

A computational model of shared fine-scale structure in the human connectome

J. Swaroop Guntupalli^{1,2}, James V. Haxby^{1,2}

1. Department of Psychological and Brain Sciences, Dartmouth College, Hanover, NH, 03755 USA
2. Center for Cognitive Neuroscience, Dartmouth College, Hanover, NH 03755 USA

Corresponding author:

James V. Haxby

Center for Cognitive Neuroscience

Department of Psychological and Brain Sciences,

6207 Moore Hall, Hanover, NH 03755

USA

james.v.haxby@dartmouth.edu

Running title: Common model of the human connectome

Abstract

Local variation in cortical connectivity profiles is typically modeled as having a coarse spatial scale that can be parcellated into cortical systems of large, interconnected brain areas. We created a common model of the human connectome using a new algorithm – “connectivity hyperalignment” – that discovers shared basis functions for connectivity topographies. Transforming individual connectivity data into the common model connectome dramatically increases intersubject correlation of connectivity profiles. More importantly, the common model connectome captures shared fine-scale spatial variation in connectivity profiles that is obscured by other methods of aggregating multi-subject data into a common space. We found similar results using fMRI data obtained in the resting state and while subjects watched a movie. Fine-scale variations in connectivity profiles are closely related to fine-scale distinctions in representations of information. These results reveal the existence of a shared fine-scale structure in the human connectome that was not incorporated in previous models.

Introduction

Functional magnetic resonance imaging (fMRI) data collected while subjects are at rest reveals patterns of functional connectivity that are used to investigate the human connectome and parcellate the brain into interconnected areas that form brain systems and can be modeled as networks (Biswal, Yetkin, Haughton, & Hyde, 1995; Glasser et al., 2016; Gordon et al., 2016; Laumann et al., 2015; Smith, Vidaurre, et al., 2013; Sporns, Chialvo, Kaiser, & Hilgetag, 2004; Thomas Yeo et al., 2011; Wig et al., 2014). Aggregating individual data into a general model is typically accomplished by aligning brains based on anatomical features. Individual variations from the general model are described at the level of areal organization and characterized in terms of the absence or anomalous locations of functional areas (Gordon, Laumann, Adeyemo, & Petersen, 2015; Laumann et al., 2015). The connectivity of a single area is modeled as a mean connectivity profile and considered to be relatively homogeneous (Cohen et al., 2008).

Cortical topography has both a coarse scale of cortical areas and a finer scale of multiplexed topographies within areas (Grill-Spector & Weiner, 2014; Guntupalli, Hanke, et al., 2016; Haxby et al., 2001, 2011). Fine-scale within-area topographies are reflected in patterns of activity that can be measured with fMRI and decoded using multivariate pattern analysis (MVPA) (Haxby et al., 2001; Haxby, Connolly, & Guntupalli, 2014; Norman, Polyn, Detre, & Haxby, 2006). For example, patterns of activity in primary visual cortex distinguish responses to natural image stimuli that reflect local receptive fields for Gabor filters of different orientations and spatial frequencies (Kay, Naselaris, Prenger, & Gallant, 2008). Patterns of activity in the fusiform face area (FFA) (Kanwisher, McDermott, & Chun, 1997) distinguish responses to different head angles, face identities, animal species, and small objects (Guntupalli, Wheeler, & Gobbini, 2016; Haxby et al., 2001, 2011), but the stimulus feature spaces that determine these patterns is not well understood. These fine-scale patterns of activity are not well-aligned across brains based on anatomical landmarks. Consequently, averaging responses across brains blurs these patterns, making them less distinctive (Haxby et al., 2011). We have developed an algorithm, response hyperalignment (RHA) (Guntupalli, Hanke, et al., 2016; Haxby et al., 2011), that models variations in response tunings with a large set of basis functions that are shared across brains. Using RHA to transform individual cortical

representational spaces into a high-dimensional common model representational space aligns the fine-scale structure of within-area functional topographies. We ask here whether local variation in functional connectivity also has a fine-scale structure, similar to the fine-scale spatial structure of response tuning profiles, and whether such variation can be captured in a common model with basis functions that are shared across brains.

We developed a new algorithm, connectivity hyperalignment (CHA), to discover basis functions that are shared across brains and model variation in connectivity profiles. We show that CHA can derive these shared basis functions based on functional connectivities derived from neural activity during the resting state and during a task, namely watching and listening to movie. Transformation of individual connectomes into the resultant common model of the human connectome greatly increases intersubject correlation of local connectivity profiles. Moreover, the common model connectome also captures a fine spatial scale in local variation of connectivity profiles with a granularity equivalent to what we observe in individual connectomes. Although this fine-scale local variability in connectivity profiles can be observed and modeled in individual data (Haak, Marquand, & Beckmann, 2016; Jbabdi, Sotiropoulos, & Behrens, 2013), its existence was not incorporated into models of the connectome because it is obscured when connectomes are averaged across brains after anatomical alignment. We show further that this local variability in functional connectivity profiles is meaningful in that it is closely related to local patterns of responses during movie viewing that encode fine-scale distinctions among representations. Our results indicate that this shared fine-scale local variation accounts for a large portion of variance in connectivity profiles and is, therefore, a dominant component that was previously overlooked in models of the human connectome.

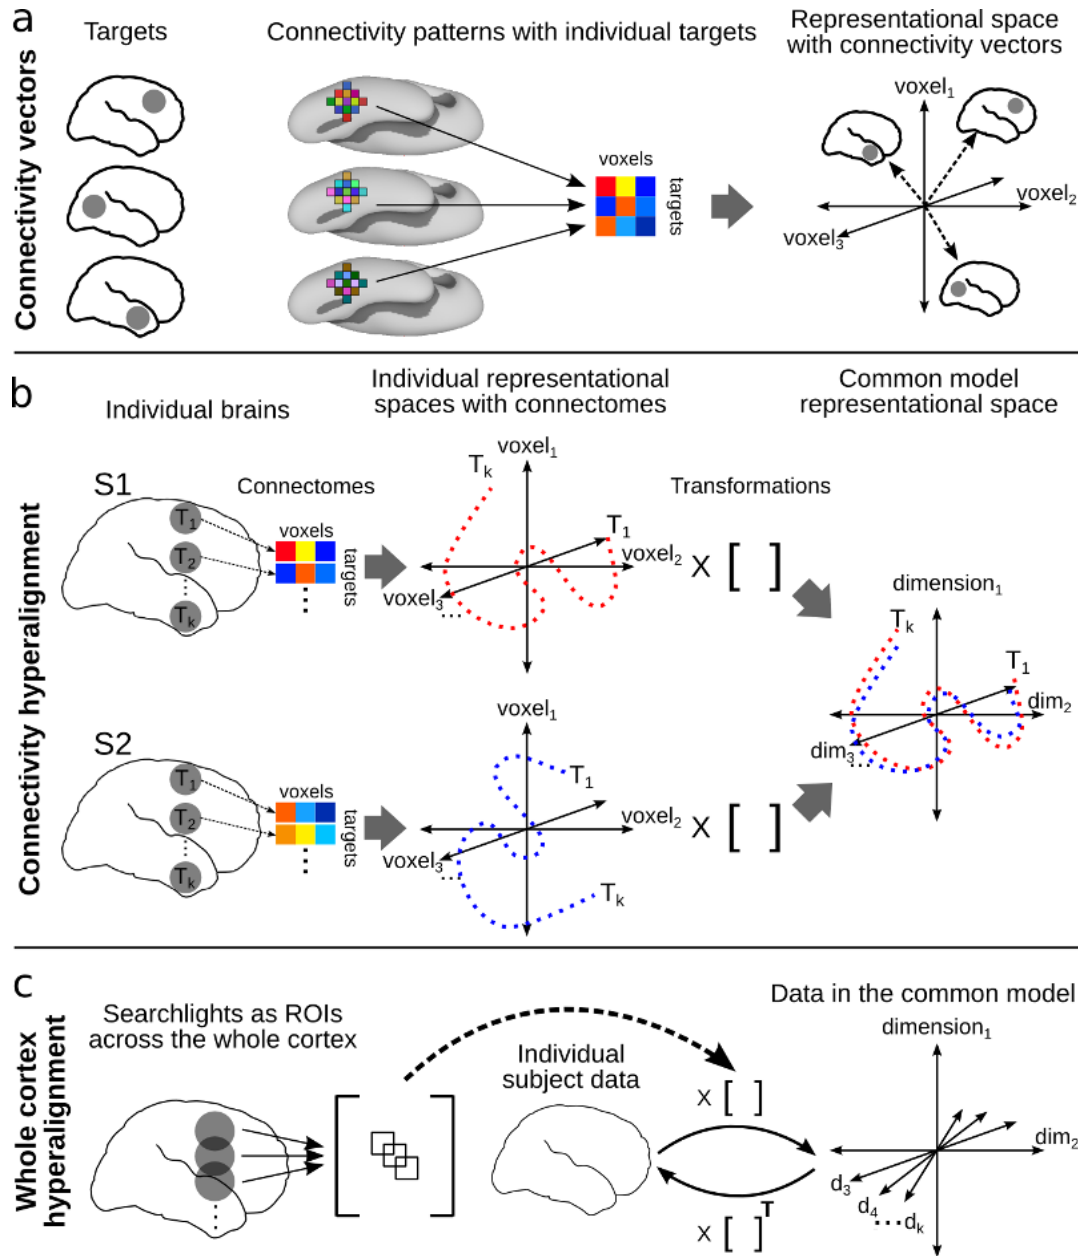


Figure 1. Schematic of connectivity hyperalignment (CHA). (a) Connectivity can be defined as any measure of similarity between a cortical locus (eg. voxel) and target region. Connectivities of loci in a region of interest or a searchlight to a target region is associated with a connectivity pattern in that region of interest or a searchlight for that target. These patterns can be thought of as connectivity vectors in a space in which each cortical locus in that region is a dimension. (b) Connectivity vectors in a region of interest or a searchlight to be hyperaligned are calculated for target regions distributed uniformly across the whole cortex. Connectivity hyperalignment rotates individual subject connectivity spaces to align these vectors across subjects. (c) Whole cortex connectivity hyperalignment extends connectivity hyperalignment to the whole cortex by performing connectivity hyperalignment in searchlights uniformly covering the cortex. Local transformations derived in each searchlight are aggregated by summation into a single transformation matrix that aligns data from the whole cortex into the common space and affords projection of vectors in common space back into individual cortical patterns by multiplication with the transpose of the transformation matrix.

Results

We derived a common model of the human connectome by applying whole cortex connectivity hyperalignment (CHA; Fig. 1) to two different fMRI datasets: 1) data collected for the HCP during task-free resting state and a battery of tasks (Barch et al., 2013; Smith, Beckmann, et al., 2013; Van Essen et al., 2013) and 2) data collected while subjects watched a full-length audio-visual feature film - “Raider of the Lost Ark” (Guntupalli, Hanke, et al., 2016; Haxby et al., 2011). We validated the common model in terms of 1) inter-subject correlation (ISC) of connectivity profiles, 2) spatial specificity of shared connectivity profiles. The results show that the common model connectivity captures shared fine-scale structure. To test whether this fine-scale structure is meaningful for the representation of information, we tested 3) ISC of task maps in the HCP data, 4) ISC of representational geometry in the movie data, and 5) between-subject multivariate pattern classification (bsMVPC) of responses to the movie.

