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

*Correspondence to: james.v.haxby@dartmouth.edu

Abstract

Variation in cortical connectivity profiles is typically modeled as having a coarse spatial scale parcellated into interconnected brain areas. We created a high-dimensional common model of the human connectome to search for fine-scale structure that is shared across brains. Projecting individual connectivity data into this new common model connectome accounts for over three times more variance in the human connectome than do previous models. This newly discovered shared structure resides in fine-scale local variation that is closely related to fine-scale distinctions in representations of information. These results reveal a shared fine-scale structure that is a markedly larger component of the human connectome than coarse-scale, areal structure. This shared fine-scale structure was not captured in previous models and was, therefore, inaccessible to analysis and study.

Author Summary

Resting state fMRI has become a ubiquitous tool for measuring connectivity in normal and diseased brains. Current dominant models of connectivity are based on coarse regional connectivity ignoring fine-scale structure within those regions. We developed a high-dimensional model common model of the human connectome that captures both coarse and fine-scale structure of connectivity shared across brains. We showed that this shared fine-scale structure is related to fine-scale distinctions in representation of information, and our model accounts for over three times more shared variance of connectivity compared to previous models. Our model opens new territory – shared fine-scale structure, a dominant but mostly unexplored component of the human connectome – for analysis and study.

Introduction

Resting state functional magnetic resonance imaging (rsfMRI) reveals patterns of functional connectivity that are used to investigate the human connectome [1-3] and parcellate the brain into interconnected areas that form brain systems and can be modeled as networks [4-11]. The connectivity of a single area is considered to be relatively homogeneous and typically is modeled as a mean connectivity profile. Cortical topography, however, has both a coarse scale of cortical areas and a finer scale of multiplexed topographies within areas [12-16]. Fine-scale within-area topographies are reflected in patterns of activity that can be measured with fMRI and decoded using multivariate pattern analysis (MVPA)[12,13,17]. Fine-scale variation in connectivity, however, has been overlooked due to poor anatomical alignment of this variation across individual brains. We ask here whether local variation in functional connectivity also has a fine-scale structure, similar to fine-scale response tuning topographies, 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 (Fig 1). The resultant common model consists of transformation matrices for each individual brain and a common model connectome space defined by shared connectivity basis functions. Individual transformation matrices transform an individual brain's connectome, in its native

anatomical coordinate space, into the common model coordinate space [13]. The shared connectivity basis functions that define common model coordinate dimensions are a matrix of mean transformed individual connectivity profiles. The individual transformation matrices and common model connectivity matrix are derived iteratively from training data. Validity testing is done on connectivity profiles and other functional parameters from independent test data that are hyperaligned into the common model connectome space. The results show that CHA can derive these shared basis functions from resting state functional connectivity, as well as from functional connectivity while watching an audiovisual movie.

The resultant common model connectome accounts for over three times more shared variance than was accounted for by previous models. This shared variance resides in fine-scale local variations in connectivity. We show further that this local variability in functional connectivity profiles is meaningful in that it is closely related to local patterns of response that encode fine distinctions among representations. Our results indicate that shared fine-scale local variation, which was not evident in previous models, is a markedly larger component of the human connectome than is shared areal structure. Our common model connectome makes this fine-scale local variation accessible for group-level study of its network properties.

