Title: Probabilistic mapping of human functional brain networks identifies regions of high group consensus

Authors: Ally Dworetsky[1,7], Benjamin A. Seitzman[2], Babatunde Adeyemo[2], Maital Neta[6], Rebecca S. Coalson[1,2], Steven E. Petersen[1-5], Caterina Gratton[7-9]

Affiliations:
Departments of Radiology[1], Neurology[2], Psychological and Brain Sciences[3], Neuroscience[4], and Biomedical Engineering[5] – Washington University School of Medicine
Department of Psychology[6] – University of Nebraska-Lincoln
Departments of Psychology[7], Neurology[8], Interdepartmental Neuroscience Program[9] – Northwestern University

Acknowledgements:
Funding was provided by NIH grant R01MH118370 (CG), the James S. McDonnell Foundation (SEP), and NIH grant R01MH111640 (MN).

ABSTRACT

Many recent developments surrounding the functional network organization of the human brain have focused on data that have been averaged across groups of individuals. While such group-level approaches have shed considerable light on the brain's large-scale distributed systems, they conceal individual differences in network organization, which recent work has demonstrated to be common and widespread. Here our goal was to leverage information about individual-level brain organization to identify locations of high inter-subject consensus. We probabilistically mapped 14 functional networks in multiple datasets with relatively high amounts of data. All networks show "core" (high-probability) regions, but differ from one another in the extent of their higher-variability components. These patterns replicate well across datasets with different scanning parameters. We produced a set of high-probability regions of interest (ROIs) from these probabilistic maps; these and the probabilistic maps are made publicly available, allowing researchers to apply information about group consistency to their own work in rest- or task-based studies.

INTRODUCTION

A key objective of functional magnetic resonance imaging (fMRI) studies has been to gain insight into how brain regions respond during tasks and how they interact with one another in distributed large-scale systems. To do so, both task- and connectivity-based analyses have typically been performed on averages across groups of subjects, to counteract noisy data from individuals. Studies using a group-average approach to examine human functional brain networks have produced robust and well-validated descriptions of, for example, typical functional network architecture (Power et al., 2011; Yeo et al., 2011).

Although the group-average brain has been useful in revealing fundamental qualities of functional network organization, recent data have suggested that averaging across subjects ignores distinct individual-specific features of cortical organization (Braga & Buckner, 2017; Finn et al., 2015; Gordon et al., 2017; Kong et al., 2019; Miranda-Dominguez et al., 2014; Mueller et al., 2013). Historically, a major barrier to producing reliable connectivity estimates at the individual level using resting-state functional connectivity (RSFC) techniques has been acquiring a sufficient quantity of data to counteract sampling variability (Gordon et al., 2017; Laumann et al., 2015). Previous work has demonstrated that the reproducibility of connectivity estimates and individual-specific features of functional brain networks is drastically improved with greater quantities of data per subject (Anderson et al., 2011; Elliott et al., 2019; Laumann et al., 2015; Noble et al., 2017). Accordingly, RSFC studies acquiring a typical 5-10 minutes of data per subject may not be sufficient to accurately reflect connectivity patterns in a given individual, or to examine individual differences in network organization. Several recent works have used higher reliability datasets to illuminate regions of high individual difference in functional network topography (Braga & Buckner, 2017; Gordon et al., 2017; Seitzman et al., 2019), outlining a geography of brain locations that show substantial variability across individuals.

Given the wealth of recent findings on individual variability in functional network architecture, an important question for group analyses emerges: how can we be confident that the regions we're measuring are functionally similar across individuals? To address this issue, our goal here was to quantify group *commonalities* (rather than individual differences) in network topography across individuals. Despite individual differences, past data have suggested that commonalities in network organization are also large and widespread, with many regions of the cortex showing substantial similarity to the typical group-average brain (Gratton et al., 2018; Kong et al., 2019; Seitzman et al., 2019). Given that most studies do not contain sufficient data to reliably map individual networks, the majority of studies have used group-average approaches in analysis. Thus, identifying regions that are consistent across individuals will determine (1) how to implement group average analyses to limit confounds from individual differences and (2) the extent to which past group average data can be extrapolated to single subjects.

In the present work, we aimed to address this need by mapping and quantifying areas of high group consensus: regions where the greatest group convergence in functional network organization is observed. Notably, we represent this group-consensus information probabilistically, and provide tools that can be directly applied in various experimental contexts to quantify the degree of certainty in network assignments. To create high-data estimates of group consensus, we focused our analyses on datasets with relatively higher amounts of data, where individual network maps achieve higher reliability. We used a template-matching procedure to identify cortical brain networks in higher-data individuals (N = 69 with > 20 min. of low-motion resting-state data per individual), and combined the resulting maps to produce a cortex-wide probabilistic estimate for each network. We replicated these findings in two other high-data datasets (the Midnight Scan Club: N = 9 with > 154 min., and the Human Connectome

Project: N = 99, with > 52 min.). Probabilistic maps are presented and quantified at various thresholds and are validated by contrasting to past independent results of high variability regions. Finally, we provide two tools for research use: (1) a set of network-specific, high-probability ROIs for use in task- and functional connectivity-based analyses and (2) a point-and-click tool allowing researchers to explore voxel-by-voxel probabilistic network estimates for regions of activation in their own data. The use of high-consensus regions may provide greater confidence in ROIs selected as priors in network-informed task and resting-state studies.

METHODS

Datasets and overview

Four independent datasets focused on young neurotypical populations were utilized in this paper (Table 1): a Washington University dataset (a subset of the participants reported in Power et al., 2012), a Dartmouth dataset (Gordon et al., 2016), the Midnight Scan Club (MSC) dataset (Gordon et al., 2017), and the Human Connectome Project (HCP) dataset (Van Essen et al., 2012). Notably, each dataset we use here consists of subjects with a relatively large amount of low-motion data, ranging from a minimum of 20 min. (for N = 69 in the Dartmouth primary mapping dataset) to upwards of 154 min. (for N = 9 in the MSC replication dataset). This large amount of data dramatically increases the reliability of functional connectivity measurements relative to more typical 5-10 min. scans (Gordon et al., 2017; Laumann et al., 2015).

Table 1: Datasets and data details. Low-motion data quantities were measured after correction for movements using framewise displacement (see *Supp. Methods*).

Dataset	**N**	**Usage**	**Amount of data collected per subject (min. – max.)**	**Amount of low-motion data retained per subject (min. – max.)**
WashU-120	120	Group-average, network discovery	7.6 – 30.2 minutes	6.3 – 29.7 minutes
WashU-24 (higher-data subset of WashU- 120)	24	Network template creation	44 – 133 minutes	35 – 124 minutes
Dartmouth	69	Template-matching, probabilistic maps	21– 60 minutes	20 – 49 minutes
MSC	9	Replication of probabilistic maps	300 minutes	154 – 281 minutes
HCP	99	Replication of probabilistic maps	60 minutes	52 – 57 minutes

The WashU datasets were used to generate group-average cortical network classifications based on networks typically found in group decompositions (Laumann et al., 2015; Power et al., 2011), and to create network templates derived from these descriptions. The larger dataset (the "WashU-120"; 60 female, average age 24.7 years) was used for network discovery, and

included subjects ranging from 6.3 to 29.7 min. of low-motion data (see *Supplementary Materials* for detailed information). A high-quality group-average connectivity template for each of the 14 canonical brain networks described in this paper was generated using a subject "subset" of the WashU-120, consisting of 24 high-data subjects (the "WashU-24"). Subjects in this subset had at least 35 min. of low-motion data when combining across additional resting-state scan sessions previously obtained from our group (see *Template Generation* below for more details). These group-average templates were then applied to subjects in the Dartmouth dataset to identify brain networks in single individuals.

The Dartmouth dataset (N = 69 subjects [56 female; average age 20.2 years]) included subjects with over 20 min. of low-motion data. Given its large sample size and its standard, single-band scanning parameters, this dataset was the primary dataset used to determine network probabilities across individuals and generate network-specific regions of high inter-subject consensus.

