Szeszko PR. 2014. Resting-state fMRI connectivity impairment in schizophrenia and bipolar disorder. Schizophr Bull. 40:100–110.

Barch DM, Burgess GC, Harms MP, Petersen SE, Schlaggar BL, Corbetta M, Glasser MF, Curtiss S, Dixit S, Feldt C et al. 2013. Function in the human connectome: task-fMRI and individual differences in behavior. NeuroImage. 80:169–189.

Barnes KA, Nelson SM, Cohen AL, Power JD, Coalson RS, Miezin FM, Vogel AC, Dubis JW, Church JA, Petersen SE et al. 2012. Parcellation in left lateral parietal cortex is similar in adults and children. Cereb Cortex. 22:1148–1158.

Beckmann CF, DeLuca M, Devlin JT, Smith SM. 2005. Investigations into resting-state connectivity using independent component analysis. Philos Trans R Soc Lond B Biol Sci. 360:1001–1013.

Beucher S, Lantuejoul C. 1979. Use of watersheds in contour detection. Presented at the international workshop on image processing: real-time edge and motion detection/estimation. Rennes, France.

Biswal BB, Mennes M, Zuo X-N, Gohel S, Kelly C, Smith SM, Beckmann CF, Adelstein JS, Buckner RL, Colcombe S et al. 2010. Toward discovery science of human brain function. Proc Natl Acad Sci. 107:4734–4739.

Biswal BB, Yetkin FZ, Haughton VM, Hyde JS. 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn Reson Med. 34:537–541.

Blumensath T, Jbabdi S, Glasser MF, Van Essen DC, Ugurbil K, Behrens TEJ, Smith SM. 2013. Spatially constrained hierarchical parcellation of the brain with resting-state fMRI. NeuroImage. 76:313–324.

Brodmann K. 1909. Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. Leipzig: Barth.

Bullmore E, Sporns O. 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. Nat Rev Neurosci. 10:186–198.

Carmichael ST, Price JL. 1994. Architectonic subdivision of the orbital and medial prefrontal cortex in the macaque monkey. J Comp Neurol. 346:366–402.

Carmichael ST, Price JL. 1996. Connectional networks within the orbital and medial prefrontal cortex of macaque monkeys. J Comp Neurol. 371:179–207.

Carp J. 2013. Optimizing the order of operations for movement scrubbing: comment on Power et al. NeuroImage. 76:436–438.

Caspers S, Geyer S, Schleicher A, Mohlberg H, Amunts K, Zilles K. 2006. The human inferior parietal cortex: cytoarchitectonic parcellation and interindividual variability. NeuroImage. 33:430–448.

Cohen AL, Fair DA, Dosenbach NUF, Miezin FM, Dierker D, Van Essen DC, Schlaggar BL, Petersen SE. 2008. Defining functional areas in individual human brains using resting functional connectivity MRI. NeuroImage. 41:45–57.

Cole MW, Reynolds JR, Power JD, Repovs G, Anticevic A, Braver TS. 2013. Multi-task connectivity reveals flexible hubs for adaptive task control. Nat Neurosci. 16:1348–1355.

Corbetta M, Shulman GL. 2002. Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci. 3:201–215.

Craddock RC, James GA, Holtzheimer PE, Hu XP, Mayberg HS. 2012. A whole brain fMRI atlas generated via spatially constrained spectral clustering. Hum Brain Mapp. 33:1914–1928.

Dale AM, Fischl B, Sereno MI. 1999. Cortical surface-based analysis: I. segmentation and surface reconstruction. NeuroImage. 9:179–194.

Dale AM, Sereno MI. 1993. Improved localizadon of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. J Cogn Neurosci. 5:162–176.

Damoiseaux JS, Rombouts SA, Barkhof F, Scheltens P, Stam CJ, Smith SM, Beckmann CF. 2006. Consistent resting-state networks across healthy subjects. Proc Natl Acad Sci USA. 103:13848–13853.

Dosenbach NUF, Fair DA, Miezin FM, Cohen AL, Wenger KK, Dosenbach RAT, Fox MD, Snyder AZ, Vincent JL, Raichle ME et al. 2007. Distinct brain networks for adaptive and stable task control in humans. Proc Natl Acad Sci. 104:11073–11078.

Eloyan A, Muschelli J, Nebel MB, Liu H, Han F, Zhao T, Barber AD, Joel S, Pekar JJ, Mostofsky SH et al. 2012. Automated diagnoses of attention deficit hyperactive disorder using magnetic resonance imaging. Front Syst Neurosci. 6:61.

Felleman DJ, Van Essen DC. 1991. Distributed hierarchical processing in the primate. Cereb Cortex. 1:1–47.

Fischl B, Rajendran N, Busa E, Augustinack J, Hinds O, Yeo BTT, Mohlberg H, Amunts K, Zilles K. 2008. Cortical folding patterns and predicting cytoarchitecture. Cereb Cortex. 18:1973–1980.

Fischl B, Sereno MI, Dale AM. 1999. Cortical surface-based analysis: II: inflation, flattening, and a surface-based coordinate system. NeuroImage. 9:195–207.