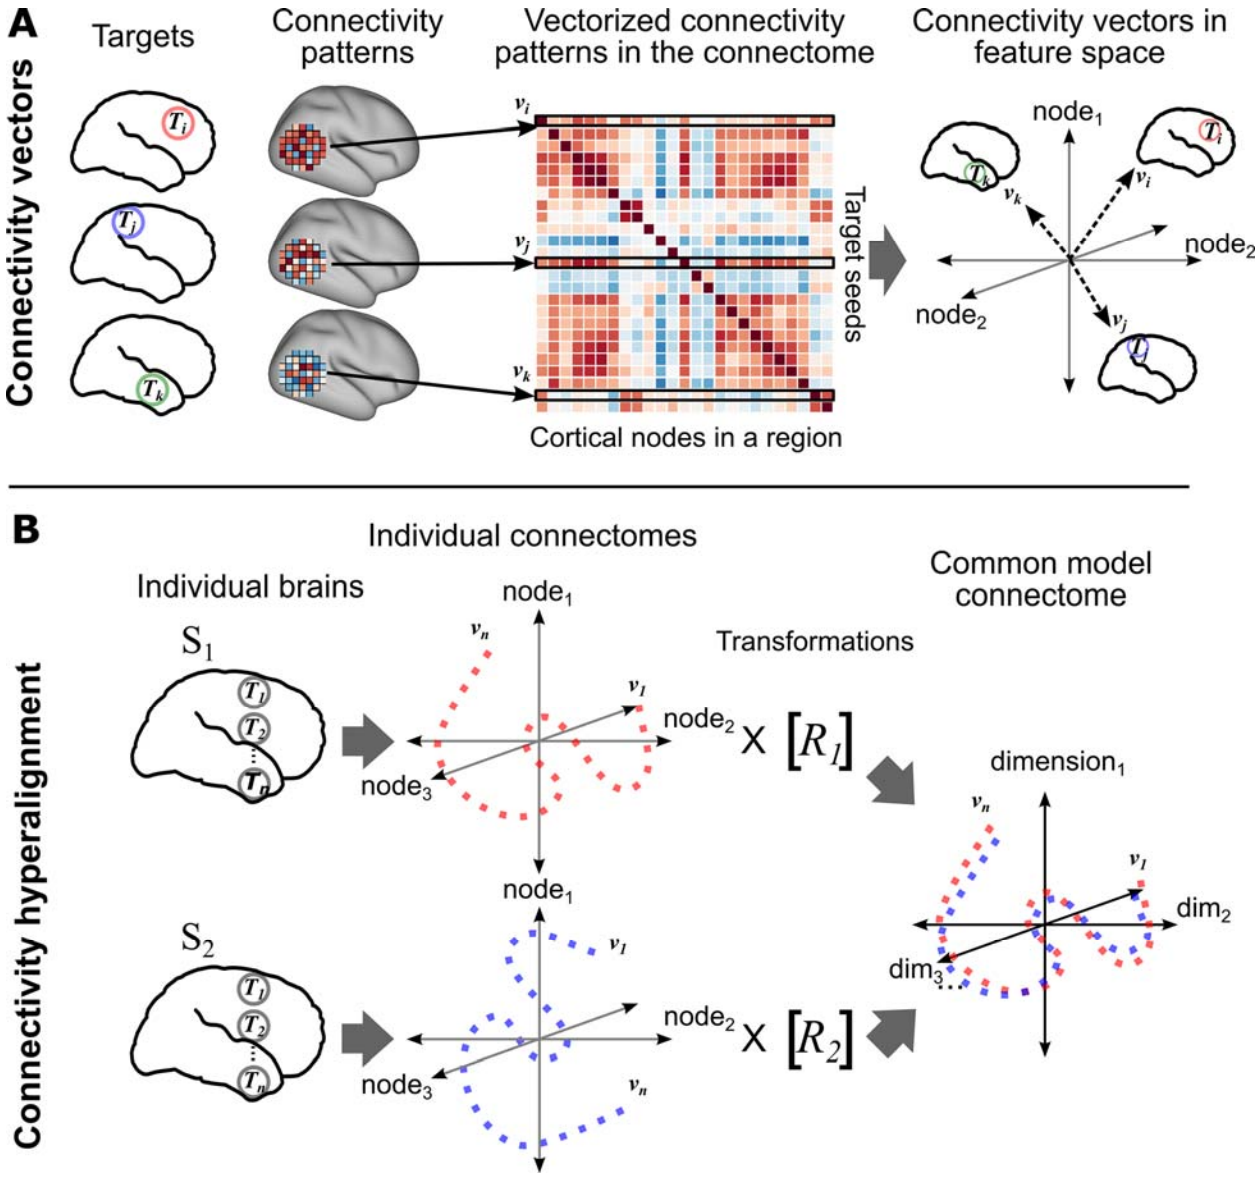

Fig. 1. Schematic of connectivity hyperalignment (CHA). A. Connectivity can be defined as any measure of similarity between a cortical locus (eg. surface node/voxel) and a target region. Connectivities of loci in a searchlight to a target region (T_i, T_j, or T_k) yield a connectivity pattern for that target in that searchlight. These patterns can be analyzed as connectivity vectors ($\boldsymbol{v}_i$, $\boldsymbol{v}_j$, $\boldsymbol{v}_k$) in a space in which each cortical locus in that region is a dimension. B. Connectivity vectors ($\boldsymbol{v}_1$ ... $\boldsymbol{v}_n$) in a region of interest or a searchlight to be hyperaligned are calculated for target regions (T_1, T_2, ..., T_n,) distributed uniformly across the whole cortex. Connectivity hyperalignment derives transformation matrices for each brain ($\boldsymbol{R}_1$, $\boldsymbol{R}_2$, ...) that align these vectors across subjects into a common high-dimensional connectivity space. Searchlight transformation matrices are aggregated into a whole brain matrix, as in [16], to produce a common model connectome for the whole cortex.

Results

We derived a common model of the human connectome by applying CHA to rsfMRI data for 20 subjects in the Human Connectome Project (HCP) database [18-20]. The common model connectome is high-dimensional with connectivity profiles for dimensions that serve as basis functions for modeling the connectivity profiles of cortical loci in individual brains. We validated the common model in terms of 1) portion of total variance that is accounted for by shared connectivity profiles, 2) intersubject correlations (ISCs) of connectivity profiles, and 3) the spatial specificity of shared connectivity profiles. To test whether this fine-scale structure is meaningful for the representation of information, we derived a common model connectome by applying CHA to fMRI data collected during viewing a naturalistic audiovisual movie and tested 4) ISCs of representational geometry, and 5) between-subject multivariate pattern classification (bsMVPC) of responses to the movie. See Supplementary Data and Figures for results of further validation tests.

Variance accounted for and intersubject correlation of connectivity profiles

Transforming functional connectivity data into the common model connectome space increased the variance accounted for (VAF) by shared cortical node-to-node connectivity profiles to 46.8% (95% confidence interval, CI: [42.4%,49.0%]), up from 14.7% [13.4%,16.0%] VAF for data in a space aligned with the HCP's MSM-All method (Multimodal Surface Matching)[21]. An areal model in which each node within a parcel is assigned the mean connectivity profile for that parcel [9] further reduces VAF to 12.7% [11.3%,14.0%]. This *3.2 to 3.7 fold*

increase in VAF, relative to anatomically aligned and areal models, is evident as widely-distributed increases in ISCs of resting state connectivity profiles across all of cortex (Fig 2). ISC at a cortical node is the correlation of the one subject's connectivity profile with the mean of other subjects' profiles, indexing how well other subjects' connectivity profiles can predict an individual's connectivity profiles. Fig 2A shows a cortical map of mean ISCs of connectivity profiles in MSM-All-aligned data and in the common model connectome space. Fig 2B is a scatterplot of mean ISCs for individuals after MSM-All alignment and CHA, which shows that CHA increased ISC for each individual and preserved individual similarity or deviance from the group. Fig 2C shows a cortical map of within-subject correlations between connectivity profiles from different resting state sessions. We quantify the increases in 26 functional ROIs, identified using a meta-analytic database, NeuroSynth [22](Fig 2D; table S1). Mean ISC of connectivity profiles across these ROIs was markedly higher in the common model connectome than in the MSM-All-aligned data (0.66 versus 0.35; 95% confidence interval, CI, for mean difference = [0.30, 0.33]). 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 = 0.64; CI for difference = [0.00, 0.05]) (Fig 2D). This latter result indicates that an individual's connectome is better predicted by the common model connectome, based on other subjects' data, than by estimates based on a typical sample of that subject's own rsfMRI data,

due to the benefit of estimating connectivity profiles based on a large number of brains and the precision of CHA.