The MSC dataset (N = 9 subjects [4 female; average age 29.3 years] with over 154 min. of low-motion rest data) and subjects from the HCP dataset (N = 99 subjects [52 female; average age 28.2 years] with at least 52 minutes of data) were used as supplemental datasets to replicate findings on the distribution of high-probability brain regions across individuals. Notably, the MSC dataset includes very highly-sampled individuals whose functional connectivity maps have been demonstrated to have high reliability and validated with functional activation studies. The HCP dataset replicates the current findings in a large dataset at high spatial and temporal resolution. See Supp. Table 1 for acquisition parameters for functional data across all datasets, and see *Supplementary Materials* for details on all preprocessing and functional connectivity (FC) processing procedures.

Template-matching and generation of high-probability ROIs

In this work, we created network maps for high-data individual subjects using a template matching approach. These network maps were then overlaid to generate a probabilistic estimate of network distributions across subjects. High-consensus ROIs were generated for research use from regions of high cross-subject agreement of network assignment. Procedures for template-matching in individuals and probabilistic network map generation are illustrated in Fig. 1 and described in more detail below.

1. Template Matching

Brain networks were identified in individual subjects by a winner-take-all procedure (similar to that employed in Gordon et al., 2017) which assigned each cortical gray matter voxel in a particular subject to one of 14 network templates. Networks used in this analysis are previously described in (Gordon et al., 2017) and (Laumann et al., 2015) and include the default mode (DMN), visual, fronto-parietal (FP), dorsal attention (DAN), ventral attention (VAN), salience, cingulo-opercular (CO), somatomotor dorsal (SMd), somatomotor lateral (SMl), auditory, temporal pole (Tpole), medial temporal lobe (MTL), parietal medial (PMN), and parieto-occipital (PON). Rather than using a data-driven community detection approach to map individualized networks, this template-matching approach was chosen based on our goal of investigating known, previously-described brain networks to allow for a reliable comparison of network structure across individuals. Template generation and individual matching are described in more detail below.

1.1. Template generation

Templates were generated based on a map of group-average network assignments in the cortex using the WashU-120 dataset. This analysis was performed in volume space with 3 mm isotropic voxels after applying a cortical gray matter mask. A voxelwise correlation matrix was calculated by correlating each cortical voxel's BOLD timeseries with the timeseries of every other cortical voxel, resulting in a 57,544 by 57,544 matrix. Correlation matrices were Fisher transformed and averaged across subjects. The inverse Fisher transform was then applied to the resulting group-average matrix.

Next, the data-driven community detection InfoMap algorithm (Rosvall & Bergstrom, 2008) was applied to the group average matrix to identify cortical brain networks across a range of edge density thresholds (from 0.5% to 5% by increments of 0.1%; closely following Gordon et al., 2017). Thresholding was applied after excluding edges to nodes within 20 mm from a given voxel. Fourteen canonical consensus networks were defined by an automated algorithm (Seitzman et al., 2019), which summarizes assignments by weighting communities across thresholds. The algorithm allows smaller networks to contribute more heavily by allotting greater weight to the sparser thresholds at which those communities are commonly detected.

A second layer of averaging was then applied, in this case using the higher-sampled WashU-24. For each of these subjects, the mean timeseries of all voxels labeled as a given network in the group-average consensus map was extracted. This average timeseries was then correlated with all other gray matter voxel timeseries, resulting in a network seedmap, which was averaged across subjects. This procedure was repeated for each network, producing a group-average template connectivity map for each brain network (see Supp. Fig.1 for a schematic of the creation of network templates and a view of all network templates).

1.2. Applying a template-matching procedure to identify networks in individuals

After the group-average templates were generated, they were used to identify individual specific network assignments in the primary Dartmouth dataset, as well as the supplementary datasets for replication. Specifically, for each individual subject, a voxelwise correlation matrix was generated by correlating each voxel's BOLD timeseries with the timeseries of every other voxel in the cortical mask, resulting in a 57,544 by 57,544 correlation matrix per subject.
This matrix was binarized to the top 5% of connectivity values across voxels. Each voxel's binarized map was iteratively compared with the 14 (also binarized, see *Supp. Fig. 1*) network templates and matched to its "best fit," resulting in a cortex-wide, individual-specific network map. Fit was measured using the Dice coefficient of overlap between the binarized voxel connectivity map and each binarized template map (Fig. 1A). This procedure was repeated for all voxels in the individuals cortical mask (Fig. 1B). Same-network clusters of less than 108 mm^3 (4 contiguous voxels) were removed from each individual's network map.

2. Creating network-specific probabilistic maps

After individual specific network maps from the Dartmouth dataset had been generated with the template mapping procedure, these maps were overlapped to produce a cross-subject probabilistic map for each network (Fig. 1C). To generate this cross-subject probabilistic map, individual network assignments at each brain location were tallied to calculate the total occurrence (in number of subjects with a given network assignment out of the total N = 69). This produced a continuous probabilistic map for each network which specified the probability of a given network assignment at every voxel within the cortical mask. Frequency values of network

assignments were divided by the number of subjects within the primary dataset and were converted to percentages to illustrate the probability of network membership at each voxel. Probabilistic maps were created in the same manner from both the MSC and HCP datasets based on the number of subjects included (9 and 99, respectively), and were compared to the results from the primary dataset. Thresholded versions of the network-specific probabilistic maps were also produced (Fig. 1D), allowing for visualization of the network assignment frequencies at various probability thresholds (e.g., in 50, 60, 70, 80, or 90 percent of subjects).