Fox MD, Raichle ME. 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nat Rev Neurosci. 8:700–711.

Friston KJ, Williams S, Howard R, Frackowiak RSJ, Turner R. 1996. Movement-related effects in fMRI time-series. Magn Reson Med. 35:346–355.

Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, Xu J, Jbabdi S, Webster M, Polimeni JR et al. 2013. The minimal preprocessing pipelines for the Human Connectome Project. NeuroImage. 80:105–124.

Glasser MF, Van Essen DC. 2011. Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI. J Neurosci. 31:11597–11616.

Glerean E, Salmi J, Lahnakoski JM, Jääskeläinen IP, Sams M. 2012. Functional magnetic resonance imaging phase synchronization as a measure of dynamic functional connectivity. Brain Connectivity. 2:91–101.

Hirose S, Watanabe T, Jimura K, Katsura M, Kunimatsu A, Abe O, Ohtomo K, Miyashita Y, Konishi S. 2012. Local signal time-series during rest used for areal boundary mapping in individual human brains. PLoS ONE. 7:e36496.

Hirose S, Watanabe T, Wada H, Imai Y, Machida T, Shirouzu I, Miyashita Y, Konishi S. 2013. Functional relevance of micromodules in the human association cortex delineated with high-resolution FMRI. Cereb Cortex. 23:2863–2871.

Honey CJ, Sporns O, Cammoun L, Gigandet X, Thiran JP, Meuli R, Hagmann P. 2009. Predicting human resting-state functional connectivity from structural connectivity. Proc Natl Acad Sci. 106:2035–2040.

Kujovic M, Zilles K, Malikovic A, Schleicher A, Mohlberg H, Rottschy C, Eickhoff SB, Amunts K. 2013. Cytoarchitectonic mapping of the human dorsal extrastriate cortex. Brain Struct Funct. 218:157–172.

Lancaster JL, Glass TG, Lankipalli BR, Downs H, Mayberg H, Fox PT. 1995. A modality-independent approach to spatial normalization of tomographic images of the human brain. Hum Brain Mapp. 3:209–223.

Lewis JW, Van Essen DC. 2000. Mapping of architectonic subdivisions in the macaque monkey, with emphasis on parieto-occipital cortex. J Comp Neurol. 428:79–111.

Long X, Goltz D, Margulies DS, Nierhaus T, Villringer A. 2014. Functional connectivity-based parcellation of the human sensorimotor cortex. Eur J Neurosci. 39:1332–1342.

Malikovic A, Amunts K, Schleicher A, Mohlberg H, Eickhoff SB, Wilms M, Palomero-Gallagher N, Armstrong E, Zilles K. 2007. Cytoarchitectonic analysis of the human extrastriate cortex in the region of V5/MT+: a probabilistic, stereotaxic map of area hOc5. Cereb Cortex. 17:562–574.

Markov NT, Ercsey-Ravasz MM, Gomes ARR, Lamy C, Magrou L, Vezoli J, Misery P, Falchier A, Quilodran R, Gariel MA et al. 2014. A weighted and directed interareal connectivity matrix for macaque. Cereb Cortex. 24:17–36.

Miezin FM, Maccotta L, Ollinger JM, Petersen SE, Buckner RL. 2000. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. NeuroImage. 11:735–759.

Morosan P, Schleicher A, Amunts K, Zilles K. 2005. Multimodal architectonic mapping of human superior temporal gyrus. Anat Embryol. 210:401–406.

Morris R, Paxinos G, Petrides M. 2000. Architectonic analysis of the human retrosplenial cortex. J Comp Neurol. 421:14–28.

Mueller S, Wang D, Fox MD, Yeo BTT, Sepulcre J, Sabuncu MR, Shafee R, Lu J, Liu H. 2013. Individual variability in functional connectivity architecture of the human brain. Neuron. 77:586–595.

Mugler JP III, Brookeman JR. 1990. Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP RAGE). Magn Reson Med. 15:152–157.

Nelson SM, Cohen AL, Power JD, Wig GS, Miezin FM, Wheeler ME, Velanova K, Donaldson DI, Phillips JS, Schlaggar BL et al. 2010. A parcellation scheme for human left lateral parietal cortex. Neuron. 67:156–170.

Nelson SM, Dosenbach NUF, Cohen AL, Wheeler ME, Schlaggar BL, Petersen SE. 2010. Role of the anterior insula in task-level control and focal attention. Brain Struct Funct. 214:669–680.

Ojemann JG, Akbudak E, Snyder AZ, McKinstry RC, Raichle ME, Conturo TE. 1997. Anatomic localization and quantitative analysis of gradient refocused echo-planar fMRI susceptibility artifacts. NeuroImage. 6:156–167.

Öngür D, Ferry AT, Price JL. 2003. Architectonic subdivision of the human orbital and medial prefrontal cortex. J Comp Neurol. 460:425–449.

Paxinos G, Huang X-F, Toga AW. 2000. The rhesus monkey brain in stereotaxic coordinates. San Diego (CA): Academic Press.

Petersen SE, Fox PT, Posner MI, Mintun M, Raichle ME. 1988. Positron emission tomographic studies of the cortical anatomy of single-word processing. Nature. 331:585–589.

Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE. 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. NeuroImage. 59:2142–2154.

Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL et al. 2011. Functional network organization of the human brain. Neuron. 72:665–678.

Power JD, Mitra A, Laumann TO, Snyder AZ, Schlaggar BL, Petersen SE. 2014. Methods to detect, characterize, and remove motion artifact in resting state fMRI. NeuroImage. 84:320–341.

Power JD, Schlaggar BL, Lessov-Schlaggar CN, Petersen SE. 2013. Evidence for hubs in human functional brain networks. Neuron. 79:798–813.

Rao SM, Binder JR, Hammeke TA, Bandettini PA, Bobholz JA, Frost JA, Myklebust BM, Jacobson RD, Hyde JS. 1995. Somatotopic mapping of the human primary motor cortex with functional magnetic resonance imaging. Neurology. 45:919–924.

Rosvall M, Bergstrom CT. 2008. Maps of random walks on complex networks reveal community structure. Proc Natl Acad Sci USA. 105:1118–1123.

Rudie JD, Brown JA, Beck-Pancer D, Hernandez LM, Dennis EL, Thompson PM, Bookheimer SY, Dapretto M. 2013. Altered functional and structural brain network organization in autism. NeuroImage: Clinical. 2:79–94.

Saleem KS, Price JL, Hashikawa T. 2007. Cytoarchitectonic and chemoarchitectonic subdivisions of the perirhinal and parahippocampal cortices in macaque monkeys. J Comp Neurol. 500:973–1006.

Scheperjans F, Eickhoff SB, Hömke L, Mohlberg H, Hermann K, Amunts K, Zilles K. 2008. Probabilistic maps, morphometry, and variability of cytoarchitectonic areas in the human superior parietal cortex. Cereb Cortex. 18:2141–2157.

Schleicher A, Palomero-Gallagher N, Morosan P, Eickhoff SB, Kowalski T, de Vos K, Amunts K, Zilles K. 2005. Quantitative architectural analysis: a new approach to cortical mapping. Anat Embryol. 210:373–386.

Ségonne F, Dale AM, Busa E, Glessner M, Salat D, Hahn HK, Fischl B. 2004. A hybrid approach to the skull stripping problem in MRI. NeuroImage. 22:1060–1075.

Ségonne F, Grimson E, Fischl B. 2005. A genetic algorithm for the topology correction of cortical surfaces. In: Christensen GE, Sonka M, editors. Information Processing in Medical Imaging. Presented at the 19th International IPMI Conference. Glenwood Springs, CO, USA: Springer. p. 393.

Sejnowski TJ, Churchland PS. 1989. Brain and cognition. In: Posner MI, editor. Foundations of cognitive science. Cambridge (MA): MIT Press. p. 888.

Shehzad Z, Kelly AMC, Reiss PT, Gee DG, Gotimer K, Uddin LQ, Lee SH, Margulies DS, Roy AK, Biswal BB et al. 2009. The resting brain: unconstrained yet reliable. Cereb Cortex. 19:2209–2229.

Shen X, Tokoglu F, Papademetris X, Constable RT. 2013. Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. NeuroImage. 82:403–415.

Smith SM, Fox PT, Miller KL, Glahn DC, Fox PM, Mackay CE, Filippini N, Watkins KE, Toro R, Laird AR et al. 2009. Correspondence of the brain's functional architecture during activation and rest. Proc Natl Acad Sci. 106:13040–13045.

Snyder AZ. 1996. Difference image versus ratio image error function forms in PET-PET realignment. In: Myer R, Cunningham VJ, Bailey DL, Jones T, editors. Quantification of brain function using PET. San Diego (CA): Academic Press. p. 131–137.

Talairach J, Tournoux P. 1988. Co-planar stereotaxic atlas of the human brain. New York: Thieme Medical Publishers, Inc.

Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M. 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. NeuroImage. 15:273–289.

Van Essen DC, Drury HA, Dickson J, Harwell J, Hanlon D, Anderson CH. 2001. An integrated software suite for surface-based analyses of cerebral cortex. J Am Med Inform Assoc. 8:443–459.

Van Essen DC, Glasser MF, Dierker DL, Harwell J, Coalson T. 2012. Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. Cereb Cortex. 22:2241–2262.

Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, Raichle ME. 2007. Intrinsic functional architecture in the anaesthetized monkey brain. Nature. 447:83–86.

Wig GS, Laumann TO, Cohen AL, Power JD, Nelson SM, Glasser MF, Miezin FM, Snyder AZ, Schlaggar BL, Petersen SE. 2014. Parcellating an individual subject's cortical and subcortical brain structures using snowball sampling of resting-state correlations. Cereb Cortex. 24:2036–2054.

Wig GS, Laumann TO, Petersen SE. 2014. An approach for parcellating human cortical areas using resting-state correlations. NeuroImage. 93:276–291.

Wig GS, Schlaggar BL, Petersen SE. 2011. Concepts and principles in the analysis of brain networks. Ann N Y Acad Sci. 1224:126–146.

Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR et al. 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J Neurophysiol. 106:1125–1165.