Spatial granularity of connectivity profile variation

We investigated the spatial specificity of the common model connectome by computing the spatial point spread functions (PSF) of ISCs of connectivity profiles [16]. The PSF of connectivity profiles was computed as the correlation of the connectivity profile in a cortical surface node for a given subject with the average connectivity profiles of other subjects in the same node and nodes at cortical distances ranging from 3 to 12 mm. We similarly calculated within-subject PSFs based on within-subject correlations (WSC) of connectivity profiles between two resting state sessions. Fig 3 shows the slopes of connectivity profile PSFs in the ROIs (Fig 3A) and the mean PSF across ROIs (ISC or WSC as a function of cortical distance, Fig 3B) in MSM-All-aligned data, in the common model connectome space, and within-subject. CHA increased the average slope of PSF across these ROIs, relative to MSM-All alignment, from 0.012 to 0.065 (difference=0.053 [0.047, 0.055]). The PSF slopes in the common model connectome space and within subject were not significantly different (slope = 0.067; difference = 0.002 [-0.002, 0.007]).

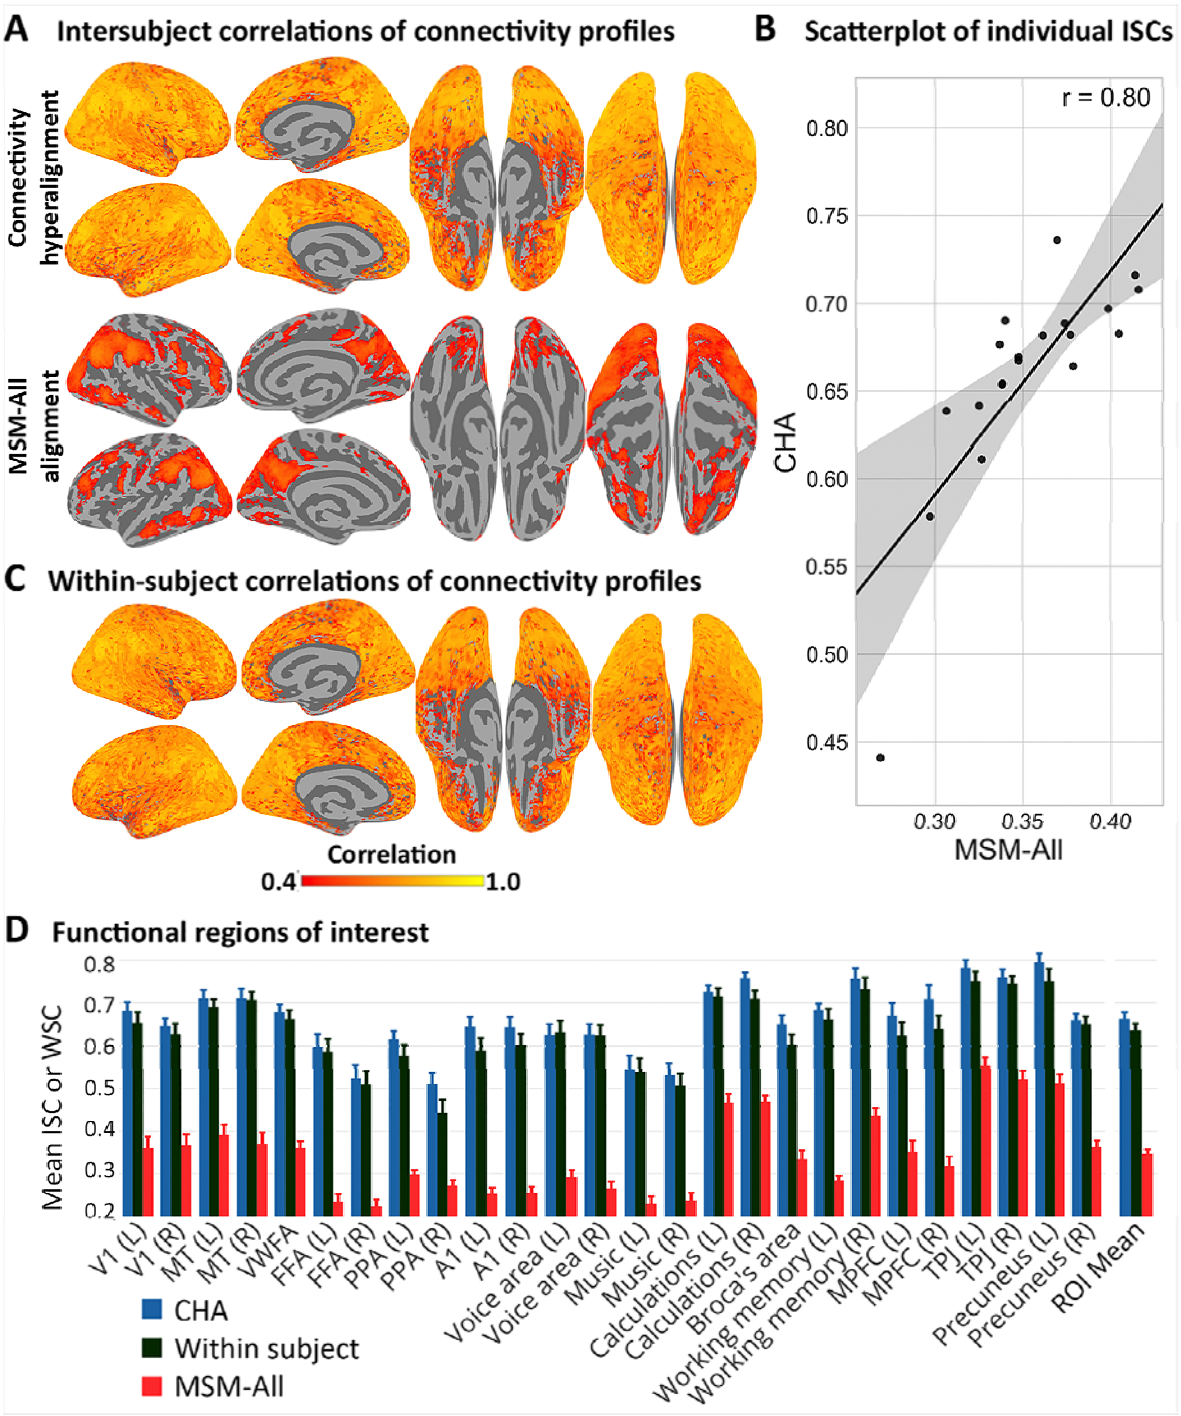

Fig. 2. ISC of connectivity profiles calculated from HCP rsfMRI data. A. Average ISC of connectivity profiles in each surface node in the common model connectome space and after surface alignment (MSM-All). B. Scatter plot of individual whole cortex mean ISCs of connectivity profiles before and after CHA with linear fit. Each subject's similarity of connectome with the group is improved by CHA while preserving similarity or deviance from others. Shaded region is the 95% CI. C. Average within-subject between-session correlations in the common space. D. Mean ISCs and WSCs of connectivity