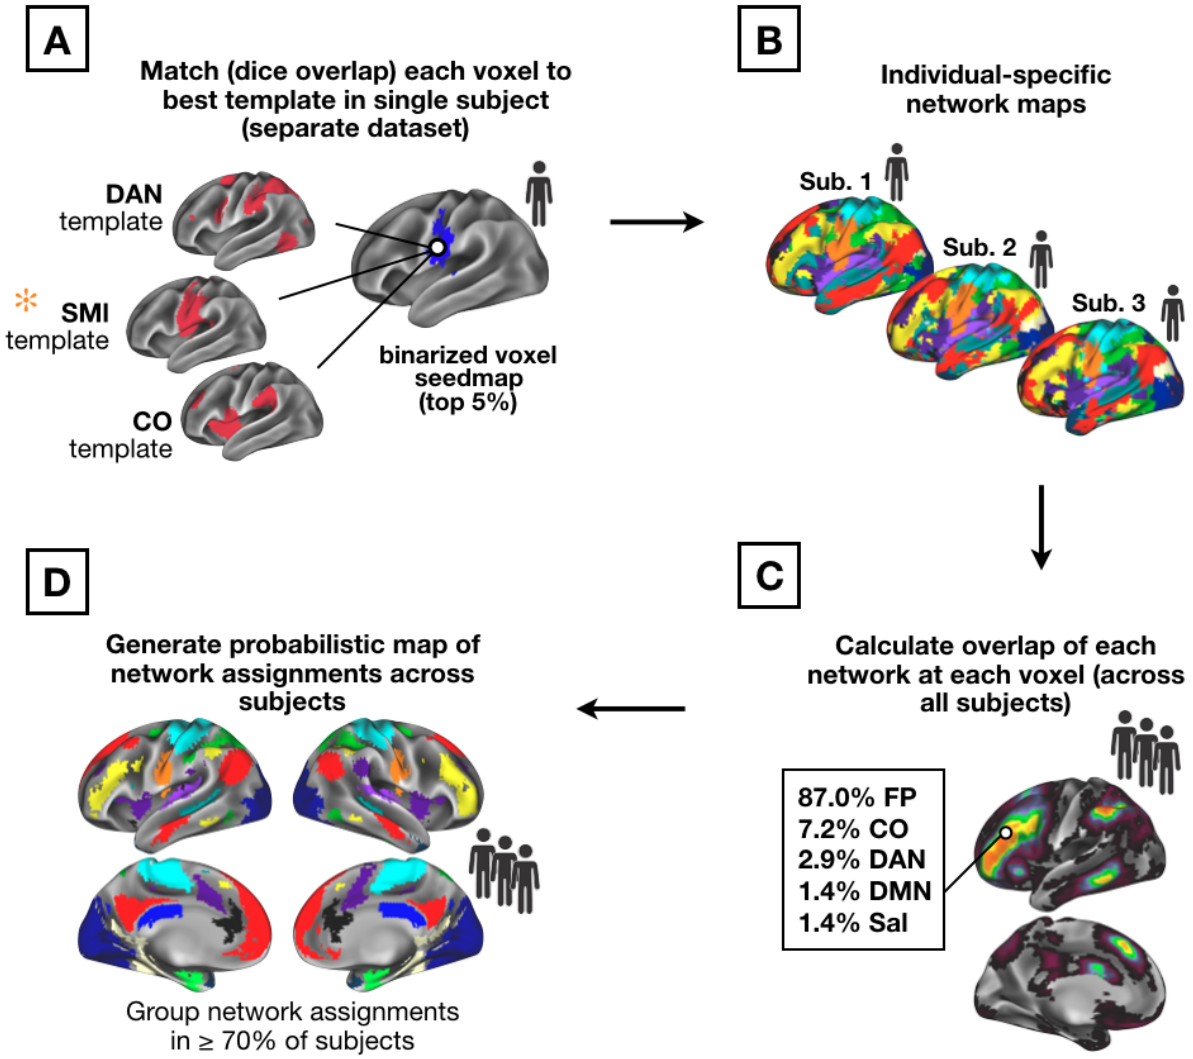

Figure 1: Template-matching procedure (*A*) and creation of probabilistic network maps (*B-D*). A set of group-average network templates were created from the WashU dataset. These group-average templates were binarized at the top 5% of connectivity values. Next, for each single individual in the Dartmouth dataset, a voxelwise correlation matrix was calculated between all gray matter voxels. Seedmaps were thresholded at the top 5% of values across voxels. The individual's voxel-level binarized map was then iteratively compared (by Dice overlap) with each group-average network template (also thresholded to the top 5% of values), and the network with the highest Dice coefficient was assigned to the voxel (*A*). Once all voxels were assigned in all subjects (*B*), the number of network assignments at each voxel were tallied across subjects (*C*) to generate probabilistic maps of networks. These probabilistic maps were then thresholded (*D*) to represent locations with network consensus in a large majority of subjects. Note that while all steps were performed in volume (Talairach) space, results are mapped onto a template surface for visualization purposes only.

3. Creating ROIs of high group consistency for studies in other modalities

Once probabilistic maps were defined, we next set out to create a set of regions of interest (ROIs) with high group consensus for use in future (and retrospective) studies. These ROIs were created by contrasting the probabilistic maps generated above from the Dartmouth dataset with 248 (of 264) ROIs of the larger set previously proposed in Power et al. (2011) found in the cerebral cortex.

Specifically, high group consensus regions were derived from the probabilistic maps of the Dartmouth dataset by identifying locations that showed consistent network assignments across a large majority (i.e., > 75%) of subjects. A spherical 7 mm diameter region was placed on each of the center coordinates reported in Power et al. (2011). ROIs were identified as "high-probability" if their average probability (across voxels) was ≥ 75%. If a region failed to meet the 75% criteria to be identified as "high-probability," it was shifted one voxel in space (i.e., 3 mm in the *x*, *y*, or *z* direction) and was retained if this shift produced an average ROI probability that met the threshold, resulting in a total of 44 ROIs shifted from their original position. ROIs that failed to meet the high consensus definition with a single voxel shift were dropped from the final group.

4. Creating point and click voxel-wise network tool

Finally, we created a tool for displaying the probability of network membership at each cortical voxel for research use. Specifically, a *scene* file was created using the Connectome Workbench software that contains each network's probabilistic map in volume space and allows for point-and-click usability to identify the probability (across subjects) that a given voxel is associated with each network.

RESULTS

Overview of results

In this work, we sought to characterize high-consensus network locations for use in analysis and interpretation of group research studies. To this end, we compiled network assignments in a large sample of relatively high-data subjects to create a reliable cross-subject probabilistic map of network definitions. We replicated these findings in two additional datasets to demonstrate their stability. Using these datasets, we explored the degree of consensus for networks across different probability thresholds. Finally, we created two tools for use in future research studies: (1) a set of "high-probability" regions of interest, and (2) a publicly available point-and-click tool for determining network probabilities in researcher-specified locations.

Estimated probabilistic maps of 14 canonical networks

As described in the Methods, we used a template-matching approach to determine a voxel-based network assignment for each individual (N = 69) in our primary Dartmouth dataset, based on templates created from the WashU cohort. We then overlapped the individuals' network maps within the Dartmouth dataset for each canonical network. This overlap was used to generate a cross-subject probabilistic map (Fig. 2; see Supp. Fig. 2 for maps for the remaining 8 canonical networks we examined). As can be seen, all networks demonstrated some regions with high-group consensus (warm colors), but also a spread of lower-consensus locations.

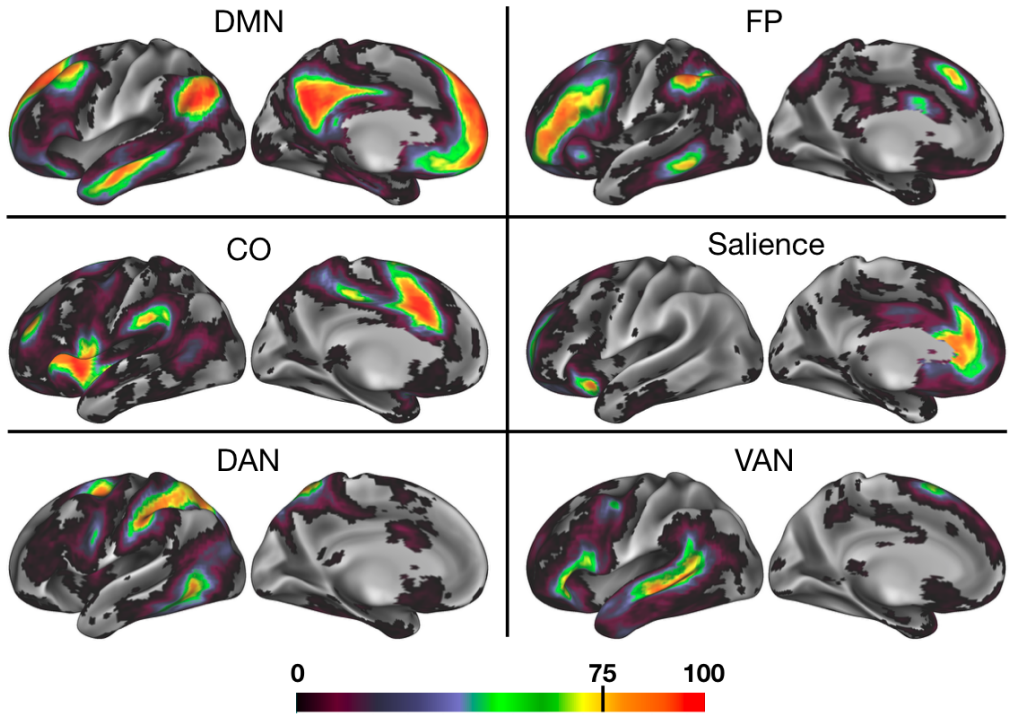

Figure 2: A probabilistic representation of 6 association networks. Cooler colors represent regions with the least confidence in network assignment across subjects, while warmer colors represent brain regions with the highest group consistency– in bright red regions, up to 100% of subjects converged on a given network assignment. (See Supp. Fig. 2 for probabilistic maps produced for both hemispheres and for all 14 networks.)