Inter-subject correlation of connectivity profiles

Resting state data. CHA based on resting state connectivity profiles greatly increased ISC of resting state connectivity profiles in independent resting state data. The connectivity profile of a feature is measured as the correlations of its time-series with the average time-series in a set of target searchlights that covered the whole cortex (Fig. 1a). The centers of these target searchlights were identical to those used for CHA. Figure 2a shows a cortical map of mean ISCs of connectivity profiles in MSMAll-aligned data and in the common model connectome space, as well as within-subject between-session correlations of connectivity profiles. ISC is measured as the correlation of the connectivity profile for a given cortical location in each subject with the mean of other subjects' profiles. Within-subject correlations are computed between connectivity profiles for different resting state sessions. We quantify the increases in 26 functional ROIs, identified using a meta-analytic database, NeuroSynth (Yarkoni, Poldrack, Nichols, Essen, & Wager, 2011) (Fig. 2b; Supplementary Table 1). CHA increased mean ISC of connectivity profiles across these ROIs, relative to anatomical alignment, from 0.56 to 0.72 (95% confidence interval, CI, for mean difference = [0.14, 0.17]). Note that data alignment before CHA of the HCP dataset uses the MSM-All (Multimodal

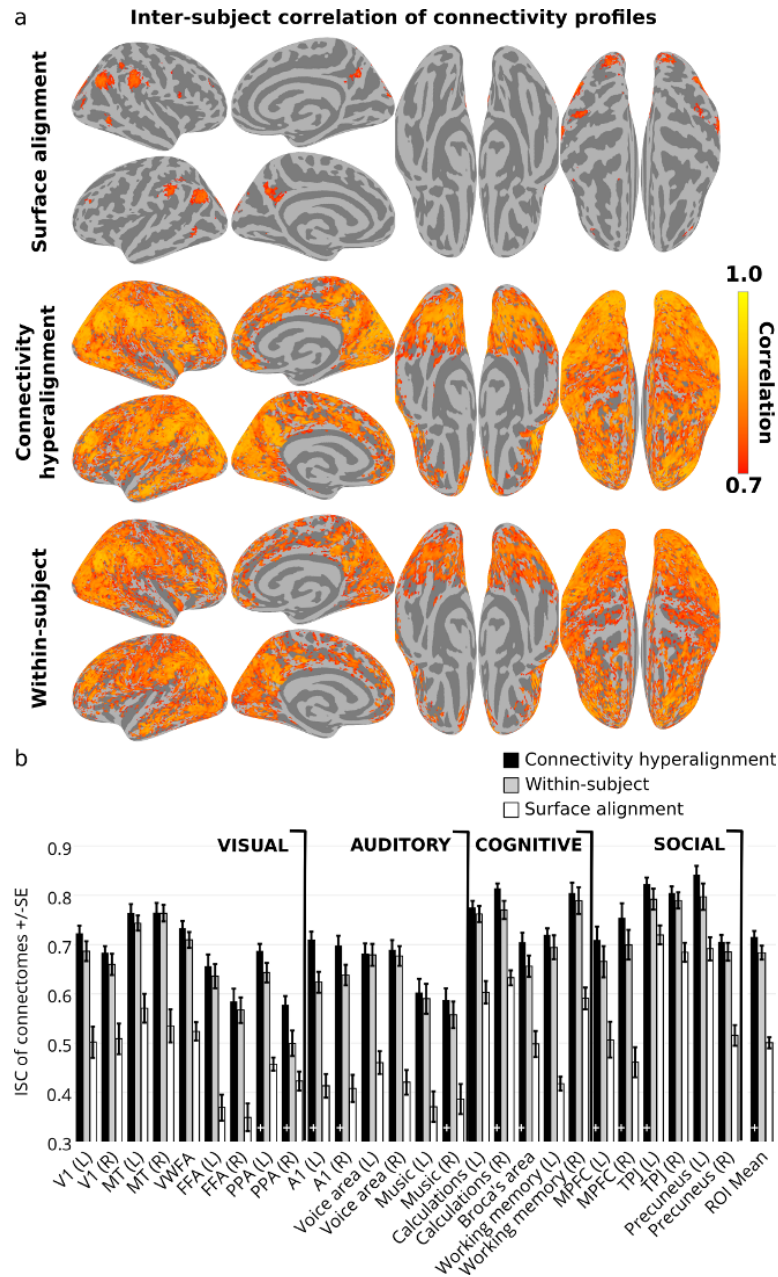


Figure 2. ISC of connectivity profiles calculated from HCP rsfMRI data. (a) Average ISC of connectivity profiles in each surface node after surface alignment (MSM-All) and in the common model connectome space, and within-subject between-session correlation in common space. (b) ISC of connectivity profiles averaged within ROIs covering visual, auditory, cognitive, and social systems across the cortex and their mean comparing the common model connectome space, within-subject between-session correlation in common space, and surface alignment. Bootstrapped testing showed significantly higher within-subject correlations and ISCs in the common model connectome space relative to anatomical alignment in all ROIs, and significantly higher ISCs in the common model connectome space than within-subject correlations in some ROIs (“+”: ISC after CHA > within-subject between-session correlation).

Surface Matching using individual cortical folding, myelin map, and rsfMRI correlation data) method which aligns subjects' data using both functional and structural properties (Robinson et al., 2014). Furthermore, ISCs of resting state connectivity profiles in the common model connectome space are slightly, but significantly, higher than within-subject correlations of resting state connectivity profiles (mean correlation = 0.68; CI for difference = [0.01, 0.06]) (Fig. 2b). This latter result indicates that after transformation into the common model space, an individual's connectome is better predicted by other subjects' connectomes (N=19) than by estimates based on that subject's own rsfMRI data, due to the benefit of aggregating data across a large sample of brains.

Movie data. CHA increased the ISC of connectivity profiles for each feature in the movie data as well. Figure 3a shows the mean ISCs of the connectivity profiles mapped onto cortical surface after anatomical alignment and CHA (see Supplementary Fig. 1 for control analysis). Figure 3b shows mean ISCs in anatomically-aligned data and in the common model connectome space in functionally-defined ROIs. Connectivity profiles of features are better aligned in the common model connectome space, relative to an anatomically aligned model (ROI means = 0.82 and 0.66, respectively; CI for the difference = [0.14, 0.17] (Figure 3b).

Local spatial granularity of the common model connectome

Resting state data. We investigated the spatial specificity of the common model connectome derived from rsfMRI data by computing spatial point spread functions (PSFs) of ISCs of connectivity profiles (Guntupalli, Hanke, et al., 2016). PSF of connectivity profiles was computed as the correlation of connectivity profile in a cortical surface node with the average connectivity profile of nodes at cortical distances ranging from 0 to 12 mm in other subjects' data. We similarly calculated within-subject PSFs between two resting state sessions. Figure 4 shows the slopes of connectivity profile PSFs in the ROIs (Fig. 4a) and the mean PSF across ROIs (Fig. 4b) in MSMAll-aligned data, in the common model connectome space, and within-subject. CHA increased the average slope of PSF across these ROIs, relative to anatomical alignment, from 0.015 to 0.066 (difference=0.051 [0.046, 0.054]). The PSF slopes in the common model

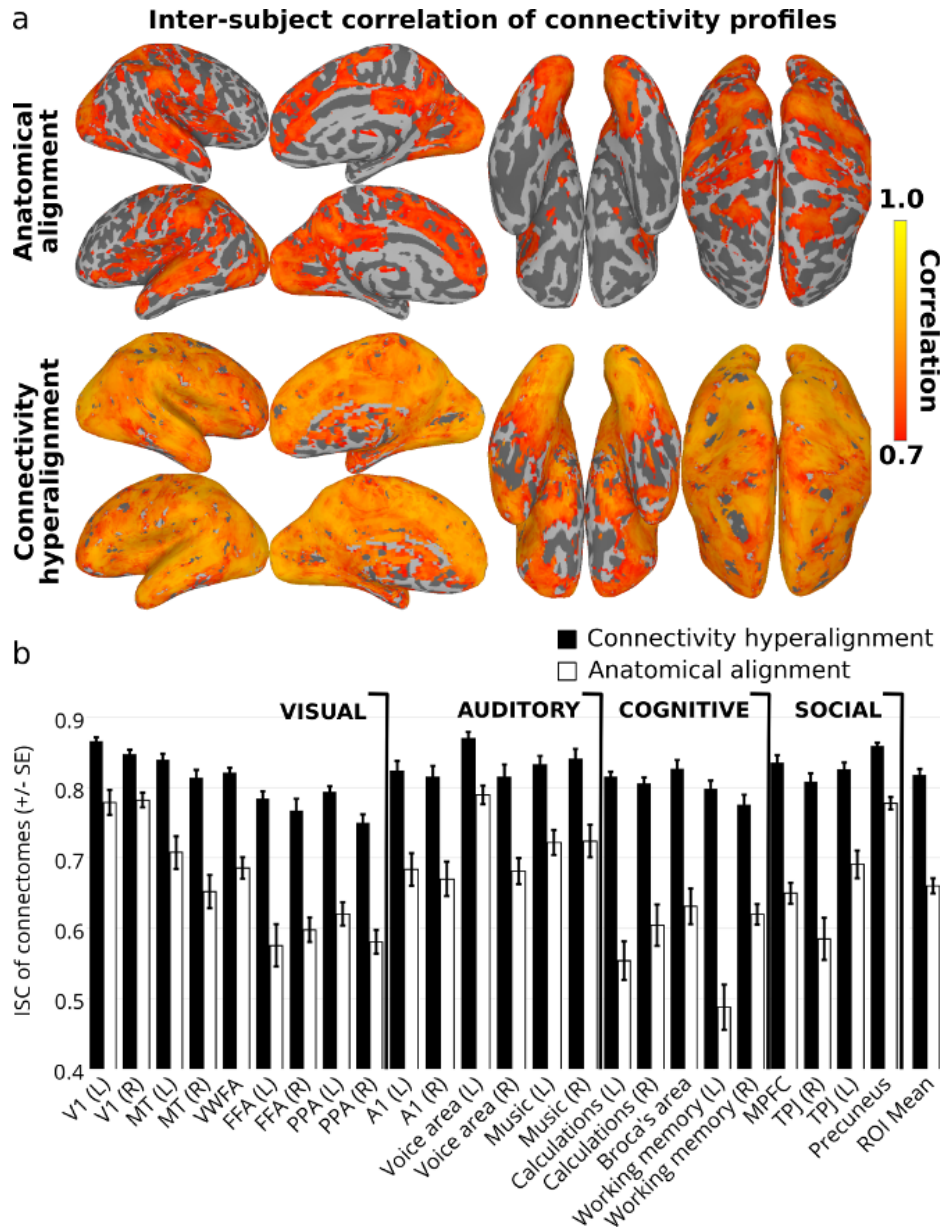


Figure 3. ISC of connectivity profiles calculated from movie data. (a) Average ISC of connectivity profiles in each surface node after anatomical alignment and in the common model connectome space. (b) ISC of connectivity profiles averaged within ROIs covering visual, auditory, cognitive, and social systems across the cortex and their mean comparing anatomical alignment and the common model connectome space. Bootstrapped testing showed significantly higher within-subject correlations and ISCs in the common model connectome space relative to anatomical alignment in all ROIs.

connectome space and within subject were not significantly different (slope = 0.070; difference = 0.003 [-0.001, 0.008]). The mean PSF across ROIs plotted as a function of distance from the cortical node clearly shows that CHA captures fine-scale variations in