profiles in functional ROIs covering visual, auditory, cognitive, and social systems comparing the common model connectome space, within-subject between-session correlation in common space, and surface alignment.

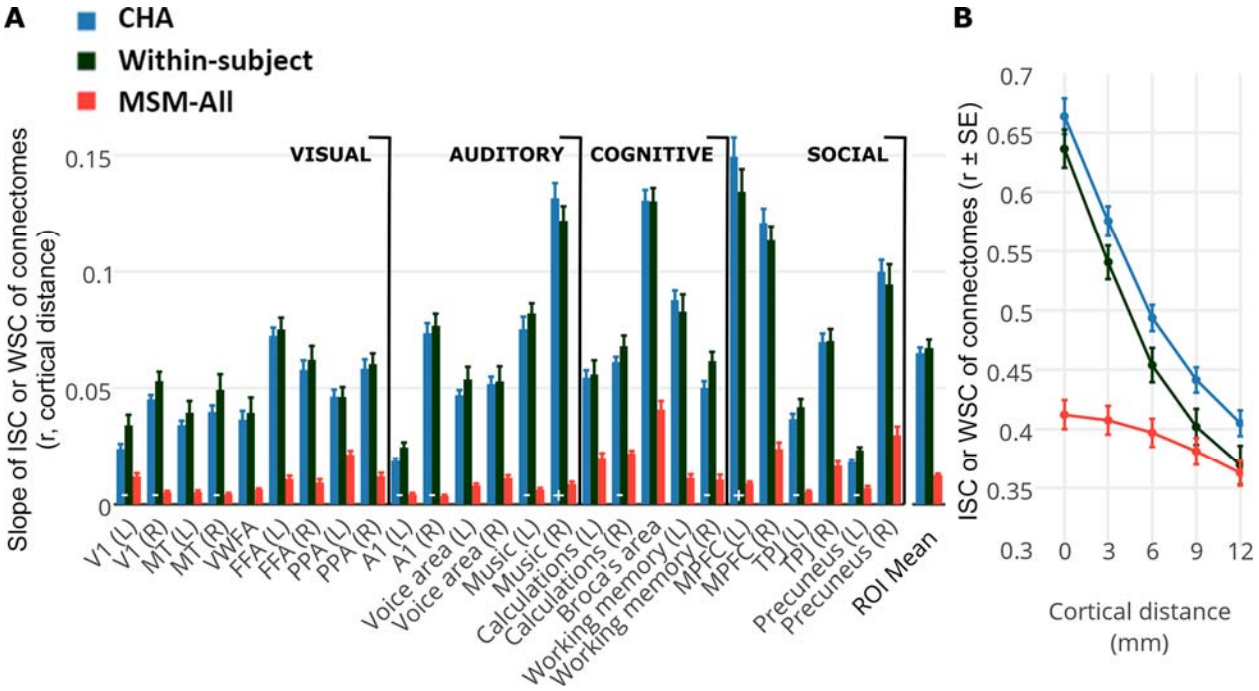

Fig. 3. Spatial granularity of shared connectivity profiles. The intersubject point spread function (PSF) of connectivity profile correlations are computed as the correlation between the connectivity profile for a cortical locus in one subject and the profiles of the same locus and its spatial neighbors in other subjects at increasing distances from that locus. 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 the same locus and its spatial neighbors from a different rsfMRI session. A. Slope is estimated in each functional ROI as the linear fit of intersubject or within subject correlations as a function of distance. B. Average PSF across all ROIs is plotted as ISC or WSC as a function of cortical distance.

The mean PSF across ROIs (Fig 3B), clearly shows that CHA captures fine-scale variations in connectivity profiles for neighboring cortical nodes across subjects that are not captured by MSM-All alignment. The ISCs of connectivity profiles for neighboring nodes in the common model connectome are substantially lower than ISCs for the same node (difference=0.09 [0.08,0.09]). Similar fine spatial granularity is seen in the within-subject between-session PSFs

(0.10 [0.09,0.10]). By contrast, ISCs for connectivity profiles in the MSM-All aligned data barely differ for nodes spaced 0 versus 3 (difference=0.004, [0.004,0.015]) and 6 mm (0.02 [0.01, 0.02]) apart, with only small decrements for ISCs of nodes spaced 9 mm (0.03 [0.03,0.03]) and even 12 mm (0.05 [0.04,0.05]) apart. This fine spatial granularity was ubiquitous in cortex, with steep PSFs in sensory-perceptual areas in occipital and temporal cortices as well as in higher-order cognitive areas in lateral and medial parietal and prefrontal cortices.