Consensus locations replicate across multiple datasets

Next, we implemented the probabilistic map procedure in two supplemental datasets (consisting of 9 MSC subjects and 99 HCP subjects). This analysis yielded results that largely replicated the findings from the Dartmouth dataset (Fig. 3). Despite differences in participant populations, scanners, and acquisition parameters (most notably in the HCP dataset; see *Supplementary Materials* for further information), probabilistic network assignments generally replicate, with results from the three test datasets visually appearing similar at the 50 percent probability threshold and experiencing similar patterns of network "dropout" as the probability of assignment increases at 70 and 90 percent. We note that more dropout is observed in the HCP dataset, perhaps due to the lower SNR associated with these scans (Ji et al., 2019; Seitzman et al., 2019)

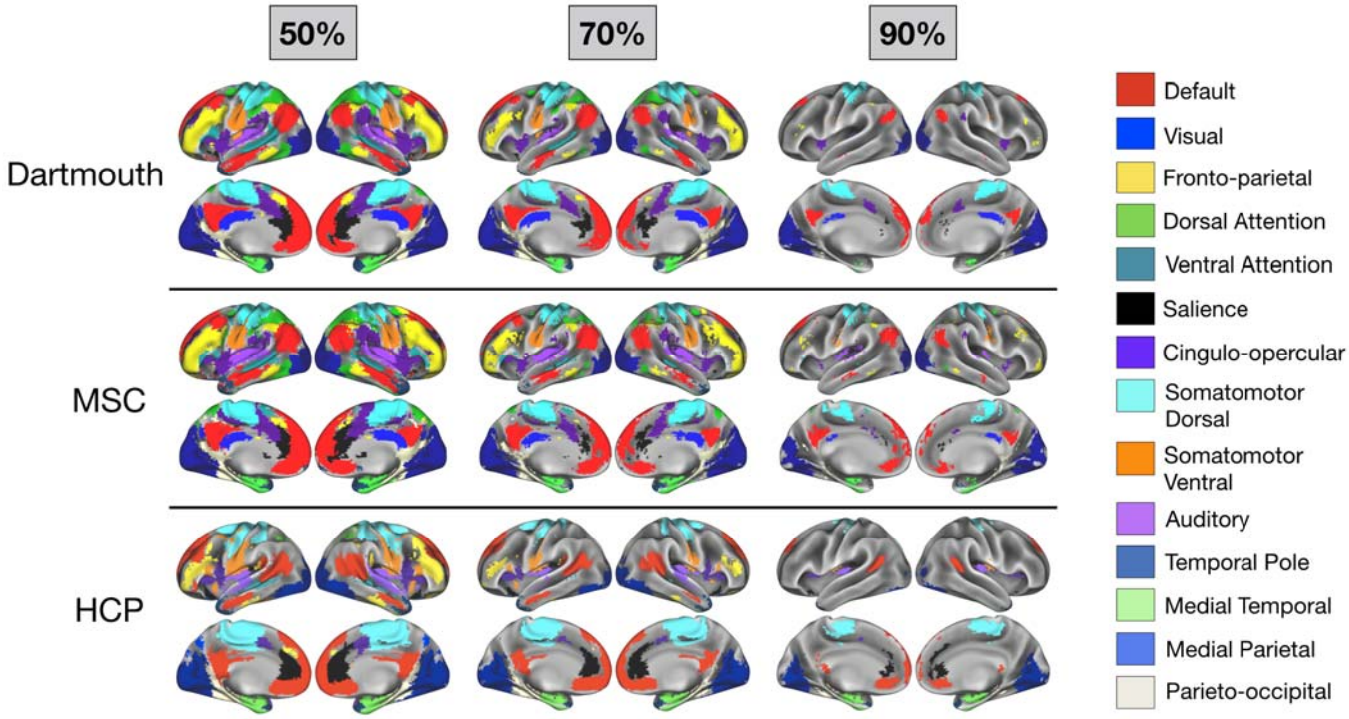

Figure 3: Thresholded probabilistic maps across 3 datasets. Probabilistic maps were generated from two additional datasets (MSC and HCP). For each dataset, network assignments consistent across 50%, 70%, and 90% of subjects are displayed.

Individual networks vary in the size of their core and the span of surrounding components

While core regions of high consensus exist in all of the canonical networks investigated here, the networks vary in the extent of their more peripheral (i.e., low consensus) regions. As shown in Fig. 4, networks retain cortical territory at varying rates as the probability threshold (i.e., consensus across subjects) increases. For example, while the visual network consistently remains the most highly-represented network across the highest probability thresholds, the inverse is true for FP: it is the third-most highly represented network across at least 50% of individuals, but when group consensus is examined at 80% or 90% of individuals, cortical representation of FP diminishes significantly.

Differences in the rate of network “dropout” seem not to be driven purely by a distinction of sensorimotor vs. association networks. While sensorimotor networks tend to have higher consensus, some association networks also maintain a relatively high group consensus across thresholds, including DMN and CO. It appears unlikely that network size alone is driving the effect (i.e., that smaller networks taper off more quickly across probability thresholds); while some smaller networks experience relatively fast dropout (e.g., VAN), others (e.g., PON and MTL) remain consistent across a high percentage of subjects. Regardless, all networks have some core regions of high inter-subject consensus, and networks vary in the cross-subject variability observed in locations surrounding these core regions.

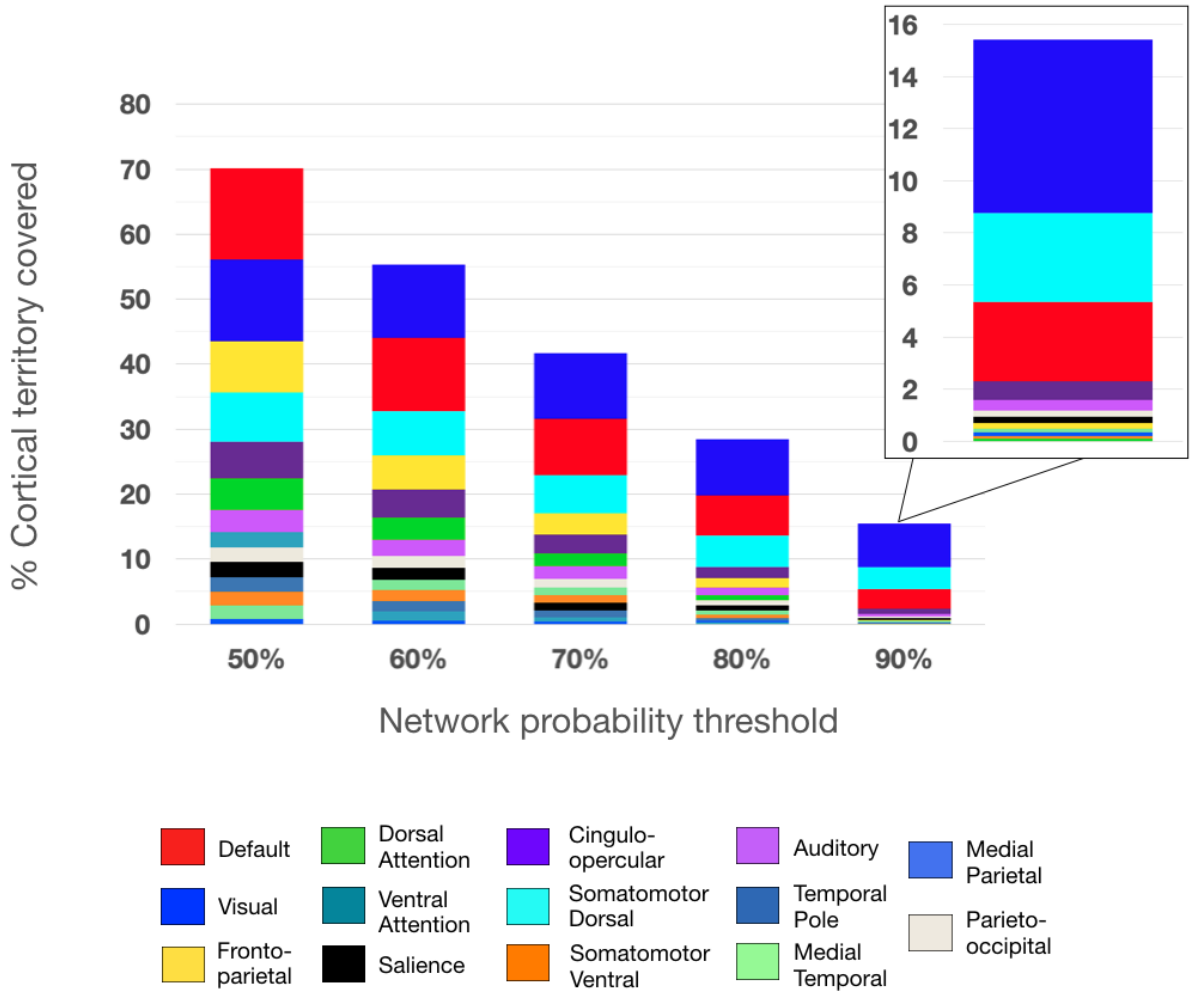

Figure 4: Representation of the proportion of cortical territory covered by each network at each probability threshold. Each bar represents the total percentage of the cortex covered at a given probability threshold, and stacks represent network coverage. Stacks are ordered by most to least coverage from top to bottom.