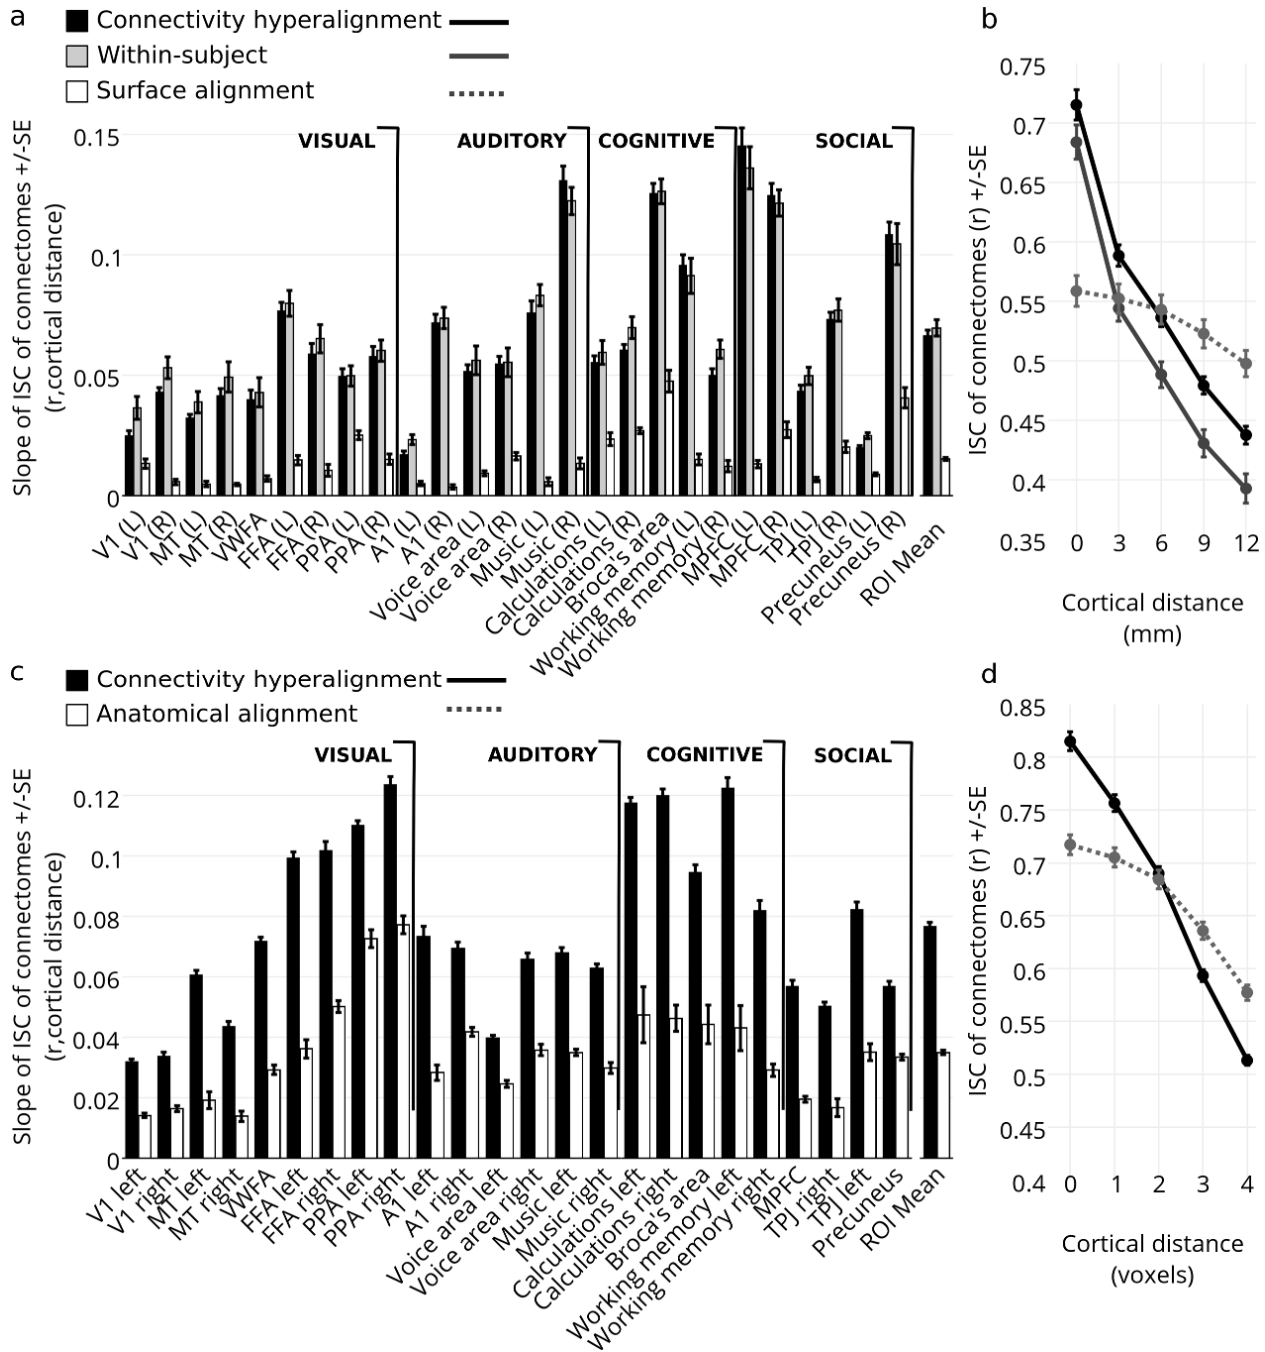


Figure 4. Spatial granularity of shared connectivity profiles. The intersubject point spread function (PSF) of connectivity profile correlations are computed as correlation between the connectivity profile for a cortical locus in one subject and the profiles of its spatial neighbors in other subjects at increasing distances from that locus (ISCs). For the HCP rsfMRI data, within-subject PSFs are computed as the correlation between the connectivity profile for a cortical locus from one rsfMRI session and the profiles of its spatial neighbors from a different rsfMRI session. Slope is estimated using a linear fit of these correlations at different distances as a function of distance in each of the ROIs (Left). Average PSF across all ROIs is plotted on the right. (a) shows PSF and slopes for HCP rsfMRI data and (b) shows the same for the movie data.




Figure 5. ISC of HCP task activation and contrast maps. Connectivity hyperalignment parameters derived from a session of resting state data were applied to the task maps and the correlation of these maps is computed between each subject and the average of others before and after hyperalignment. Hyperalignment improved correlations on average across all tasks and in all but two (labeled ns) individual task maps.

connectivity profiles for neighboring cortical nodes across subjects that are not captured after the MSM-All alignment provided by the HCP.

Movie data. Figure 4c shows the slopes of connectivity profile PSFs for the movie data in the ROIs, and Fig. 4d shows the mean PSF across ROIs. ISCs of connectivity profiles decrease as the cortical distance between features increases. Inter-subject cortical PSF slopes are much steeper in the common model connectome space, relative to anatomical alignment. Mean of connectivity profile PSF slopes was 0.077 in the common model connectome space and 0.035 after anatomical alignment (difference = 0.042 [0.039, 0.044]). See Supplementary Fig. 2 for comparison with RHA.

Inter-subject correlation of task activations

CHA based on resting state connectivity profiles increased ISC of task activation and contrast maps (Barch et al., 2013). Figure 5 shows the average correlation between

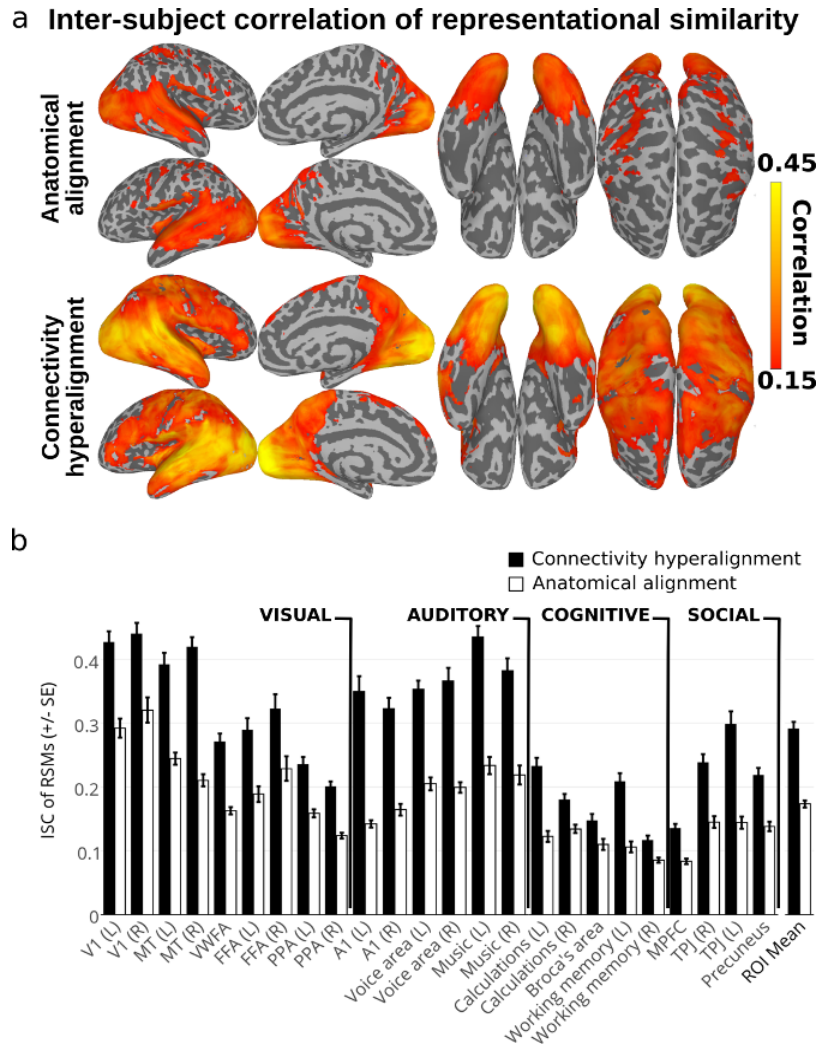


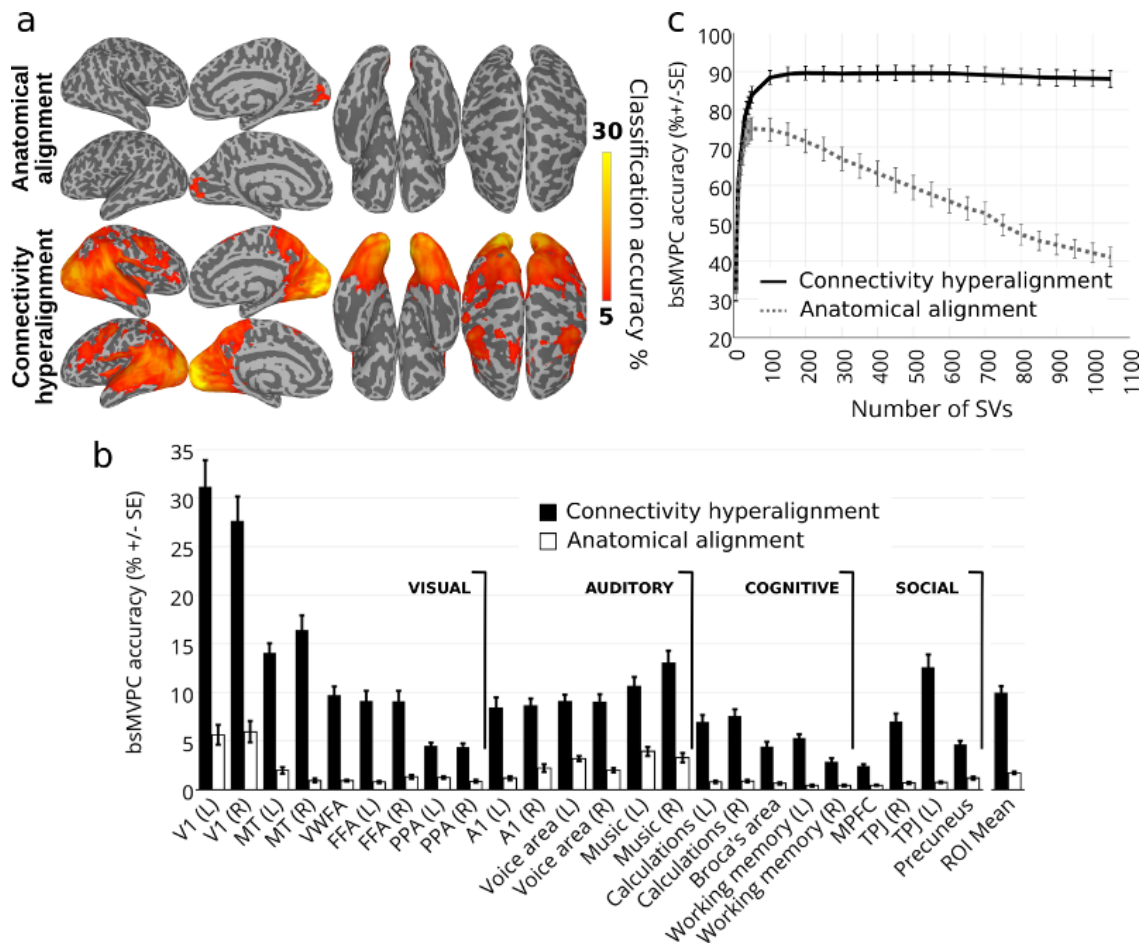
Figure 6. ISC of representational geometry of responses to movie time-points. (a) ISC of representational geometry in each voxel mapped onto cortical surfaces. (b) ISC of representational geometries in 24 ROIs covering visual, auditory, cognitive, and social systems across the cortex and their mean after anatomical alignment and in the common model connectome space. Bootstrapped testing showed significantly higher ISCs in the common model connectome space, as compared to anatomical alignment in all ROIs.

each subject's task map and the average task map of all other subjects. This was done for each of the 46 task maps independently. ISC of all but two (Face-Shapes and Body-Average) task maps were higher after CHA, and the average correlation across task maps increased from 0.58 to 0.65 (mean difference = 0.07 [0.06, 0.08]).