Generalization to fine-scale patterns in response tuning

Next we asked if this shared variance in fine-scale local variation in connectivity profiles carries meaning by testing whether it reflects fine-scale variations in response tuning topographies that carry fine-grained distinctions in representation. We built a common model of the connectome based on functional connectivity in fMRI data collected while 11 subjects watched an audiovisual movie and tested whether projecting movie response data into the common connectome space afforded better alignment of representational geometry for movie time-points and better bsMVPC of movie time segments.

Results show that shared fine-scale structure in the common model connectome is closely related to fine distinctions in representations. Fig 4A shows a cortical map of mean ISCs of local representational geometry after anatomical alignment and in the common model connectome. Representational geometry is the matrix of all pairwise similarities between patterns of response to different time-points in the movie. Fig 4B shows a cortical

map of mean bsMVPC accuracies for 15 s movie time-segments in searchlights after anatomical alignment and CHA. CHA greatly increased both ISCs of representational geometry and bsMVPC accuracies. Quantification of these effects in selected functional ROIs is illustrated in Fig S3. CHA significantly increased ISCs of representational geometry and bsMVPC accuracies in all ROIs (ROI mean ISCs = 0.291 and 0.173 after CHA and anatomical alignment, respectively, difference=0.118 [0.103, 0.129]; ROI mean bsMVPC accuracies = 9.93% and 1.73% after CHA and anatomical alignment, respectively, difference = 8.19% [6.83%, 9.14%]).

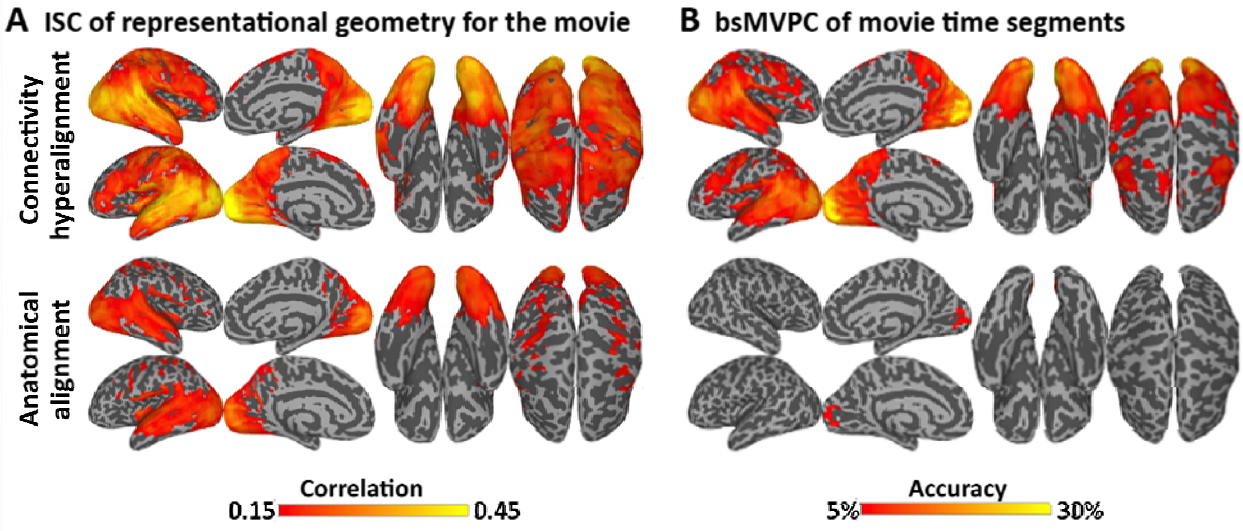

Fig 4. A. ISC of representational geometry in each voxel mapped onto the cortical surface. B. Accuracies for bsMVPC of 15 s movie segments. Classification was performed within each movie half separately, and the accuracies are then averaged across the two halves. Parameters for hyperalignment and SVD are derived from the half that was not used for classification.

Comparison of CHA and Response Hyperalignment (RHA)

Since CHA aligned fine-scale patterns of response tuning functions across subjects better than anatomy-based alignment, we asked how well it compares

to our previously published response-based hyperalignment (RHA) [16]. We used the audio visual movie data for this comparison as RHA requires responses that are synchronized across subjects in time. We compare CHA and RHA on 1) ISC of connectivity profiles, 2) ISC of response profiles, 3) ISC of representational geometry, and 4) bsMVPC of 15 s movie segments.

Results showed that both CHA and RHA increased ISCs and bsMVPC classification accuracies significantly over anatomy-based alignment, but each algorithm achieves better alignment for the information that it uses to derive a common model, namely connectivity profiles and patterns of response, respectively. A common model based on CHA has slightly higher ISCs of connectivity profiles than a common model based on RHA (ROI mean ISCs = 0.82 and 0.77, respectively; CHA-RHA difference=0.05 [0.04, 0.07]) (Fig 5). By contrast, RHA slightly outperforms CHA on validations based on response tuning functions, namely ISCs of response profiles (ROI means = 0.181 and 0.169 respectively; RHA-CHA difference = 0.012 [0.008, 0.016]), ISCs of representational geometry (ROI means = 0.306 and 0.291, respectively; RHA-CHA difference = 0.015 [0.009, 0.019]), and bsMVPC of movie segments (ROI mean accuracies = 12.95% and 9.93%, respectively; RHA-CHA difference=3.02% [2.52%, 3.40%]).