Non-core areas overlap with previously-described locations of network variants

Next, we sought to confirm our findings by examining how consensus regions compared with previously identified locations of individual variability in functional network organization. Transparent white regions in Fig. 5 show “network variant” locations across 752 HCP subjects from (Seitzman et al., 2019) where a given individual’s correlation pattern differs significantly from the group-average. Despite differences in methodological approaches used for identifying consensus probabilistic assignments (via template-matching) and individual variants (via low spatial correlations), we find that there is a good contrasting correspondence between these two methods. As would be expected, regions of high consensus lie mostly outside of the boundaries of network variants, and appear to fill in gaps where there is the greatest inter-subject variability in network assignment (e.g., temporoparietal junction, lateral frontal cortex).

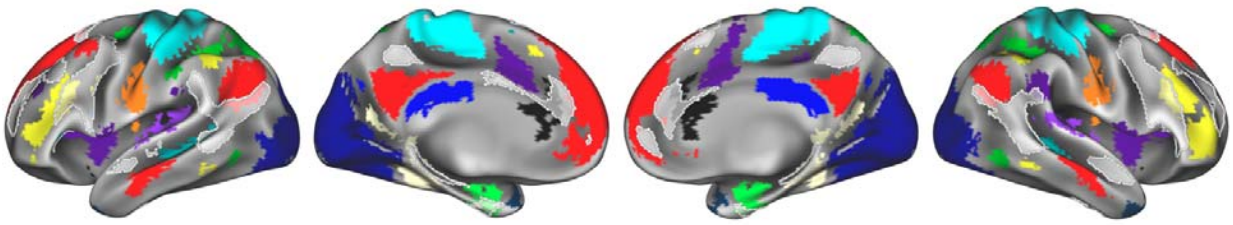

Figure 5: The spatial distribution of network variants across 752 HCP subjects (as identified in Seitzman et al., 2019) is displayed in transparent white, overlaid on the network map at 75% probability. The distribution displayed here is thresholded to show variant locations exhibited by at least 11% of subjects. Notably, the variants distribution appears to fill in gaps where there is the most inter-subject variability in network assignment, including temporoparietal junction and the left and right lateral frontal cortex.

Generation of a high-probability set of ROIs and point-and-click tool

We sought to refine previous group-average ROI definitions based on these probabilistic network assignments to generate a set of high-consensus ROIs for future research. To this end, we began with the 248 cortical ROIs from the commonly used 264 regions from Power et al. (2011). We then restricted this set to regions where the average network assignment probability was ≥ 75% within the 7 mm diameter ROI. This resulted in 153 cortical ROIs. Thirteen of the 14 canonical networks were represented (no ROIs were retained for the temporal pole network), although the quantity of high-probability ROIs kept varied by network (see Fig. 6A for locations and network descriptions of ROIs). While the regions cover much of the cortex, some higher-variability areas such as the temporo-parietal junction and the lateral frontal cortex are more sparsely represented, as expected (e.g., see Gordon et al., 2017; Laumann et al., 2015; Mueller et al., 2013; Seitzman et al., 2019). ROIs with the highest peak probabilities were identified largely in dorsal somatomotor and visual regions, with relatively lower peaks in lateral frontal and orbitofrontal regions (Fig. 6B).

Additionally, we provide each network's probabilistic maps as a series of downloadable volume images, available with the 153 ROIs at https://github.com/GrattonLab. For researchers using Connectome Workbench, a *scene* file was created to allow researchers to explore network probabilities at every cortical voxel. Fig. 7 displays an example of this tool's utility by exploring the DMN map, including an "Information" window with probabilities listed across all networks.

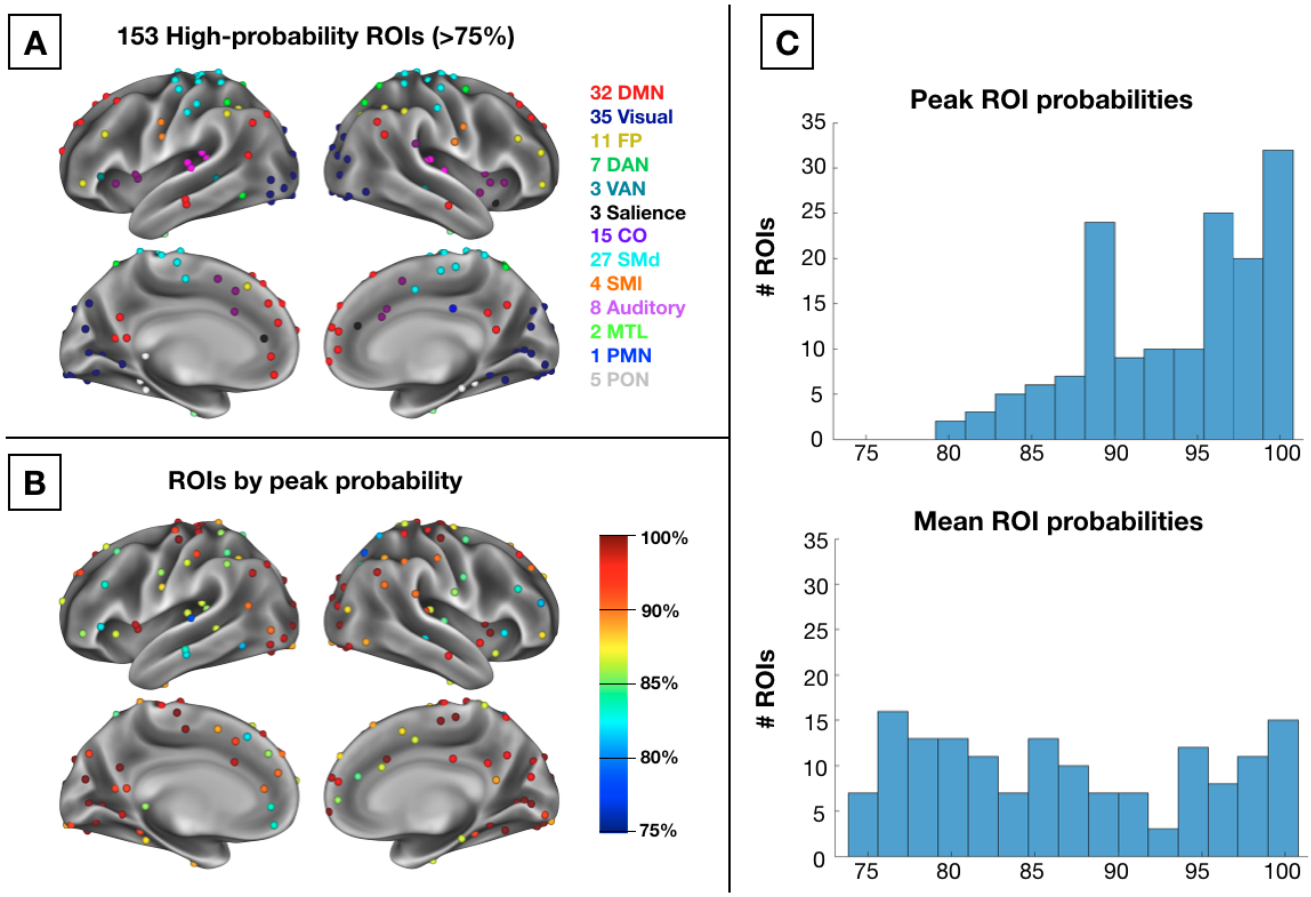

Figure 6: (A) 153 high-probability ROIs colored by their network description. (B) ROIs colored by peak probability across voxels within the ROI. (C) Histograms of peak probability values across all 153 ROIs (top) and mean probability values (bottom).