Intersubject correlation of representational geometry

Figure 6 shows the mean correlations of local representational geometry in each subject with the mean of other subjects' geometry after anatomical alignment and CHA.

Representational geometry is indexed as the similarity of patterns of response to different time-points in the movie. CHA increased the ISC of representational geometry in all of the ROIs (ROI mean ISC=0.291) relative to anatomical alignment (ISC=0.173) (difference=0.118 [0.103, 0.129]) (Fig. 6b). See Supplementary Fig. 3 for comparison with RHA.



Between-subject multivariate pattern classification of movie segments

Successful bsMVPC of movie segments reflects the degree to which patterns of activity representing the content of the movie stimulus are aligned across subjects. Figure 7a shows the mean classification accuracies in searchlights distributed across the whole cortex mapped to their centers. CHA greatly increased bsMVPC accuracies. Accuracies in the 24 ROIs (Fig. 7b), were nearly six times higher after CHA than after anatomical alignment (mean accuracies across ROIs were 9.93% and 1.73%, respectively). Classifying movie segments using the information distributed across the whole cortex also reflected a similar trend (Fig. 7c). Peak accuracy after CHA is 89.61% with more than 200 dimensions, whereas the peak accuracy with anatomical alignment is 74.85% with 50 dimensions (mean difference between CHA and anatomical alignment = 14.76% [12.17%, 17.12%]). See Supplementary Fig. 4 for comparison with RHA.

Discussion

We developed and present here a computational model of the human connectome that captures fine-scale local variation in connectivity profiles with basis functions that are shared across brains. We developed a new algorithm, connectivity hyperalignment (CHA), to derive this model. The common model connectome is a matrix of connectivity profiles for 59,412 cortical loci (HCP data) across 3852 connectivity targets. CHA calculates transformation matrices that project individual connectomes into the common model connectome space. Individual connectomes in the common model connectome space have markedly increased point-by-point similarity of connectivity profiles, as reflected in higher ISCs in the common model connectome space than in MSMAll- and anatomically-aligned data, and preserved local variability at a fine scale that matched the granularity of local variability observed within-subject. The intersubject PSFs of resting state connectivity profiles in the common model connectome were over four times steeper than PSFs in the MSMAll-aligned models. Whereas the PSF in the common model connectome shows significant change in connectivity profiles at a cortical distance of 3 mm, corresponding to a granularity of a single voxel, PSFs in the MSMAll- and anatomically-aligned models show negligible changes over distances of 6

mm, corresponding to a smoothness of three or more voxels. This local fine-scale variation was not incorporated into previous models of the connectome because its common structure across brains was obscured in anatomically-aligned and MSMAll-aligned data.

The common model connectome accounts for a large share of the variance of connectivity profiles across cortex, approximately 80% more than is accounted for by MSMAll-alignment of HCP data and 77% more than is accounted for by anatomical-alignment of movie data. In the HCP rsfMRI data, the common model connectome is a better predictor of an individual brain's connectome than is an independent rsfMRI data session from the same subject, due to the benefit of aggregating data over a large number of brains at the precision afforded by CHA. The results indicate that this added shared variance in the model is meaningful, affording substantially higher ISCs of representational geometry and dramatically higher bsMVPC for movie time segments. The results also found significantly higher ISCs of task maps and contrasts in the HCP data, even though these maps reflect coarse functional distinctions with functional topographies that are less dependent on fine-scale local variations.

Resting state fMRI connectivity is used to map the brain as a network of interconnected areas that form multi-area systems (Glasser et al., 2016; Gordon et al., 2016; Laumann et al., 2015; Smith, Vidaurre, et al., 2013; Sporns et al., 2004; Thomas Yeo et al., 2011; Wig et al., 2014). Homogeneous connectivity profiles within areas and sharp changes in connectivity profiles are used to define functional areas and their boundaries (Gordon et al., 2016; Wig et al., 2014). The existence of meaningful fine-scale local variation within areas is not inconsistent with such parcellation of cortex into areas. In fact, aligning cortical functional topography with CHA may sharpen the changes at borders, affording more precise delineation of their locations. Shared fine-scale variation within areas may reveal finer parcellations or how within-area functional topographies, such as retinotopy and orientation selectivity, are recapitulated and transformed from one area to another.

CHA aligns local high-dimensional connectivity spaces across subjects. Whole cortex CHA produces a common model of functional connectivity for all of cortex using a searchlight approach that we used previously for RHA (Guntupalli, Hanke, et al., 2016).

CHA, in contrast to RHA, does not rely on synchronized stimulation across subjects, as it finds basis functions for shared patterns of within-brain correlations of response profiles (connectivity profiles), rather than for synchronous response profiles. Consequently, CHA can be applied to fMRI data that are not synchronized across subjects. We show that CHA can derive valid common model connectomes based on rsfMRI data with asynchronous response profiles as well as on movie fMRI data with synchronous responses. We defined the connectivity profile of a cortical locus as its functional connectivity to a set of target regions distributed across the cortex. We index functional connectivity of a cortical locus as the correlation of its time-series with the time-series of the connectivity target regions. These target regions can be individual voxels or parcels covering the cortex. Our implementation of CHA here uses a set of searchlights that uniformly cover the cortex as target regions, but the algorithm can be applied to any other choice of targets and to other connectivity measures, including diffusion weighted imaging (DWI). To capture multivariate information within a target region, we used responses of the top three principal components instead of mean response in each region as our target.

CHA analyzes connectivity profiles in a searchlight or ROI as a set of vectors in which each cortical locus is a dimension and the connectivities of all ROI loci with each target is a vector in a connectivity space. These vectors reflect the “geometry” of functional connectivity for that ROI and are the basis for its hyperalignment. By defining the targets as sizeable regions or parcels, we can assume a coarse correspondence of targets across subjects. CHA aligns connectivity spaces across subjects by aligning these connectivity vectors across subjects with a high-dimensional transformation that is specific to each individual. This is analogous to how RHA aligns representational spaces across subjects by aligning response vectors during synchronized stimulation (Guntupalli, Wheeler, et al., 2016; Haxby et al., 2011). Synchronized stimulation allowed us to assume correspondence of responses at each time-point across subjects, and similarly, the top PCs for sizeable target regions, which are aligned to have similar connectivities with other target regions, allows us to assume correspondence of connectivity vectors for each target PC across subjects.

Successful alignment of fine-scale patterns of neural activity using CHA, as shown by increased ISCs of representational geometry and bsMVPC of movie time segments, suggests that the response tuning function of a cortical locus is closely related to its connectivity profile to the rest of the cortex. Although we predicted this relationship, it was not clear that it need exist. It was possible that different cortical loci in a given brain area have different response tuning functions but communicate with the same constellation of areas in other parts of the brain (Saygin et al., 2012, 2016). In other words, the connectivity of cortical loci within an area could be driven by the portion of response tuning variance that is shared across those loci, yielding a homogeneous connectivity profile, or by the portion of response tuning variance that varies across loci, yielding fine-scale local variation in connectivity profiles. Our results show the fine-scale local variation in response tunings is strongly associated with fine-scale local variation in connectivity profiles.

Recent studies have explored another avenue to extract more information from connectivity in the temporal domain. Most studies, including this one, use a single consensus connectivity computed across the whole session of an experiment or resting state. But connectivity between regions varies across time and tasks (Allen et al., 2014; Geerligs, Rubinov, Cam-CAN, & Henson, 2015; Hutchison et al., 2013; Shine et al., 2016). CHA could be extended to multiple connectomes extracted from different tasks or states in the same subject. Using connectivity vectors from a set of brain states to derive hyperalignment, rather than assuming a single stationary brain state, could potentially improve the generalization of alignment to other experiments.

Using connectivity profiles for hyperalignment offers important practical advantages. Since connectivity is computed within subject, synchronized responses to the same stimuli are no longer required for hyperalignment. Consequently, connectivity hyperalignment can be based on data from experiments probing planning and decision making tasks, which might unfold at different speeds across subjects, more free-form tasks such as social interactions, and data from task-free resting state sessions. Resting state data in particular are simple to acquire, already ubiquitous and acquired routinely by many labs for participants in varied experiments, and can be acquired in patient populations in which demanding tasks may affect compliance. Many large

shared datasets acquired on normal subjects, such as the HCP (Van Essen et al., 2013) and ADHD-200 (“The ADHD-200 Consortium,” 2012), and on patient populations (ADHD-200, ADNI-GO (Jack Jr. et al., 2010), UCLA Neuropsychiatric Phenomics (Poldrack et al., 2016)) already contain resting state data.

Acquisition of resting state data is now widespread in fMRI research, but such data may not be optimal for building a common model of high-dimensional representational spaces. rsfMRI datasets typically are limited in size – usually five to ten, at most twenty, minutes – and more data afford better estimates of the large number of parameters in transformation matrices for increased general validity (Guntupalli, Hanke, et al., 2016). Furthermore, a recent study has suggested that connectivity profiles during movie viewing, as compared to connectivity profiles during the resting state, may be more reliable across sessions and may enhance measures of individual differences in a normal population (Vanderwal et al., 2016). Further work will be necessary to quantify the precision and general validity afforded by different sizes and types of fMRI datasets for building a common model connectome using hyperalignment.

In summary, we present here a common model of the human connectome that captures the local fine-scale structure of functional connectivity in shared basis functions. The model is derived with a new hyperalignment algorithm, CHA, that enables application of hyperalignment to a greatly expanded array of experimental data. The common model connectome captures the functional architecture of human cortex at a fine spatial scale. Application of CHA to resting state rsfMRI data afforded the discovery of a shared fine-scale spatial granularity in local variation of connectivity profiles that was previously unrecognized. Capturing this fine-scale structure in the common model connectome can afford better models of the areal structure of cortical topography and study of how within-area topographies are recapitulated and transformed between areas. CHA can be used to project data from large public databases, such as the HCP, into a common model space to align functional connectivities and functional topographies at a fine spatial scale and, thereby, increase the power of these databases for investigation of neural representation, connectivity, individual differences, and parcellation of cortex into functional areas.

Materials and Methods

Resting state data: Human Connectome Project

In the HCP database (Van Essen et al., 2013), we found unrelated subjects of age ≤ 35 with at least four resting state scans, yielding a list of 64 subjects. We chose the first 20 of these subjects in the sorted order of subject IDs for our analysis.

For each subject, we used their cortical surfaces and fMRI data aligned to the group using MSMAll (Robinson et al., 2014) with 32K nodes in each hemisphere as provided by the HCP. We used data from one resting state session (Smith, Beckmann, et al., 2013) (“rfMRI_REST1_LR”) to derive CHA parameters and validated it on a different resting state session (“rfMRI_REST2_LR”), and task fMRI sessions (Barch et al., 2013) (EMOTION, GAMBLING, LANGUAGE, MOTOR, RELATIONAL, SOCIAL, and WM). Resting state data were acquired for 1200 TRs with a TR of 0.720s in each session (total time=14 min 33 s). The data used to derive the CHA parameters and common model and the resting state data used for validation tests used the same phase-encoding direction (LR). We used a single session of rsfMRI for alignment to mimic a typical resting state data acquisition which usually varies from 10-20 mins of scanning. See Smith et al. (2013) for more details about the acquisition and preprocessing pipelines.

We masked the data to include only the left and right cortices (Cortex_Left and Cortex_Right), removing all the non-zero nodes that correspond to the medial subcortical regions, resulting in 59,412 nodes across both hemispheres. We generated a coarser surface grid to match the ico8 surface in SUMA (Saad, Reynolds, Argall, Japee, & Cox, 2004) with 1284 nodes (10.7 mm spacing between nodes). We found the closest matching nodes on the 32K surface to the nodes on the ico8 surface, and used those as our searchlight centers for connectivity target searchlights. For validation of task fMRI, we used all of the maps provided by the HCP after removing redundancies (such as FACE-AVG and AVG-FACE), which resulted in 46 maps.