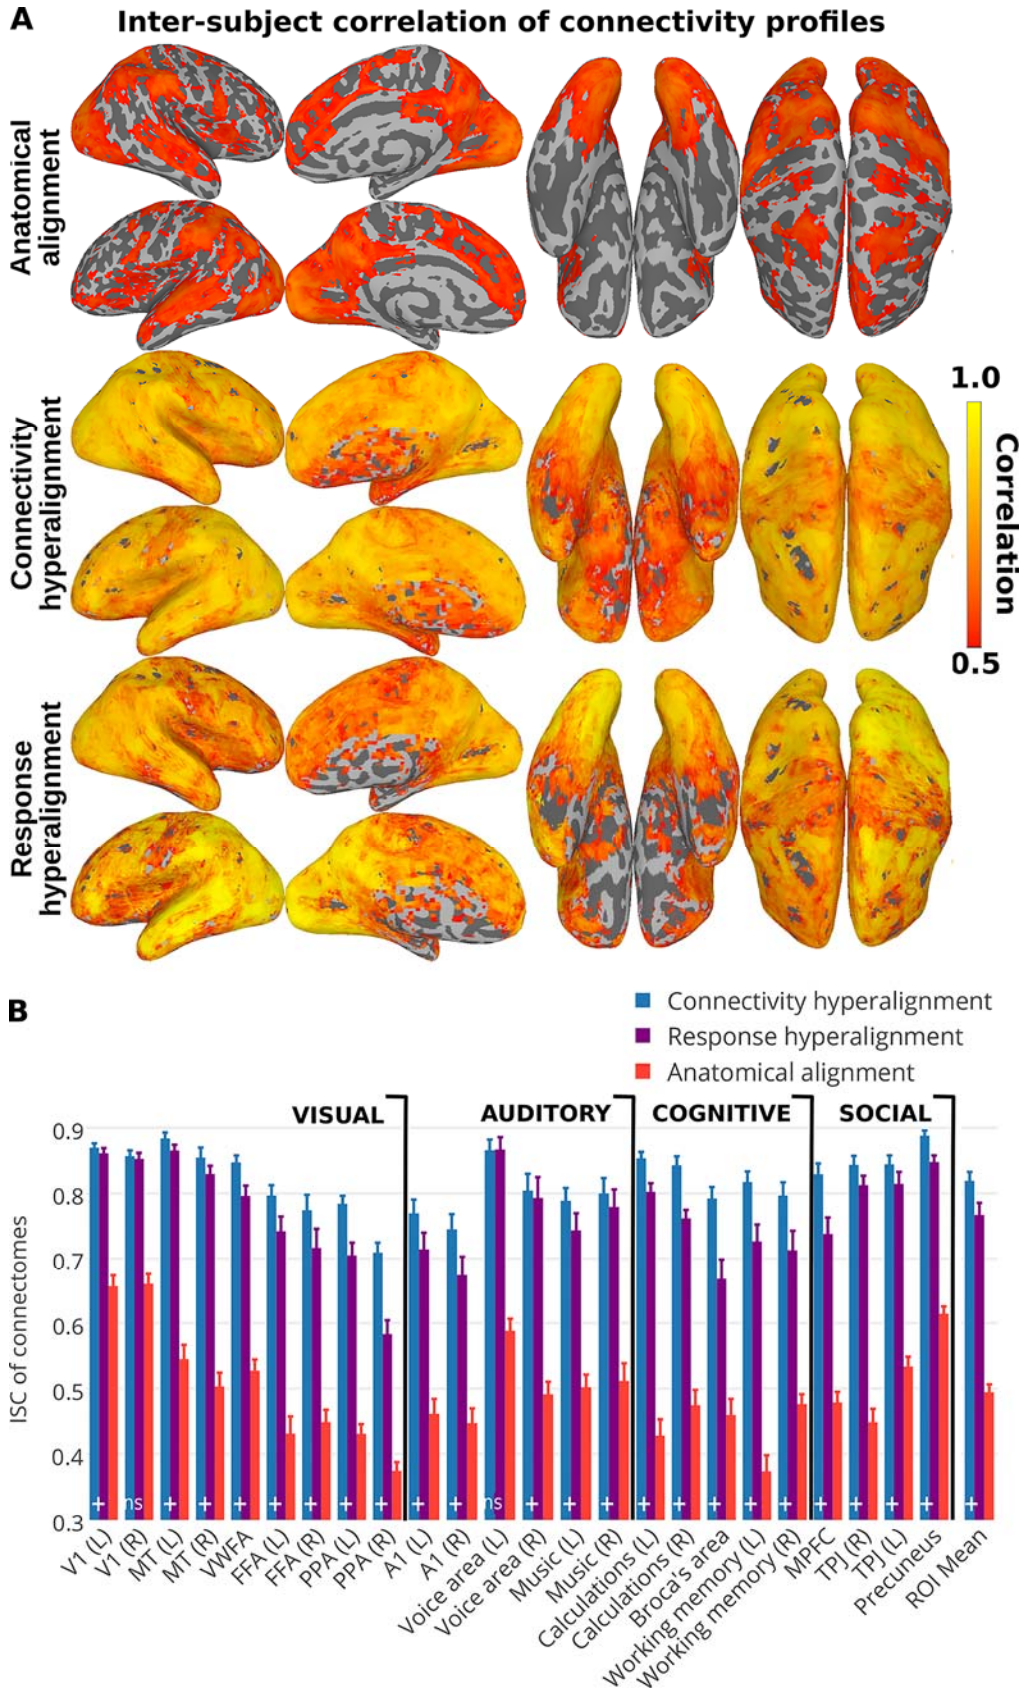

Fig 5. ISC of connectivity profiles calculated from movie data. (A) Average ISCs of connectivity profiles in each surface node after anatomical alignment, CHA, and RHA. (B) ISCs of connectivity profiles averaged within functional ROIs. Bootstrapped testing showed significantly higher ISCs after both 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 ("ns" : no significant difference between ISCs after CHA and RHA; "+" : CHA>RHA).

Discussion

These results show that fine-scale local variation in connectivity profile is a major component of the human connectome that can be modeled with shared connectivity basis functions. This shared fine-scale variation is a ubiquitous characteristic of all of human cortex and is a larger component of the human connectome than is shared coarse, areal variation.