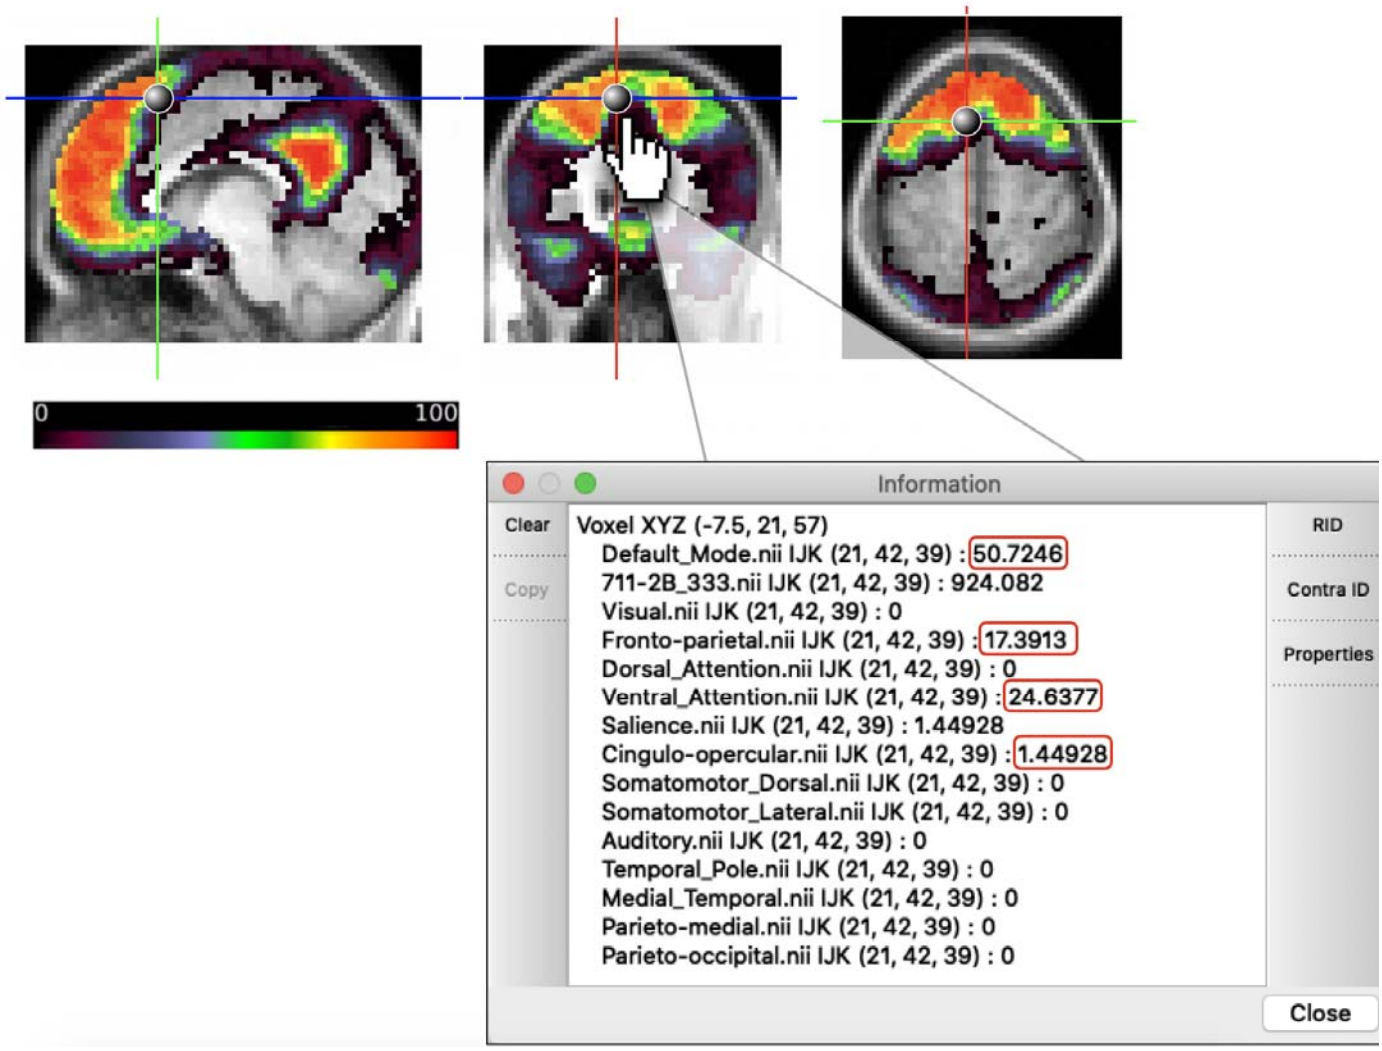

Figure 7: Schematic of publicly-available research tool for exploring network probabilities. The DMN map is displayed, and probabilities of network membership to all 14 networks for the given voxel are listed in the "Information" window with non-zero probabilities outlined in red.

DISCUSSION

Here, we probabilistically mapped functional networks across a group of high-data individuals. We found that there are "core" locations of high group consensus within each network. Networks vary in the extent and peak probability of their core regions, suggesting that networks with a higher group consensus may be more amenable to group-level analyses. The ability to identify locations with high group consensus allows for better-informed group studies, using either task or functional connectivity approaches. To facilitate this process, we provide a set of voxelwise probability maps for each of 14 canonical networks. In addition, we provide two tools for research use: (1) a set of network-specific, high-probability ROIs for use in task- and functional connectivity-based analyses and (2) a point-and-click tool allowing researchers to explore voxel-by-voxel probabilistic network estimates for regions of activation in their own data.

Probabilistic approaches in imaging

In imaging literature, probabilistic atlases have been utilized as a way to quantify spatial distributions of anatomical structures or functional areas to pinpoint locations of high consensus across a group. Many popular probabilistic atlases of the brain are based on anatomical data – e.g., the cerebellum (Diedrichsen et al., 2009), subcortical nuclei (Pauli et al., 2018), the basal ganglia (Keuken & Forstmann, 2015), tissue type, lobes, and sulci (Mazziotta et al., 1995) – to provide references for cross-subject comparisons. However, functional areas (at least in the cortex) do not necessarily conform well to anatomical definitions (Eickhoff et al., 2018; Gordon et al., 2016), suggesting that anatomical atlases are less well-suited for definition of functional ROIs in task-based or resting-state fMRI. The current cortical probabilistic atlas based on functional network mapping fills this gap.

Utility of a probabilistic mapping approach to functional networks

We have adopted a probabilistic approach in this study given past evidence for both individual differences and group consensus in functional neuroanatomy (Gordon et al., 2017; Gordon et al., 2017; Gratton et al., 2018; Power et al., 2011; Yeo et al., 2011) . It became increasingly apparent in our own work that, rather than qualitative statements about the magnitude of variability or the extent of similarity, it would be useful to have a quantitative probabilistic view of the variability associated with each cortical location and each network to evaluate the consistency of findings. Identifying regions of group consensus provides a wealth of opportunity for more well-informed future research on brain networks.

The network probabilistic maps that we have produced can be used by researchers in a number of ways. First, these maps can be thresholded to create regions of interest for future (or retrospective) group analyses. For example, these maps may be thresholded to select regions of the frontoparietal and cinguloopercular system where we have high confidence in group consensus. This would allow for a re-analysis of past task dissociations between these two systems (Dubis et al., 2016; Gratton et al., 2017; Neta et al., 2015; Power & Petersen, 2013), but now accounting for potential individual variability in network assignments. To facilitate this application, we have provided a set of 153 ROIs that identify high-consensus regions within 13 of the 14 networks examined. Should a researcher wish to perform task-based or rest-based analyses at the group level, high-probability ROIs would be crucial in ensuring that the brain regions being analyzed are those which are most consistent across individuals; researchers can be more assured that a majority of individuals are providing information from the same network.