Movie data: Raiders of the Lost Ark

We scanned 11 healthy young right-handed participants (4 females; Mean age: 24.6+/- 3.7 years) with no history of neurological or psychiatric illness during movie viewing. All had normal or corrected-to-normal vision. Informed consent was collected in accordance with the procedures set by the local Committee for the Protection of Human Subjects. Subjects were paid for their participation.

Stimuli and design Stimuli consisted of the full-length feature movie – “Raiders of the Lost Ark” – divided into eight parts of approximately 14 min 20 s duration. Video was projected onto a rear projection screen with an LCD projector which the subject viewed through a mirror on the head coil. The video image subtended a visual angle of approximately 22.7° horizontally and 17° vertically. Audio was presented through MR Confon’s MRI-compatible headphones. Subjects were instructed to pay attention to the movie and enjoy. See Guntupalli et al. (2016) for details.

fMRI protocol Subjects were scanned in a Philips Intera Achieva 3T scanner with an 8 channel head coil at the Dartmouth Brain Imaging Center. T1-weighted anatomical scans were acquired at the end of each session (MPRAGE, TR=9.85 s, TE=4.53 s, flip angle=8°, 256 x 256 matrix, FOV=240 mm, 160 1 mm thick sagittal slices). The voxel resolution was 0.938 mm x 0.938 mm x 1.0 mm. Functional scans of the whole brain were acquired with an echo planar imaging sequence (TR=2.5 s, TE=35 ms, flip angle=90°, 80 x 80 matrix, FOV=240 mm x 240 mm) every 2.5 s with whole brain coverage (41 3 mm thick interleaved axial slices, giving isotropic 3 mm x 3 mm x 3 mm voxels). We acquired a total of 2718 functional scans with 1350 TRs in four runs during the first session and 1368 TRs in four runs during the second session.

fMRI data preprocessing fMRI data were preprocessed using AFNI software (Cox, 1996) (<http://afni.nimh.nih.gov>). Functional data were first corrected for the order of slice acquisition and head motion by aligning to the last volume of the last functional run. Any spikes in the data were removed using 3dDespike in AFNI. Data were then filtered using 3dBandpass in AFNI to remove any temporal signal variation slower than 0.00667 Hz, faster than 0.1 Hz, and that correlated with the whole brain average signal or the head movement parameters. Each subject’s anatomical volume was first aligned to the motion corrected average EPI volume and then to the MNI 152 brain template in AFNI.

Functional EPI BOLD data were then aligned to the MNI 152 brain template using nearest neighbor resampling by applying the transformation derived from the alignment of the anatomical volume to the template. Functional data were spatially smoothed with a 4 mm full-width at half-max Gaussian filter (see Guntupalli et al. (2016) for the details of how smoothing affects hyperalignment and subsequent validations). Data acquired during the overlapping movie segments were discarded resulting in a total of 2662 TRs with 1326 TRs in the first session and 1336 TRs in the second session. We derived a gray matter mask by segmenting the MNI_avg152T1 brain provided in AFNI and removing any voxel that was outside the cortical surface by more than twice the thickness of the gray matter at each surface node. It included 54,034 3 mm isotropic voxels across both hemispheres. We used this mask for all subsequent analyses of all subjects.

Connectivity Hyperalignment

CHA uses functional connectivity profiles to compute transformation matrices and derive a common model of the human connectome. Connectivity profiles consist of the correlations of the response profiles of cortical loci in an ROI or searchlight with the response profiles of targets distributed across the cortex. CHA finds a transformation that minimizes the distances between the connectivity vectors for each target in one subject with the vectors for matched targets in other subjects. In our implementations of hyperalignment, we have used the Procrustes transformation (Schönemann, 1966) that derives an orthogonal matrix for the optimal improper rotation of one subject's connectivity vectors to the mean of others'. Other hyperalignment algorithms have used alternatives for calculating the transformation matrices, such as regularized canonical correlation and probabilistic estimation (Chen et al., 2015; Xu, Lorberty, Ramadge, Guntupalli, & Haxby, 2012). The RHA algorithm uses the same procedure but minimizes the distances between response vectors for the same time-points or stimuli across subjects (Guntupalli, Hanke, et al., 2016; Haxby et al., 2011). A common model connectome for an ROI is derived iteratively by first aligning a subject's connectivity vectors to those of another subject, then aligning a third subject's vectors to the mean of those subjects, and then aligning each additional subject to the mean of the subjects

that have already been hyperaligned. On a second pass, each subject's connectivity vectors are aligned to the mean vectors for all-but-that subject calculated in the first pass. The mean connectivity vectors that have been re-hyperaligned in the second pass are the common model connectome. It is a matrix of connectivities between cortical loci (tagged to the reference subject's anatomy) as columns and connectivity targets as rows. Thus, each cortical locus in the common model connectome has a connectivity profile that is common across brains. The cortical loci in an individual brain are transformed into the common model connectome space with a matrix which has the common model connectome's loci as columns and the individual brain's loci as rows. Cortical loci in the common model connectome space can be transformed into topographies in an individual brain, using the columns in these transformation matrices, that are individual-specific topographic basis functions that capture how the common connectivity profiles are expressed in individual brains as multiplexed topographies.

Each connectivity target is a surface-searchlight of 13 mm radius centered on a node of a coarse surface with a total of 1284 covering both hemispheres. Thus, neighboring searchlights for connectivity are overlapping. To capture more variance of response profiles in each target, we used 3 principal components to represent the response profiles in that searchlight instead of using the mean of response profiles of all features as is typically done for computing connectomes. CHA rests on the assumption that the top target components reflect the same connectivity profiles across subjects (similar to the correspondence of time points in a movie across subjects for RHA). Therefore, we used a 2 step process to make sure these components are aligned across subjects in terms of their connectivity to other targets.

In the first step, we hyperaligned the features (voxels/surface nodes) in each target searchlight across subjects using the correlation of responses between each feature in that target and the mean response profiles of all the features in each of the target searchlights. For each target searchlight with N_s features (voxels or surface nodes), we computed a $1284 \times N_s$ correlation matrix for each subject, and performed hyperalignment of these matrices to align those N_s features. This aligns the features in a particular target searchlight across subjects based on their connectomes – feature connectivity profiles to other target searchlight means. We averaged these aligned

connectomes across subjects and did a singular value decomposition (SVD) to compute the top three components that explained the most shared variance. Each component is a weighted sum of features in that target searchlight for each subject such that its connectivity profile is aligned across subjects. Each individual subject's responses for the top three components were then used as target response profiles for CHA. This step gave us $1284 \times 3 = 3852$ target response profiles in each subject's cortex.

In the second step, we hyperaligned the larger searchlights centered on a dense grid. For the HCP data, searchlights were centered on each of the 59,412 nodes and we selected all surface nodes within each searchlight of 20 mm radius. The mean number of surface nodes in these large searchlights was 337. For the movie data, we used a surface-searchlight (Oosterhof, Wiestler, Downing, & Diedrichsen, 2011) approach as described previously in Guntupalli et al. (2016) with searchlights centered on each of 20,484 cortical nodes with 2.9 mm average spacing between the nodes. We performed voxel selection within these searchlights to reduce the contribution from noisy or non-gray matter voxels that were included due to this dilation (Guntupalli, Hanke, et al., 2016). The mean number of selected voxels in these searchlights was 235. We then computed the connectivities profile between each surface node (or voxel for the movie data) in each searchlight of 20 mm radius to the 3852 response profiles for the connectivity targets. The 3852 connectivity vectors in an ROI or searchlight, one for each connectivity target, are used as the basis for CHA of that ROI or searchlight, producing transformation matrices that are then aggregated into a single, whole-cortex transformation matrix using the algorithm for whole-cortex searchlight RHA described in Guntupalli et al. (2016).

Whole cortex hyperalignment aggregates the orthogonal matrices from hyperalignment of each searchlight by adding all the weights estimated for a given node-to-feature mapping in surface-searchlights centered on different cortical nodes, resulting in an $N \times N$ transformation matrix for each subject, where N is 59,412 surface nodes for HCP data and 54,034 voxels in the gray matter mask for movie data. The reverse mapping matrix is the transpose of this matrix for that subject (Fig. 1c). Whole cortex hyperalignment transformation matrices for an individual were applied to any new dataset in that same individual by first normalizing the data in each voxel by z-scoring

and multiplying that normalized dataset with the transformation matrix. The same procedure can be followed to apply reverse mapping (Guntupalli, Hanke, et al., 2016). CHA of HCP data was based on one session of resting state data (~15 min, 1200 TRs) and a second session of resting state data was used for validation tests. CHA of movie data was based on one half of the movie data (~55 min, ~1300 TRs) and a second half of the movie data was used for validation tests.

Validation tests and statistical analyses

Functional ROIs. In addition to analyzing the results of validation tests in each feature or searchlight across the whole cortex, we also examined the results of validation tests in 24 functional ROIs associated with different sensory, perceptual, and cognitive functions to assess the general validity of the common model (Guntupalli, Hanke, et al., 2016). We searched for terms and cortical areas implicated in visual, auditory, cognitive, and social functions in NeuroSynth (Yarkoni et al., 2011) and took the coordinates for the peak location associated with the respective term (Supplementary Table 1). For validation testing using the HCP dataset, we found the closest surface node corresponding to each peak locus and used a surface searchlight with a 10 mm radius around that surface node as the functional ROI. Functional ROIs that were medial and encompassing both hemispheres in the volume space were split into left and right ROIs in the surface space resulting in 26 ROIs for tests on the HCP data. For validation testing using the movie dataset, we used volume searchlights centered around those peak loci with a radius of 3 voxels as our functional ROIs. For analyses of ISCs and PSFs of connectivity profiles in functional ROIs, we calculated the mean ISC or PSF across all cortical loci within the ROI searchlights (Figs. 2b, 3b, and 4).

Statistics. We used bootstrapping to test for significance of the effects by sampling subjects 10,000 times to compute 95% CIs using BootES (Kirby & Gerlanc, 2013). We did this for each ROI and for the mean of all ROIs separately. We used the same bootstrapping procedure for all validation tests unless specified otherwise.

Intersubject correlation (ISC) of connectivity profiles. Connectivity profile of a feature (or a cortical node) was defined as the correlation of its time-series to that of average time-series in a set of target regions covering the whole cortex. For HCP data, targets were surface searchlights with a radius of 10 mm with centers uniformly distributed on a coarse surface grid (ico8). There were 1284 targets in total. This gave vector of 1284 values for each cortical locus. ISCs of connectivity profiles were computed between each subject's connectivity profile and the average connectivity profile of all other subjects in each cortical locus. ISCs of connectivity profiles were computed for session REST2 before and after CHA, while the session REST1 was used for deriving CHA.

We also computed within-subject between-session (REST1 and REST2) correlation of connectivity profiles. This also was performed after CHA to account for any filtering effects but to a single common space as used for our main analyses. For the movie data, ISCs of connectivity profiles were computed within each movie half separately and before and after CHA based on the other half of the movie. Each target is a surface disc searchlight of radius 13 mm with centers uniformly distributed on a coarse surface grid (ico8). Correlation values were Fisher transformed before averaging across both halves of the movie in each voxel. These were then averaged across all subjects and inverse Fisher transformed before mapping onto the cortical surface for visualization.

Spatial point spread function. To investigate the spatial granularity of representation, we computed a spatial point spread function (PSF) of ISCs of connectivity profiles. We computed the correlation of connectivity profiles in each cortical locus (voxel or surface node) with the average connectivity profiles of cortical loci at varying cortical distances in other subjects' data. To account for the effect of filtering, we did this analyses with data after hyperalignment with each subject aligned to a different common space as proxy for anatomical alignment and after hyperalignment with each subject aligned to the same reference subject for hyperalignment (Guntupalli, Hanke, et al., 2016). We computed similar PSFs for connectivity profiles within-subject between-sessions

(REST1 and REST2). This was also performed after hyperalignment to account for any filtering effects but to a single common space as used for our main analyses.