The existence and importance of fine-scale connectivity is well-recognized [23-25] but previously could not be modeled in a common computational framework and, consequently, was largely overlooked. Attempts to model within-area topographies of connectivity either were limited mostly to within-subject analyses or coarser within-area topographies that could be captured with anatomy-based alignment of group data [25]. Consequently, within-area variations in connectivity profiles were usually analyzed as coarse gradients that have a single cycle in a cortical area, such as retinotopy or somatotopy [23-25].

Cortical functional architecture, however, has multiplexed topographies at multiple spatial scales. In primary visual cortex, retinotopy is multiplexed with ocular dominance columns, edge orientation, spatial frequency, motion

direction, and motion velocity, among other low-level visual attributes [26,27]. Primary visual cortex sends coherent projections to other visual areas where these topographies are recapitulated and transformed, affording the emergence of more complex features, such as curvature, texture, shape, color constancy, and biological motion, and, subsequently, even higher-order attributes such as object categories, view-invariant face identity, and species-invariant attributes of animals such as action categories and dangerousness [12,15,28-32]. Similar transformations of topographies characterize other sensory modalities and, undoubtedly, supramodal cognitive operations. Modeling inter-areal communication as a single value of connectivity strength sheds no light on how information is transformed along cortical processing pathways to allow high-order information to be disentangled from confounding attributes [33].

Multiplexed cortical topographies at multiple spatial scales can be modeled with individual-specific topographic basis functions that have shared tuning profiles [13,16] and shared connectivity profiles (as shown here). No previous model captured multiple spatial scales of connectivity topographies with connectivity profiles that are shared across brains. By capturing coarse- and fine-scale connectivity topographies with shared basis functions, the common model connectome casts a bright light on the dominant role of fine-scale connectivity patterns in the human connectome and opens new territory for investigation of the network properties of cortical connectivity at finer levels of detail. With this new perspective, inter-areal connectivity can be modeled as

more than a simple replication of global activity, as is the assumption underlying existing approaches to modeling the connectome, but, instead, as information processing operations in which functional topographies are transformed by projections between areas.

Methods

Resting state data: Human Connectome Project

In the HCP database [20], 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 MSM-All [21] with 32K nodes in each hemisphere as provided by the HCP. We used data from one resting state session [19](“rfMRI_REST1_LR”) to derive CHA parameters and validated it on a different resting state session (“rfMRI_REST2_LR”), and task fMRI sessions [18](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 [19] 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 [34] 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 [16] 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 [35](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 [16] 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 [36] 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 [37,38]. The RHA algorithm uses the same procedure but minimizes the distances between response vectors for the same time-points or stimuli across subjects [13,16]. 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 x 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 x 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 [39] approach as described previously in [16] 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 [16]. 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 [16].

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 x 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 [16]. 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 [16]. We searched for terms and cortical areas implicated in visual, auditory, cognitive, and social functions in NeuroSynth [22] 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 [40]. 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 all other cortical nodes. 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. 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.

Variance accounted for by shared connectivity profiles. We computed the variance of feature (or cortical node) connectivity profiles explained by other subject's connectivity profiles using the ISCs of connectivity profiles. We squared the ISCs of feature connectivity profiles to compute the variance of each feature's connectivity profile explained by the average of connectivity profiles

of features at the same cortical location in other subjects. We computed this for session REST2 before and after CHA. We averaged this variance in all features in all subjects to get a mean value of VAF.

Spatial point spread function. To investigate the spatial granularity of representation, we computed a spatial point spread function (PSF) of ISCs or WSCs of connectivity profiles. We computed the correlation of connectivity profiles in each cortical locus (surface node or voxel) 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 MSM-All or anatomical alignment and after CHA with each subject aligned to the same reference subject [16]. 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.

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 patterns of response for 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 [13,16]. 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 a singular vector (SV) space derived from 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 [41](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 [42](http://neuro.debian.net).

Acknowledgements

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.

References

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

2. Smith SM, Vidaurre D, Beckmann CF, Glasser MF, Jenkinson M, Miller KL, et al. Functional connectomics from resting-state fMRI. Trends in Cognitive Sciences. 2013;17: 666–682.

3. Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, et al. Dynamic functional connectivity: Promise, issues, and interpretations. NeuroImage. 2013;80: 360–378.

4. Sporns O, Chialvo DR, Kaiser M, Hilgetag CC. Organization, development and function of complex brain networks. Trends in Cognitive Sciences. 2004;8: 418–425.

5. Thomas Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of Neurophysiology. 2011;106: 1125–1165.

6. Wig GS, Laumann TO, Cohen AL, Power JD, Nelson SM, Glasser MF, et al. Parcellating an Individual Subject's Cortical and Subcortical Brain Structures Using Snowball Sampling of Resting-State Correlations. Cereb Cortex. 2014;24: 2036–2054.

7. Laumann TO, Gordon EM, Adeyemo B, Snyder AZ, Joo SJ, Chen M-Y, et al. Functional System and Areal Organization of a Highly Sampled Individual Human Brain. Neuron. 2015;87: 657–670.

8. Gordon EM, Laumann TO, Adeyemo B, Petersen SE. Individual Variability of the System-Level Organization of the Human Brain. Cereb Cortex. 2017;27: 386–399.

9. Glasser MF, Coalson TS, Robinson EC, Hacker CD, Harwell J, Yacoub E, et al. A multi-modal parcellation of human cerebral cortex. Nature. 2016;536: 171–178.

10. Gordon EM, Laumann TO, Adeyemo B, Huckins JF, Kelley WM, Petersen SE. Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. Cereb Cortex. 2016;26: 288–303.

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

12. Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. Science. 2001;293: 2425–2430.