Secondly, these regions may be used to help interpret ambiguous results in group studies. For instance, a region which is assigned to the CO network in 70% of subjects and the salience network in the other 30% of subjects may serve as a meaningful distinction from a region which is assigned to CO in 70% of subjects but FP in the other 30%. Thus, while the high-probability ROIs focus on regions of group similarity, useful information on the locations and forms of individual variability can also be gleaned from the point and click probabilistic tool. In the future, researchers may use the probabilities associated with this paper to provide quantitative estimates for the typical (and atypical) network assignments associated with findings of interest.

Third, probabilistic network mapping may deepen our understanding of the clinical utility of mapping functional brain networks by providing reliable quantitative priors about the network assignments of each region. This probabilistic approach may provide a basis for more precisely identifying network deviations in individuals with specific diagnoses, as well as network changes across development. For example, one possible future investigation may be to examine whether

individuals with a given clinical diagnosis vary predictably from the probability map of any network of interest; perhaps in clinical groups there will be more variability in higher-probability regions.

Group consensus in core regions within large-scale networks

Our probabilistic maps demonstrated that each network was comprised of a set of "core" regions exhibiting very little or, in some cases, no variability (*note that our use of "core" is based on anatomical location, separate from the graph theoretical connotation of the word*). This suggests that the core areas of each network are relatively fixed across individuals, with little possibility for variation, and these regions complement previously described locations of high individual variability (see Fig. 5). The consensus areas of each network were larger in sensorimotor than association systems, consistent with the idea that association systems are more variable across individuals, maturation, and evolution, which has been suggested to be due to a lack of genetically encoded tethering markers in these areas (Buckner & Krienen, 2013). However, we found a consistent core in each of the association systems as well, which would appear to be at odds with a strong interpretation that association networks lack fixed constraints (Buckner & Krienen, 2013). Indeed, the consistency of association networks differed markedly between systems with, e.g., relatively robust consensus in the DMN and CO and high variability in the FP, despite their similar overall sizes and complex "high level" natures. Exploring the basis for commonalities and plasticity in association networks will be an interesting avenue for future work.

Limitations

The findings presented here have several limitations that are worth noting. First, in an effort to optimize the tradeoff between data quantity and the number of subjects retained for our probabilistic estimates, the amount of data required per subject was set to a minimum of 20 minutes of low-motion data. While this represents relatively higher-data subjects than a majority of group studies (which collect 5-10 min. of data), most of these subjects did not reach the 30-45 minute threshold that is ideal to produce asymptotic individual-subject reliability (Laumann et al., 2015),. However, we were able to repeat the probabilistic analyses within the smaller but high-data Midnight Scan Club (MSC) dataset, which produced comparable results.

Further, while there was general agreement in the probabilistic maps from the 3 datasets examined here, there were some differences, which may be driven by differences in scan parameters or dataset size/quality. This was particularly the case in the probabilistic map generated from the Human Connectome Project (HCP) dataset relative to the other two datasets. The probabilistic maps displayed in Fig. 3 reveal that some high-consensus regions that are conserved across probability thresholds in the Dartmouth and MSC datasets show a lower degree of consensus in the HCP dataset. Such differences might be driven by the smaller voxel size and higher spatial and temporal resolution of the HCP dataset, which may lead to a lower signal-to-noise ratio (SNR). Thus, the extent to which the probabilistic assignments replicate in datasets using similar acquisitions as the HCP is less certain, and may require further investigation. Given this potential concern, we have also separately released the HCP-specific probabilistic network maps to allow those using acquisition parameters similar to the HCP dataset to have a reference for comparison.

Lastly, we note that probabilistic assignments were calculated at the level of 14 canonical functional networks and not at an areal level. A consistent network assignment across individuals is not a guarantee that a region belongs to the same brain area across those

individuals. We know from past work based on functional localizers that there is variability across subjects at the areal level as well (e.g., Kanwisher et al., 1997; Wang et al., 2015); variation at the areal level may also carry information about individual differences, and will be important in studies requiring area-level precision. Still, an understanding of the probabilistic composition of functional networks (particularly association networks in frontal and parietal cortex, in light of the more prominent individual variability that has been observed there) will serve as an important improvement to the interpretation of many task and resting-state studies.

CONCLUSIONS

Here, we produce a probabilistic representation of distributions of functional network assignments across a group of high-data subjects. While individual networks vary in the span of their “core” high-probability locations, all networks examined showed regions of high group consensus, and core regions replicated reasonably across datasets. High-probability ROIs produced from this analysis will allow researchers to interrogate task or resting-state data in locations that have been demonstrated to show high degrees of group consensus.

REFERENCES

Anderson, J. S., Ferguson, M. A., Lopez-Larson, M., & Yurgelun-Todd, D. (2011). Reproducibility of single-subject functional connectivity measurements. *American Journal of Neuroradiology*. https://doi.org/10.3174/ajnr.A2330

Braga, R. M., & Buckner, R. L. (2017). Parallel Interdigitated Distributed Networks within the Individual Estimated by Intrinsic Functional Connectivity. *Neuron*. https://doi.org/10.1016/j.neuron.2017.06.038

Buckner, R. L., & Krienen, F. M. (2013). The evolution of distributed association networks in the human brain. In *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2013.09.017

Diedrichsen, J., Balsters, J. H., Flavell, J., Cussans, E., & Ramnani, N. (2009). A probabilistic MR atlas of the human cerebellum. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2009.01.045

Dubis, J. W., Siegel, J. S., Neta, M., Visscher, K. M., & Petersen, S. E. (2016). Tasks Driven by Perceptual Information Do Not Recruit Sustained BOLD Activity in Cingulo-Opercular Regions. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhu187

Eickhoff, S. B., Yeo, B. T. T., & Genon, S. (2018). Imaging-based parcellations of the human brain. In *Nature Reviews Neuroscience*. https://doi.org/10.1038/s41583-018-0071-7

Elliott, M. L., Knodt, A. R., Cooke, M., Kim, M. J., Melzer, T. R., Keenan, R., Ireland, D., Ramrakha, S., Poulton, R., Caspi, A., Moffitt, T. E., & Hariri, A. R. (2019). General functional connectivity: Shared features of resting-state and task fMRI drive reliable and heritable individual differences in functional brain networks. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2019.01.068

Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., Papademetris, X., & Constable, R. T. (2015). Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nature Neuroscience*. https://doi.org/10.1038/nn.4135

Gordon, E. M., Laumann, T. O., Adeyemo, B., Gilmore, A. W., Nelson, S. M., Dosenbach, N. U. F., & Petersen, S. E. (2017). Individual-specific features of brain systems identified with resting state functional correlations. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2016.08.032

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*. https://doi.org/10.1093/cercor/bhu239

Gordon, E. M., Laumann, T. O., Adeyemo, B., & Petersen, S. E. (2017). 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

Gordon, E. M., Laumann, T. O., Gilmore, A. W., Newbold, D. J., Greene, D. J., Berg, J. J., Ortega, M., Hoyt-Drazen, C., Gratton, C., Sun, H., Hampton, J. M., Coalson, R. S., Nguyen, A. L., McDermott, K. B., Shimony, J. S., Snyder, A. Z., Schlaggar, B. L., Petersen, S. E., Nelson, S. M., & Dosenbach, N. U. F. (2017). Precision Functional Mapping of Individual Human Brains. *Neuron*. https://doi.org/10.1016/j.neuron.2017.07.011

Gratton, C., Neta, M., Sun, H., Ploran, E. J., Schlaggar, B. L., Wheeler, M. E., Petersen, S. E., & Nelson, S. M. (2017). Distinct Stages of Moment-to-Moment Processing in the Cinguloopercular and Frontoparietal Networks. *Cerebral Cortex (New York, N.Y.□: 1991)*.

https://doi.org/10.1093/cercor/bhw092

Gratton, Caterina, Laumann, T. O., Nielsen, A. N., Greene, D. J., Gordon, E. M., Gilmore, A. W., Nelson, S. M., Coalson, R. S., Snyder, A. Z., Schlaggar, B. L., Dosenbach, N. U. F., & Petersen, S. E. (2018). Functional Brain Networks Are Dominated by Stable Group and Individual Factors, Not Cognitive or Daily Variation. *Neuron*. https://doi.org/10.1016/j.neuron.2018.03.035

Ji, J. L., Spronk, M., Kulkarni, K., Repovš, G., Anticevic, A., & Cole, M. W. (2019). Mapping the human brain's cortical-subcortical functional network organization. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2018.10.006