Inter-subject correlation of task activations. We evaluated if CHA derived from the resting state data improves alignment of task maps across subjects. We applied the CHA transformations derived from one session of the resting state data to the task maps for the respective subjects provided by the HCP (Barch et al., 2013). We used 46 maps by keeping only one map for each contrast. We then computed the spatial correlation of these maps across subjects in the MSM-All aligned data and hyperaligned data for each task map.

ISC of representational geometry. ISCs of similarity structures were computed within each movie half separately using a searchlight of 3 voxel radius. In each searchlight, similarity structure was computed as a matrix of correlation coefficients between responses at every pair of time-points from that movie half for each subject. The flattened upper triangle of this matrix excluding the diagonal was extracted as the profile of representational geometry at each searchlight for each subject. ISC of representational geometry in each searchlight was computed as the correlation between each subject's similarity profile and the average of all other subjects' similarity profiles for that searchlight. Correlation values were Fisher transformed before averaging across both movie halves in each voxel. These were then averaged across all subject-average pairs and inverse Fisher transformed before mapping onto the cortical surface for visualization. The same steps were performed to compute inter-subject correlation of representational similarity before and after hyperalignment.

Between-subject multivariate pattern classification (bsMVPC). bsMVPC of 15 s movie time segments (6 TRs) was computed within each movie half separately using a searchlight of 3 voxel radius. bsMVPC was performed using a one-nearest neighbor classifier based on correlation distance (Guntupalli, Hanke, et al., 2016; Haxby et al., 2011). Each 15 s (6TR) sequence of brain data for an individual was compared to average 15 s sequences for all other subjects for that sequence and other sequences

using a sliding time window (chance < 0.1%). Classification accuracies at each voxel were averaged across both halves in each subject before mapping the subject mean onto the cortical surface for visualization. bsMVPC of movie segments using voxels from the whole cortex was performed after projecting the test datasets into the singular vector (SV) space derived from the the same training data as used to derive hyperalignment. These SVs were sorted in the descending order of their singular values. bsMVPC was performed for multiple sets of top singular vectors in each movie half and the accuracies for the two halves of the movie were averaged for each subject at each set size.

We implemented our methods and ran our analyses in PyMVPA (Hanke et al., 2009) (<http://www.pymvpa.org>) unless otherwise specified. All preprocessing and analyses were carried out on a 64-bit Debian 7.0 (wheezy) system with additional software from NeuroDebian (Halchenko & Hanke, 2012) (<http://neuro.debian.net>).

Acknowledgments

We would like to thank Yaroslav O. Halchenko, Michael Hanke, Feilong Ma, Sam Nastase, and Nick O. Oosterhof for discussion and software support. This work was supported by grants from the National Institute of Mental Health (R01MH075706) and the National Science Foundation (NSF1129764 and NSF1607845). Resting state and task fMRI data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

Competing interests

We declare no competing financial or non-financial interests.

References

- Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T., & Calhoun, V. D. (2014). Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cerebral Cortex*, *24*(3), 663–676. <https://doi.org/10.1093/cercor/bhs352>
- Barch, D. M., Burgess, G. C., Harms, M. P., Petersen, S. E., Schlaggar, B. L., Corbetta, M., ... Van Essen, D. C. (2013). Function in the human connectome: Task-fMRI and individual differences in behavior. *NeuroImage*, *80*, 169–189. <https://doi.org/10.1016/j.neuroimage.2013.05.033>
- Biswal, B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine: Official Journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, *34*(4), 537–541.
- Chen, P.-H., Chen, J., Yeshurun, Y., Hasson, U., Haxby, J. V., & Ramadge, P. J. (2015). A Reduced-dimension fMRI Shared Response Model. In *Proceedings of the 28th International Conference on Neural Information Processing Systems* (pp. 460–468). Cambridge, MA, USA: MIT Press. Retrieved from <http://dl.acm.org/citation.cfm?id=2969239.2969291>
- Cohen, A. L., Fair, D. A., Dosenbach, N. U. F., Miezin, F. M., Dierker, D., Van Essen, D. C., ... Petersen, S. E. (2008). Defining functional areas in individual human brains using resting functional connectivity MRI. *NeuroImage*, *41*(1), 45–57. <https://doi.org/10.1016/j.neuroimage.2008.01.066>
- Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research, an International Journal*, *29*(3), 162–173.
- Geerligs, L., Rubinov, M., Cam-CAN, & Henson, R. N. (2015). State and Trait Components of Functional Connectivity: Individual Differences Vary with Mental State. *Journal of Neuroscience*, *35*(41), 13949–13961. <https://doi.org/10.1523/JNEUROSCI.1324-15.2015>
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., ... Van

- Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615), 171–178. <https://doi.org/10.1038/nature18933>
- Gordon, E. M., Laumann, T. O., Adeyemo, B., Huckins, J. F., Kelley, W. M., & Petersen, S. E. (2016). Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cerebral Cortex*, 26(1), 288–303. <https://doi.org/10.1093/cercor/bhu239>
- Gordon, E. M., Laumann, T. O., Adeyemo, B., & Petersen, S. E. (2015). Individual Variability of the System-Level Organization of the Human Brain. *Cerebral Cortex (New York, N.Y.: 1991)*. <https://doi.org/10.1093/cercor/bhv239>
- Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews Neuroscience*, 15(8), 536–548. <https://doi.org/10.1038/nrn3747>
- Guntupalli, J. S., Hanke, M., Halchenko, Y. O., Connolly, A. C., Ramadge, P. J., & Haxby, J. V. (2016). A Model of Representational Spaces in Human Cortex. *Cerebral Cortex*, 26(6), 2919–2934. <https://doi.org/10.1093/cercor/bhw068>
- Guntupalli, J. S., Wheeler, K. G., & Gobbini, M. I. (2016). Disentangling the Representation of Identity from Head View Along the Human Face Processing Pathway. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/bhw344>
- Haak, K. V., Marquand, A. F., & Beckmann, C. F. (2016). Connectopic mapping with resting-state fMRI. *arXiv:1602.07100 [Q-Bio]*. Retrieved from <http://arxiv.org/abs/1602.07100>
- Halchenko, Y. O., & Hanke, M. (2012). Open is Not Enough. Let's Take the Next Step: An Integrated, Community-Driven Computing Platform for Neuroscience. *Frontiers in Neuroinformatics*, 6. <https://doi.org/10.3389/fninf.2012.00022>
- Hanke, M., Halchenko, Y. O., Sederberg, P. B., Hanson, S. J., Haxby, J. V., & Pollmann, S. (2009). PyMVPA: A Python toolbox for multivariate pattern analysis of fMRI data. *Neuroinformatics*, 7(1), 37–53. <https://doi.org/10.1007/s12021-008-9041-y>

- Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding Neural Representational Spaces Using Multivariate Pattern Analysis. *Annual Review of Neuroscience*, 37(1), 435–456. <https://doi.org/10.1146/annurev-neuro-062012-170325>
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430. <https://doi.org/10.1126/science.1063736>
- Haxby, J. V., Guntupalli, J. S., Connolly, A. C., Halchenko, Y. O., Conroy, B. R., Gobbini, M. I., ... Ramadge, P. J. (2011). A Common, High-Dimensional Model of the Representational Space in Human Ventral Temporal Cortex. *Neuron*, 72(2), 404–416. <https://doi.org/10.1016/j.neuron.2011.08.026>
- Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., ... Chang, C. (2013). Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*, 80, 360–378. <https://doi.org/10.1016/j.neuroimage.2013.05.079>
- Jack Jr., C. R., Bernstein, M. A., Borowski, B. J., Gunter, J. L., Fox, N. C., Thompson, P. M., ... Weiner, M. W. (2010). Update on the Magnetic Resonance Imaging core of the Alzheimer's Disease Neuroimaging Initiative. *Alzheimer's & Dementia*, 6(3), 212–220. <https://doi.org/10.1016/j.jalz.2010.03.004>
- Jbabdi, S., Sotiropoulos, S. N., & Behrens, T. E. (2013). The topographic connectome. *Current Opinion in Neurobiology*, 23(2), 207–215. <https://doi.org/10.1016/j.conb.2012.12.004>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *The Journal of Neuroscience*, 17(11), 4302–4311.
- Kay, K. N., Naselaris, T., Prenger, R. J., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, 452(7185), 352–355. <https://doi.org/10.1038/nature06713>
- Kirby, K. N., & Gerlanc, D. (2013). BootES: An R package for bootstrap confidence intervals on effect sizes. *Behavior Research Methods*, 45(4), 905–927. <https://doi.org/10.3758/s13428-013->

0330-5

Laumann, T. O., Gordon, E. M., Adeyemo, B., Snyder, A. Z., Joo, S. J., Chen, M.-Y., ... Petersen, S. E. (2015). Functional System and Areal Organization of a Highly Sampled Individual Human Brain. *Neuron*, *87*(3), 657–670. <https://doi.org/10.1016/j.neuron.2015.06.037>

Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, *10*(9), 424–430. <https://doi.org/10.1016/j.tics.2006.07.005>

Oosterhof, N. N., Wiestler, T., Downing, P. E., & Diedrichsen, J. (2011). A comparison of volume-based and surface-based multi-voxel pattern analysis. *NeuroImage*, *56*(2), 593–600. <https://doi.org/10.1016/j.neuroimage.2010.04.270>

Poldrack, R. A., Congdon, E., Triplett, W., Gorgolewski, K. J., Karlsgodt, K. H., Mumford, J. A., ... Bilder, R. M. (2016). A phenome-wide examination of neural and cognitive function. *Scientific Data*, *3*. <https://doi.org/10.1038/sdata.2016.110>

Robinson, E. C., Jbabdi, S., Glasser, M. F., Andersson, J., Burgess, G. C., Harms, M. P., ... Jenkinson, M. (2014). MSM: A new flexible framework for Multimodal Surface Matching. *NeuroImage*, *100*, 414–426. <https://doi.org/10.1016/j.neuroimage.2014.05.069>

Saad, Z. S., Reynolds, R. C., Argall, B., Japee, S., & Cox, R. W. (2004). SUMA: an interface for surface-based intra- and inter-subject analysis with AFNI. In *IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2004* (p. 1510–1513 Vol. 2). <https://doi.org/10.1109/ISBI.2004.1398837>

Saygin, Z. M., Osher, D. E., Koldewyn, K., Reynolds, G., Gabrieli, J. D. E., & Saxe, R. R. (2012). Anatomical connectivity patterns predict face selectivity in the fusiform gyrus. *Nature Neuroscience*, *15*(2), 321–327. <https://doi.org/10.1038/nn.3001>

Saygin, Z. M., Osher, D. E., Norton, E. S., Youssoufian, D. A., Beach, S. D., Feather, J., ... Kanwisher, N. (2016). Connectivity precedes function in the development of the visual word form area. *Nature Neuroscience*, *advance online publication*. <https://doi.org/10.1038/nn.4354>

Schönemann, P. H. (1966). A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1), 1–10. <https://doi.org/10.1007/BF02289451>

Shine, J. M., Bissett, P. G., Bell, P. T., Koyejo, O., Balsters, J. H., Gorgolewski, K. J., ... Poldrack, R. A. (2016). The Dynamics of Functional Brain Networks: Integrated Network States during Cognitive Task Performance. *Neuron*, 92(2), 544–554. <https://doi.org/10.1016/j.neuron.2016.09.018>

Smith, S. M., Beckmann, C. F., Andersson, J., Auerbach, E. J., Bijsterbosch, J., Douaud, G., ... Glasser, M. F. (2013). Resting-state fMRI in the Human Connectome Project. *NeuroImage*, 80, 144–168. <https://doi.org/10.1016/j.neuroimage.2013.05.039>

Smith, S. M., Vidaurre, D., Beckmann, C. F., Glasser, M. F., Jenkinson, M., Miller, K. L., ... Van Essen, D. C. (2013). Functional connectomics from resting-state fMRI. *Trends in Cognitive Sciences*, 17(12), 666–682. <https://doi.org/10.1016/j.tics.2013.09.016>

Sporns, O., Chialvo, D. R., Kaiser, M., & Hilgetag, C. C. (2004). Organization, development and function of complex brain networks. *Trends in Cognitive Sciences*, 8(9), 418–425. <https://doi.org/10.1016/j.tics.2004.07.008>