13. Haxby JV, Guntupalli JS, Connolly AC, Halchenko YO, Conroy BR, Gobbini MI, et al. A Common, High-Dimensional Model of the Representational Space

in Human Ventral Temporal Cortex. Neuron. 2011;72: 404–416.

14. Haxby JV, Connolly AC, Guntupalli JS. Decoding Neural Representational Spaces Using Multivariate Pattern Analysis. Annual Review of Neuroscience. 2014;37: 435–456.

15. Grill-Spector K, Weiner KS. The functional architecture of the ventral temporal cortex and its role in categorization. Nat Rev Neurosci. 2014;15: 536–548.

16. Guntupalli JS, Hanke M, Halchenko YO, Connolly AC, Ramadge PJ, Haxby JV. A Model of Representational Spaces in Human Cortex. Cereb Cortex. 2016;26: 2919–2934.

17. Norman KA, Polyn SM, Detre GJ, Haxby JV. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends in Cognitive Sciences. 2006;10: 424–430.

18. Barch DM, Burgess GC, Harms MP, Petersen SE, Schlaggar BL, Corbetta M, et al. Function in the human connectome: Task-fMRI and individual differences in behavior. NeuroImage. 2013;80: 169–189.

19. Smith SM, Beckmann CF, Andersson J, Auerbach EJ, Bijsterbosch J, Douaud G, et al. Resting-state fMRI in the Human Connectome Project. NeuroImage. 2013;80: 144–168.

20. Van Essen DC, Smith SM, Barch DM, Behrens TEJ, Yacoub E, Ugurbil K. The WU-Minn Human Connectome Project: An overview. NeuroImage. 2013;80: 62–79.

21. Robinson EC, Jbabdi S, Glasser MF, Andersson J, Burgess GC, Harms MP, et al. MSM: A new flexible framework for Multimodal Surface Matching. NeuroImage. 2014;100: 414–426.

22. Yarkoni T, Poldrack RA, Nichols TE, Essen DCV, Wager TD. Large-scale automated synthesis of human functional neuroimaging data. Nature Methods. 2011;8: 665–670.

23. Heinzle J, Kahnt T, Haynes J-D. Topographically specific functional connectivity between visual field maps in the human brain. NeuroImage. 2011;56: 1426–1436.

24. Jbabdi S, Sotiropoulos SN, Behrens TE. The topographic connectome. Current Opinion in Neurobiology. 2013;23: 207–215.

25. Haak KV, Marquand AF, Beckmann CF. Connectopic mapping with resting-state fMRI. arXiv:160207100 [q-bio]. 2016; Available: http://arxiv.org/abs/1602.07100. Cited 13 April 2017.

26. Yacoub E, Harel N, Uğurbil K. High-field fMRI unveils orientation columns in humans. PNAS. 2008;105: 10607–10612.

27. Nishimoto S, Vu AT, Naselaris T, Benjamini Y, Yu B, Gallant JL. Reconstructing Visual Experiences from Brain Activity Evoked by Natural Movies. Current Biology. 2011;21: 1641–1646.

28. Kanwisher N. Functional specificity in the human brain: A window into the

functional architecture of the mind. PNAS. 2010;107: 11163–11170.

29. Freiwald WA, Tsao DY. Functional Compartmentalization and Viewpoint Generalization Within the Macaque Face-Processing System. Science. 2010;330: 845–851.

30. Guntupalli JS, Wheeler KG, Gobbini MI. Disentangling the Representation of Identity from Head View Along the Human Face Processing Pathway. Cereb Cortex. 2017;27: 46–53.

31. Connolly AC, Sha L, Guntupalli JS, Oosterhof N, Halchenko YO, Nastase SA, et al. How the Human Brain Represents Perceived Dangerousness or "Predacity" of Animals. J Neurosci. 2016;36: 5373–5384.

32. Nastase SA, Connolly AC, Oosterhof NN, Halchenko YO, Guntupalli JS, Visconti di Oleggio Castello M, et al. Attention selectively reshapes the geometry of distributed semantic representation. Cereb Cortex. 2017; Forthcoming.

33. DiCarlo JJ, Cox DD. Untangling invariant object recognition. Trends in Cognitive Sciences. 2007;11: 333–341.

34. Saad ZS, Reynolds RC, Argall B, Japee S, Cox RW. SUMA: an interface for surface-based intra- and inter-subject analysis with AFNI. IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2004. 2004. p. 1510–1513 Vol. 2.

35. Cox RW. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput Biomed Res. 1996;29: 162–173.

36. Schönemann PH. A generalized solution of the orthogonal procrustes problem. Psychometrika. 1966;31: 1–10.

37. Chen P-H, Chen J, Yeshurun Y, Hasson U, Haxby JV, Ramadge PJ. A Reduced-dimension fMRI Shared Response Model. Proceedings of the 28th International Conference on Neural Information Processing Systems. Cambridge, MA, USA: MIT Press; 2015. pp. 460–468.

38. Xu H, Lorbert A, Ramadge PJ, Guntupalli JS, Haxby JV. Regularized hyperalignment of multi-set fMRI data. 2012 IEEE Statistical Signal Processing Workshop (SSP). IEEE; 2012. pp. 229–232.

39. Oosterhof NN, Wiestler T, Downing PE, Diedrichsen J. A comparison of volume-based and surface-based multi-voxel pattern analysis. NeuroImage. 2011;56: 593–600.

40. Kirby KN, Gerlanc D. BootES: An R package for bootstrap confidence intervals on effect sizes. Behav Res. 2013;45: 905–927.

41. Hanke M, Halchenko YO, Sederberg PB, Hanson SJ, Haxby JV, Pollmann S. PyMVPA: A Python toolbox for multivariate pattern analysis of fMRI data. Neuroinformatics. 2009;7: 37–53.

42. Halchenko YO, Hanke M. Open is Not Enough. Let's Take the Next Step: An

Integrated, Community-Driven Computing Platform for Neuroscience. Front Neuroinform. 2012;6.