The ADHD-200 Consortium: A Model to Advance the Translational Potential of Neuroimaging in Clinical Neuroscience. (2012). *Frontiers in Systems Neuroscience*, 6. <https://doi.org/10.3389/fnsys.2012.00062>

Thomas Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., ... Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3), 1125–1165. <https://doi.org/10.1152/jn.00338.2011>

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., & Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80, 62–79. <https://doi.org/10.1016/j.neuroimage.2013.05.041>

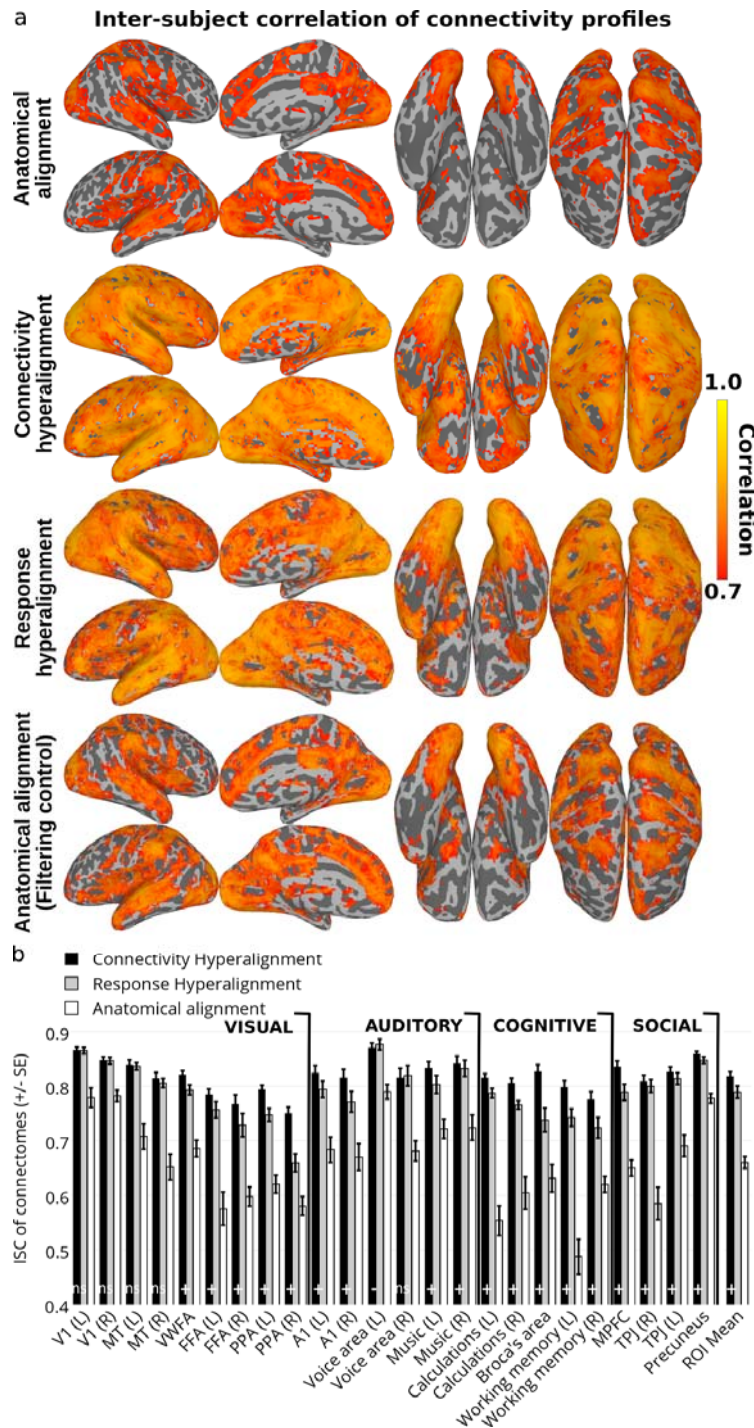
Vanderwal, T., Eilbott, J., Finn, E. S., Craddock, R. C., Turnbull, A. G., & Castellanos, F. X. (2016). Individual differences in functional connectivity during naturalistic viewing conditions. *bioRxiv*, 84665. <https://doi.org/10.1101/084665>

Wig, G. S., Laumann, T. O., Cohen, A. L., Power, J. D., Nelson, S. M., Glasser, M. F., ... Petersen, S. E. (2014). Parcellating an Individual Subject's Cortical and Subcortical Brain Structures Using Snowball Sampling of Resting-State Correlations. *Cerebral Cortex*, *24*(8), 2036–2054. <https://doi.org/10.1093/cercor/bht056>

Xu, H., Lorbett, A., Ramadge, P. J., Guntupalli, J. S., & Haxby, J. V. (2012). Regularized hyperalignment of multi-set fMRI data. In *2012 IEEE Statistical Signal Processing Workshop (SSP)* (pp. 229–232). IEEE.

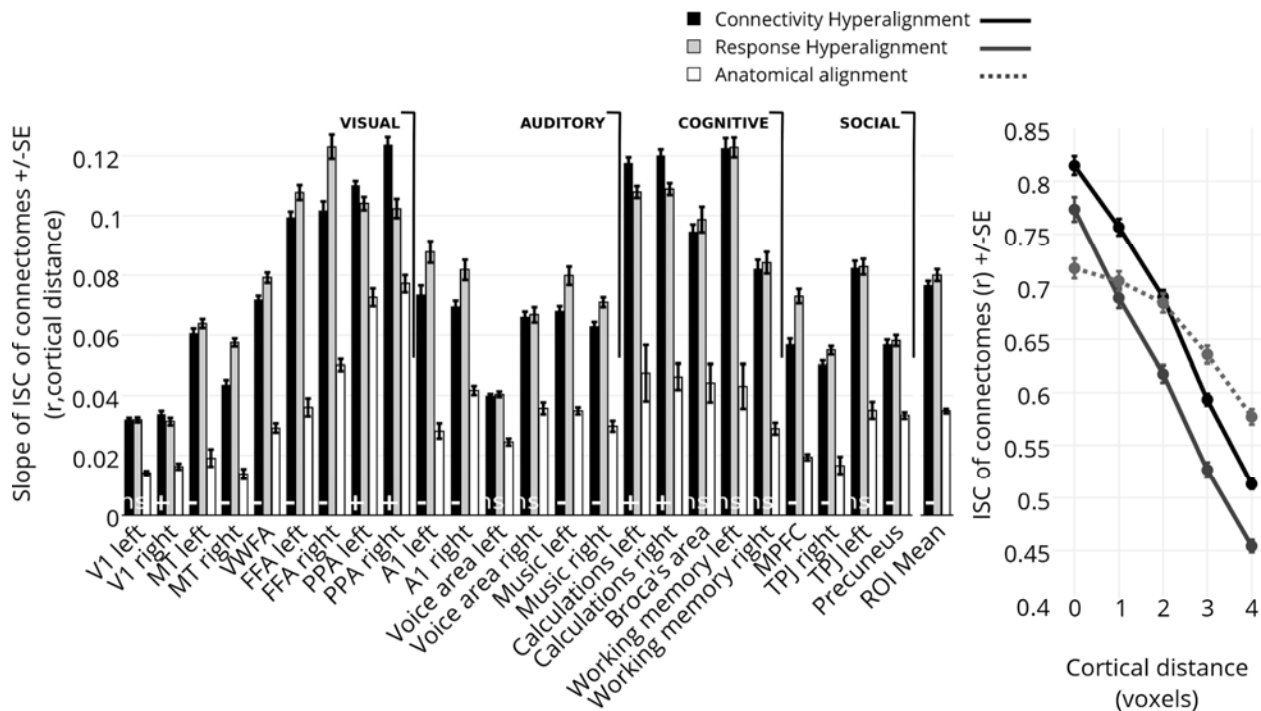
Yarkoni, T., Poldrack, R. A., Nichols, T. E., Essen, D. C. V., & Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, *8*(8), 665–670. <https://doi.org/10.1038/nmeth.1635>

Supplementary Figures

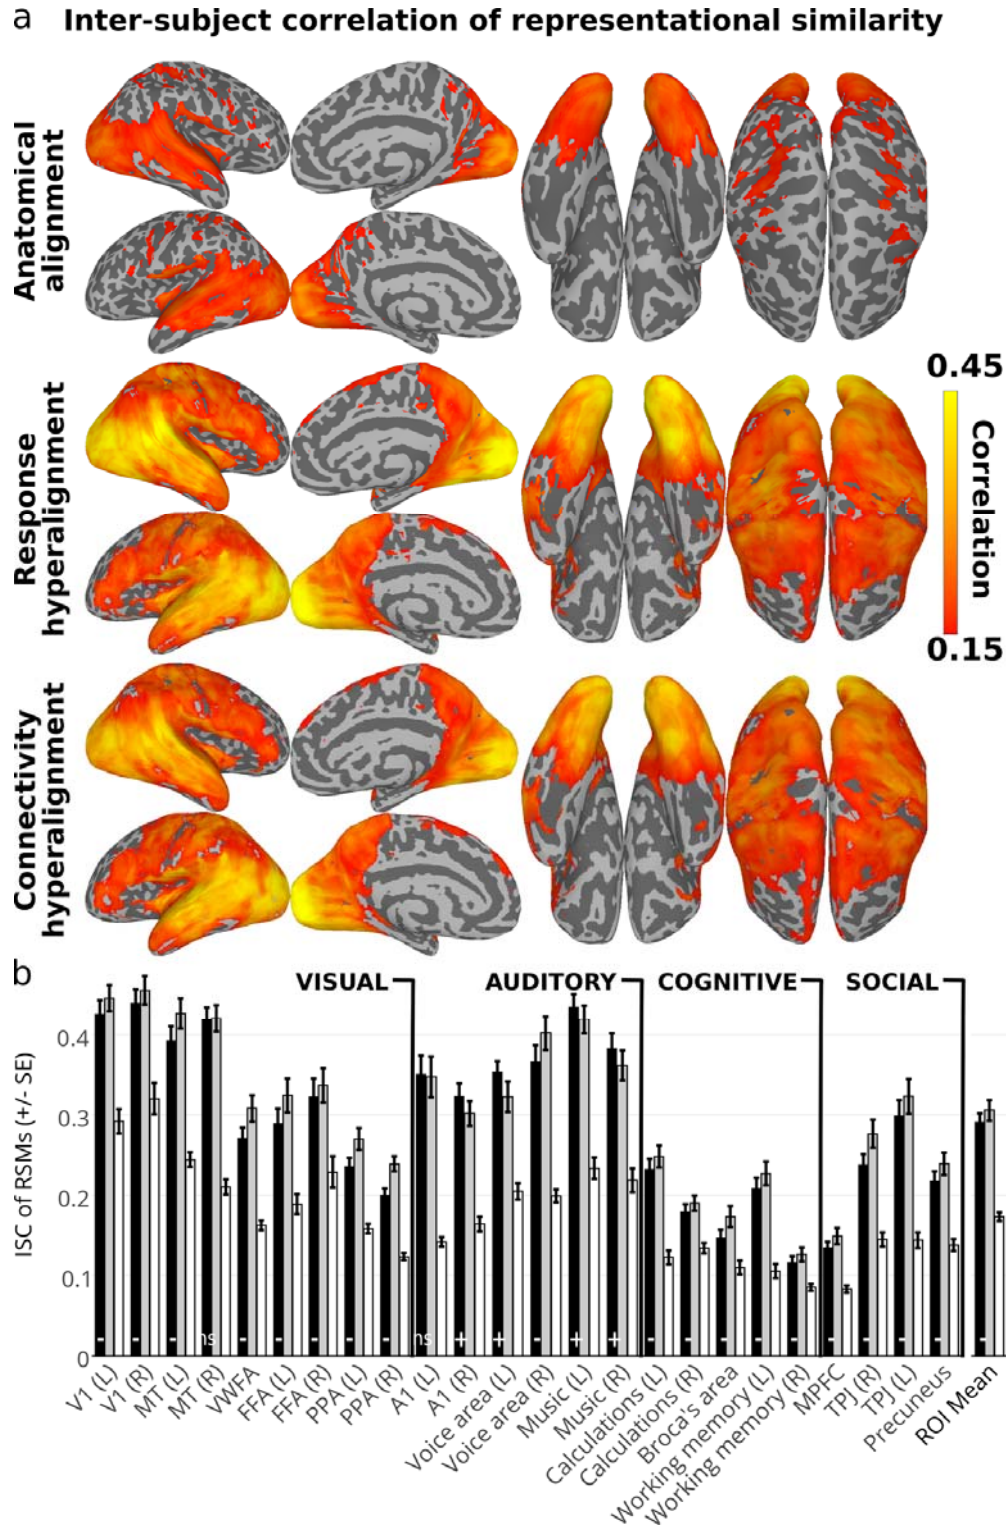


Supplementary Figure 1. ISC of connectivity profiles in the movie data with RHA and control analysis. (a) ISC of connectivity profiles mapped onto cortical surface after anatomical alignment, RHA, CHA, and filtering control. We controlled for the filtering effect of hyperalignment transformation by applying a transformation on each subject that projects that subject's data into a common model space derived with a different reference subject. That way, even though all subjects are projected to common model space, their correspondence is only based on anatomy. Increase in ISC of connectivity profiles

after CHA cannot be explained by the filtering effect. (b) ISC of connectivity profiles averaged within ROIs covering visual, auditory, cognitive, and social systems across the cortex and their mean comparing anatomical alignment, RHA and CHA. Bootstrapped testing showed significantly higher ISCs after both CHA and RHA relative to anatomical alignment in all ROIs, and slightly but significantly higher ISCs after CHA than after RHA in all but two ROIs (“-” : CHA<RHA; “ns” : no significant difference between ISCs after CHA and RHA; “+” : CHA>RHA). CHA afforded higher ISCs of connectivity profiles across ROIs than did RHA (ROI means = 0.82 and 0.79, respectively; difference = 0.03 [0.02, 0.04]).

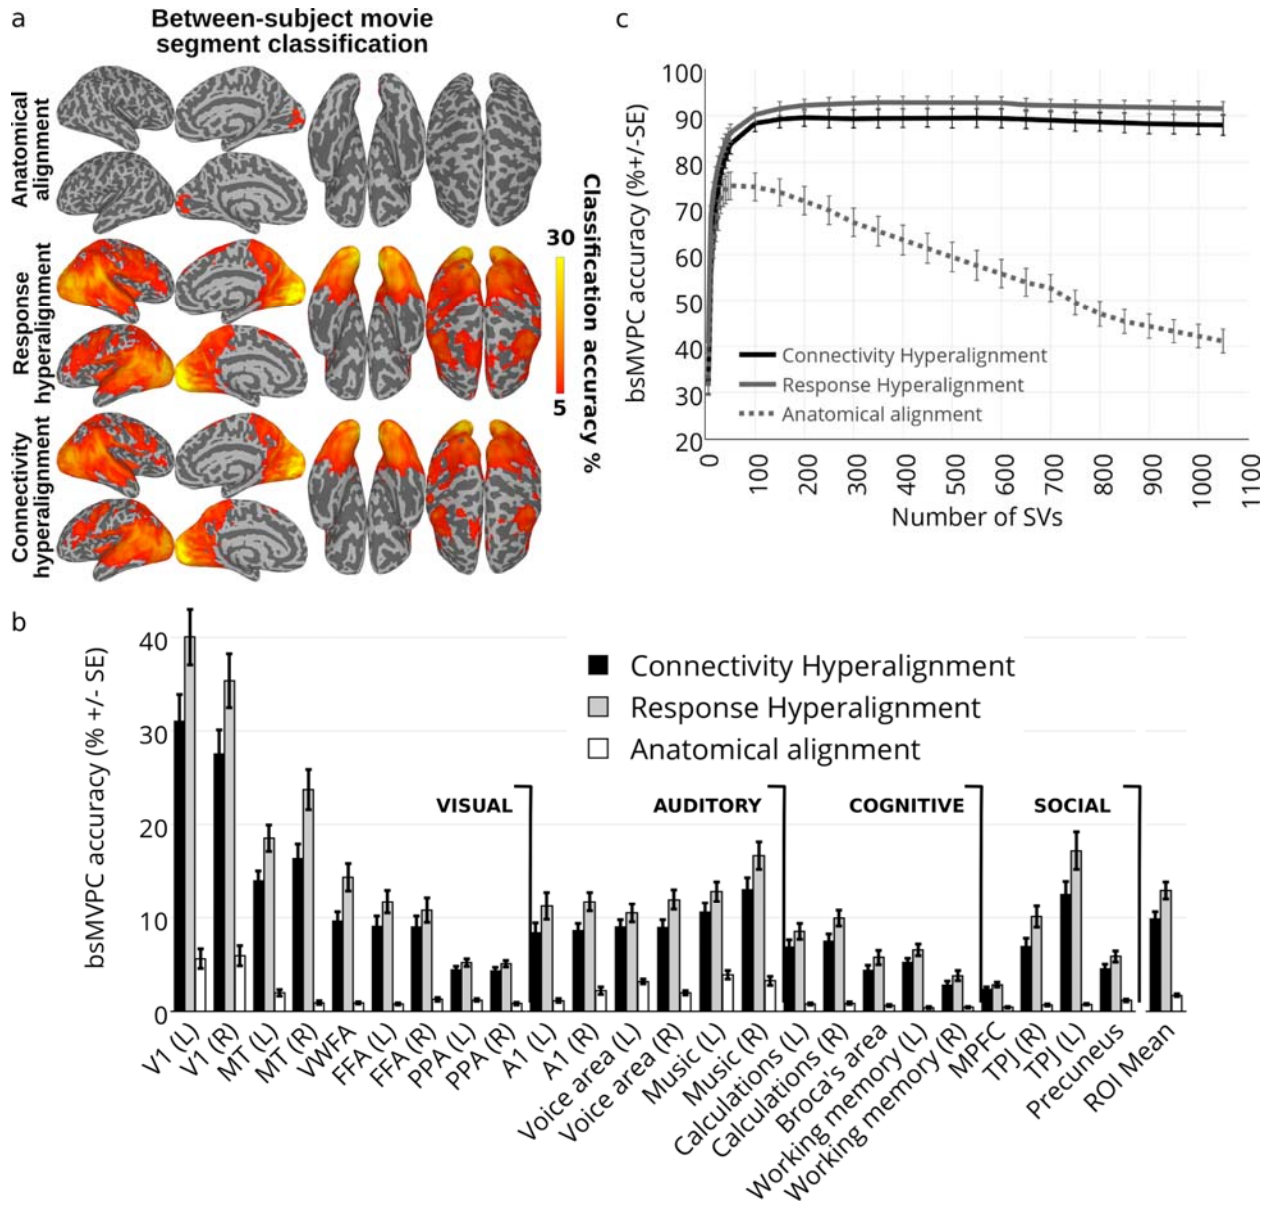


Supplementary Figure 2. Spatial granularity of connectivity profiles in the movie data. The intersubject point spread function (PSF) of connectivity profile correlations are computed as correlation between the connectivity profile for a cortical locus in one subject and the profiles of its spatial neighbors in other subjects at increasing distances from that locus (ISCs). Slope is estimated using a linear fit of these correlations at different distances as a function of distance in each of the ROIs (Left). Average PSF across all ROIs is plotted on the right. Bootstrapped testing showed significantly higher slopes of PSFs after both CHA and RHA relative to anatomical alignment in all ROIs, and some differences after CHA and RHA (“-” : CHA<RHA; “ns” : no significant difference between ISCs after CHA and RHA; “+” : CHA>RHA). Mean ROI PSF slopes for connectivity profiles were 0.077 and 0.080 after CHA and RHA, respectively, as compared to 0.035 after anatomical alignment (difference in mean slope between CHA and RHA = 0.003 [0.002, 0.004]).



Supplementary Figure 3. ISC of representational geometry in responses to movie time-points. (a) ISC of representational geometry in each voxel mapped onto cortical surfaces after anatomical alignment, RHA, and CHA. (b) ISC of representational geometries in 24 ROIs covering visual, auditory, cognitive, and social systems across the cortex and their mean after anatomical alignment and in the common model connectome. Bootstrapped testing showed significantly higher ISCs of representational geometry

after both CHA and RHA relative to anatomical alignment in all ROIs, and some differences after CHA and RHA (“-” : CHA<RHA; “ns” : no significant difference between ISCs after CHA and RHA; “+” : CHA>RHA). CHA increased the ISC of representational geometry in all of the ROIs (ROI mean ISC=0.291) relative to anatomical alignment (ISC=0.173) (difference=0.118 [0.103, 0.129]), but the increase is slightly, but significantly, less than that provided by RHA (ISC=0.306) (difference=0.015 [0.009, 0.019]).



Supplementary Figure 4. Between-subject classification of movie segments. Analysis procedure was identical to Fig. 7 with results after RHA included for comparison. (a) Classification accuracies in each searchlight mapped on cortical surfaces after anatomical alignment, RHA, and CHA. (b) Classification accuracies in 24 ROIs covering visual, auditory, cognitive, and social systems across the cortex and their mean after anatomical alignment, RHA, and CHA. Bootstrapped testing showed significantly higher accuracies after both the hyperalignment methods relative to anatomical alignment, and RHA slightly better in all ROIs relative to CHA. Bootstrapped testing showed significantly higher accuracies after RHA (ROI mean = 12.95%) than after CHA (ROI mean = 0.93%; mean difference = 3.02% [2.52%, 3.40%]) (c) Classification accuracies using information from multiple systems across the whole cortex. Dimensionality of the data is reduced using SVD and classification is performed with different set sizes of top singular vectors. Peak accuracy is reached after 200 dimensions for hyperaligned data and at 50 dimensions for anatomically aligned data. Peak accuracy after RHA is 92.98% and after CHA is 89.61% (mean difference = 3.37% [2.32%, 4.99%]).

Supplementary Tables

Supplementary Table 1 Selected cortical loci implicated in visual, auditory, cognitive, and social functions from Neurosynth.

Search Term	Hemisphere	MNI Coordinates (x, y, z)		
Visual				
V1	Left	-4	-82	-4
	Right	10	-94	2
MT	Left	-42	-72	2
	Right	44	-66	2
Visual word form area	Left	-46	-60	-14
FFA	Left	-42	-52	-20
	Right	40	-50	-22
Scenes (PPA)	Left	-26	-46	-10
	Right	34	-38	-12
Auditory				
Primary Auditory (A1)	Left	-44	-30	10
	Right	52	-14	4
Voice	Left	-60	-14	0
	Right	60	-4	-12
Music	Left	-52	-14	0
	Right	60	-20	4
Cognitive				

Calculations	Left	-30	-66	38
	Right	34	-64	44
Broca's area	Left	-52	14	12
Working memory	Left	-44	26	24
	Right	44	42	26
Social				
MPFC		6	54	14
TPJ	Left	50	-52	22
	Right	-52	-56	22
Precuneus		0	-56	40

V1 - primary visual cortex, MT - middle temporal visual motion area, FFA - fusiform face area, PPA - parahippocampal place area, MPFC - medial prefrontal cortex, TPJ - temporoparietal junction.

Supplementary Table 2 Task maps used from the HCP data.

Experiment/Task	Maps
MOTOR	CUE, LF, LH, RF, RH, T, CUE-AVG, LF-AVG, LH-AVG, RF-AVG, RH-AVG, T-AVG
LANGUAGE	MATH, STORY, MATH-STORY
WM (Working Memory)	2BK_BODY, 2BK_FACE, 2BK_PLACE, 2BK_TOOL, 0BK_BODY, 0BK_FACE, 0BK_PLACE, 0BK_TOOL, 2BK, 0BK, 2BK-0BK, BODY, FACE, PLACE, TOOL, BODY-AVG, FACE-AVG, PLACE-AVG, TOOL-AVG
RELATIONAL	MATH, REL, MATCH-REL
EMOTION	FACES, SHAPES, FACES-SHAPES
SOCIAL	RANDOM, TOM, TOM-RANDOM
GAMBLING	PUNISH, REWARD, PUNISH-REWARD

LF - Left Finger, RH - Right Hand, 2BK - Two-Back, TOM - Theory Of Mind