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*. https://doi.org/10.1523/jneurosci.17-11-04302.1997

Keuken, M. C., & Forstmann, B. U. (2015). A probabilistic atlas of the basal ganglia using 7 T MRI. *Data in Brief*. https://doi.org/10.1016/j.dib.2015.07.028

Kong, R., Li, J., Orban, C., Sabuncu, M. R., Liu, H., Schaefer, A., Sun, N., Zuo, X. N., Holmes, A. J., Eickhoff, S. B., & Yeo, B. T. T. (2019). Spatial Topography of Individual-Specific Cortical Networks Predicts Human Cognition, Personality, and Emotion. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhy123

Laumann, T. O., Gordon, E. M., Adeyemo, B., Snyder, A. Z., Joo, S. J., Chen, M. Y., Gilmore, A. W., McDermott, K. B., Nelson, S. M., Dosenbach, N. U. F., Schlaggar, B. L., Mumford, J. A., Poldrack, R. A., & Petersen, S. E. (2015). Functional System and Areal Organization of a Highly Sampled Individual Human Brain. *Neuron*. https://doi.org/10.1016/j.neuron.2015.06.037

Mazziotta, J. C., Toga, A. W., Evans, A., Fox, P., & Lancaster, J. (1995). A probabilistic atlas of the human brain: theory and rationale for its development. The International Consortium for Brain Mapping (ICBM). *NeuroImage*.

Miranda-Dominguez, O., Mills, B. D., Carpenter, S. D., Grant, K. A., Kroenke, C. D., Nigg, J. T., & Fair, D. A. (2014). Connectotyping: Model based fingerprinting of the functional connectome. *PLoS ONE*. https://doi.org/10.1371/journal.pone.0111048

Mueller, S., Wang, D., Fox, M. D., Yeo, B. T. T., Sepulcre, J., Sabuncu, M. R., Shafee, R., Lu, J., & Liu, H. (2013). Individual Variability in Functional Connectivity Architecture of the Human Brain. *Neuron*. https://doi.org/10.1016/j.neuron.2012.12.028

Neta, M., Miezin, X. M., Nelson, S. M., Dubis, J. W., Dosenbach, N. U. F., Schlaggar, B. L., & Petersen, S. E. (2015). Spatial and temporal characteristics of error-related activity in the human brain. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.1313-14.2015

Noble, S., Spann, M. N., Tokoglu, F., Shen, X., Constable, R. T., & Scheinost, D. (2017). Influences on the Test-Retest Reliability of Functional Connectivity MRI and its Relationship with Behavioral Utility. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhx230

Pauli, W. M., Nili, A. N., & Michael Tyszka, J. (2018). Data Descriptor: A high-resolution probabilistic in vivo atlas of human subcortical brain nuclei. *Scientific Data*. https://doi.org/10.1038/sdata.2018.63

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2011.10.018

Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., Vogel, A. C., Laumann, T. O., Miezin, F. M., Schlaggar, B. L., & Petersen, S. E. (2011). Functional Network Organization of the Human Brain. *Neuron*. https://doi.org/10.1016/j.neuron.2011.09.006

Power, J. D., & Petersen, S. E. (2013). Control-related systems in the human brain. In *Current Opinion in Neurobiology*. https://doi.org/10.1016/j.conb.2012.12.009

Rosvall, M., & Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences of the United States of America*. https://doi.org/10.1073/pnas.0706851105

Seitzman, B.A., Gratton, C., Laumann, T. O., Gordon, E. M., Adeyemo, B., Dworetsky, A., Kraus, B. T., Gilmore, A. W., Berg, J. J., Ortega, M., Nguyen, A., Greene, D. J., McDermott, K. B., Nelson, S. M., Lessov-Schlaggar, C. N., Schlaggar, B. L., Dosenbach, N. U. F., & Petersen, S. E. (2019). Trait-like variants in human functional brain networks. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(45). https://doi.org/10.1073/pnas.1902932116

Seitzman, Benjamin A., Gratton, C., Marek, S., Raut, R. V., Dosenbach, N. U. F., Schlaggar, B. L., Petersen, S. E., & Greene, D. J. (2019). A set of functionally-defined brain regions with improved representation of the subcortex and cerebellum. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2019.116290

Thomas Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., Fisch, B., Liu, H., & Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.00338.2011

Van Essen, D. C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T. E. J., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S. W., Della Penna, S., Feinberg, D., Glasser, M. F., Harel, N., Heath, A. C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., … Yacoub, E. (2012). The Human Connectome Project: A data acquisition perspective. In *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2012.02.018

Wang, L., Mruczek, R. E. B., Arcaro, M. J., & Kastner, S. (2015). Probabilistic maps of visual topography in human cortex. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhu277

SUPPLEMENTARY MATERIALS

Supplemental Methods

Preprocessing and FC processing of BOLD data

WashU, Dartmouth, MSC datasets

All functional data were preprocessed to remove noise and artifacts, following Miezin et al. (2000). In the WashU, Dartmouth, and MSC datasets, slice timing correction was performed using sinc interpolation to account for temporal misalignment in slice acquisition time. Next, whole-brain intensity values across each BOLD run were normalized to achieve a mode value of 1000. Motion correction was performed within and across BOLD runs via a rigid body transformation. Next, subject data were aligned to a template atlas (Lancaster et al., 1995) conforming to Talairach stereotactic atlas space (Talairach & Tournoux, 1988) using an affine transformation. The atlas transformation was performed in the WashU, Dartmouth, and HCP datasets by registering the BOLD volumes directly to a T1 structural image. In the MSC, functional data were registered first to the T2 image and then to the T1. All transformations were simultaneously applied along with resampling to 3 mm isotropic voxels.

Following Power et al. (2014), additional denoising was applied to the resting-state data for FC analysis. Temporal masks for each subject's timeseries were created in the WashU and MSC datasets by censoring all frames with a framewise displacement (FD; (Power et al., 2012)) greater than 0.2 mm, and in the Dartmouth dataset by censoring frames with FD greater than 0.25 mm. Across all datasets, segments with fewer than 5 contiguous frames were censored. FreeSurfer 5.0 segmentation using each subject's T1 image generated a white matter and a cerebrospinal fluid nuisance mask per individual. After BOLD data was demeaned and detrended, Regression of nuisance signals was implemented, regressing out global signal, cerebrospinal fluid, and white matter, as well as the six rigid-body motion regressors and their expansion terms (Friston et al., 1996). Data from high-motion frames were interpolated over via a spectra-matching interpolation technique. Data were then bandpass temporally filtered between 0.009 Hz to 0.08 Hz was. Finally, the data were spatially smoothed at FWHM (6 mm).

HCP dataset

Preprocessing and FC processing of HCP subjects were carried out similarly to the other datasets with several differences. First, slice-timing correction was not performed, following the recommendations of the minimal preprocessing pipeline guidelines (Glasser et al., 2013). Second, prior to censoring high-motion frames, motion parameters were low-pass filtered at 0.1 Hz to mitigate effects of respiratory artifacts on motion estimates attributable largely to the multi-band, fast-TR data acquisition (Fair et al., 2020; Siegel et al., 2017). Following this, a filtered FD threshold of 0.1 mm was applied to censor frames. Data were originally processed in MNI atlas space with 2 mm isotropic voxels, and were transformed into Talairach space with 3 mm isotropic voxels in a single step prior to spatial smoothing as described above.

Supplemental Tables

Supplemental Table 1

Dataset	Scanner	TR (s)	TE (ms)	Flip angle	Num. slices	Voxel size (mm)
WashU-120 and WashU-24	Siemens Trio 3T	2.5	27	90º	32-36	4 x 4 x 4
Dartmouth	Philips Achieva 3T	2.5	35	90º	36	3 x 3 x 3.5
MSC	Siemens Trio 3T	2.2	27	90º	36	4 x 4 x 4
HCP	Custom HCP Skyra 3T	0.72; MB = 8	33	52º	72	2 x 2 x 2

Supp. Table 1: Acquisition parameters for functional MRI runs.

Supplemental Figures

Supplemental Figure 1

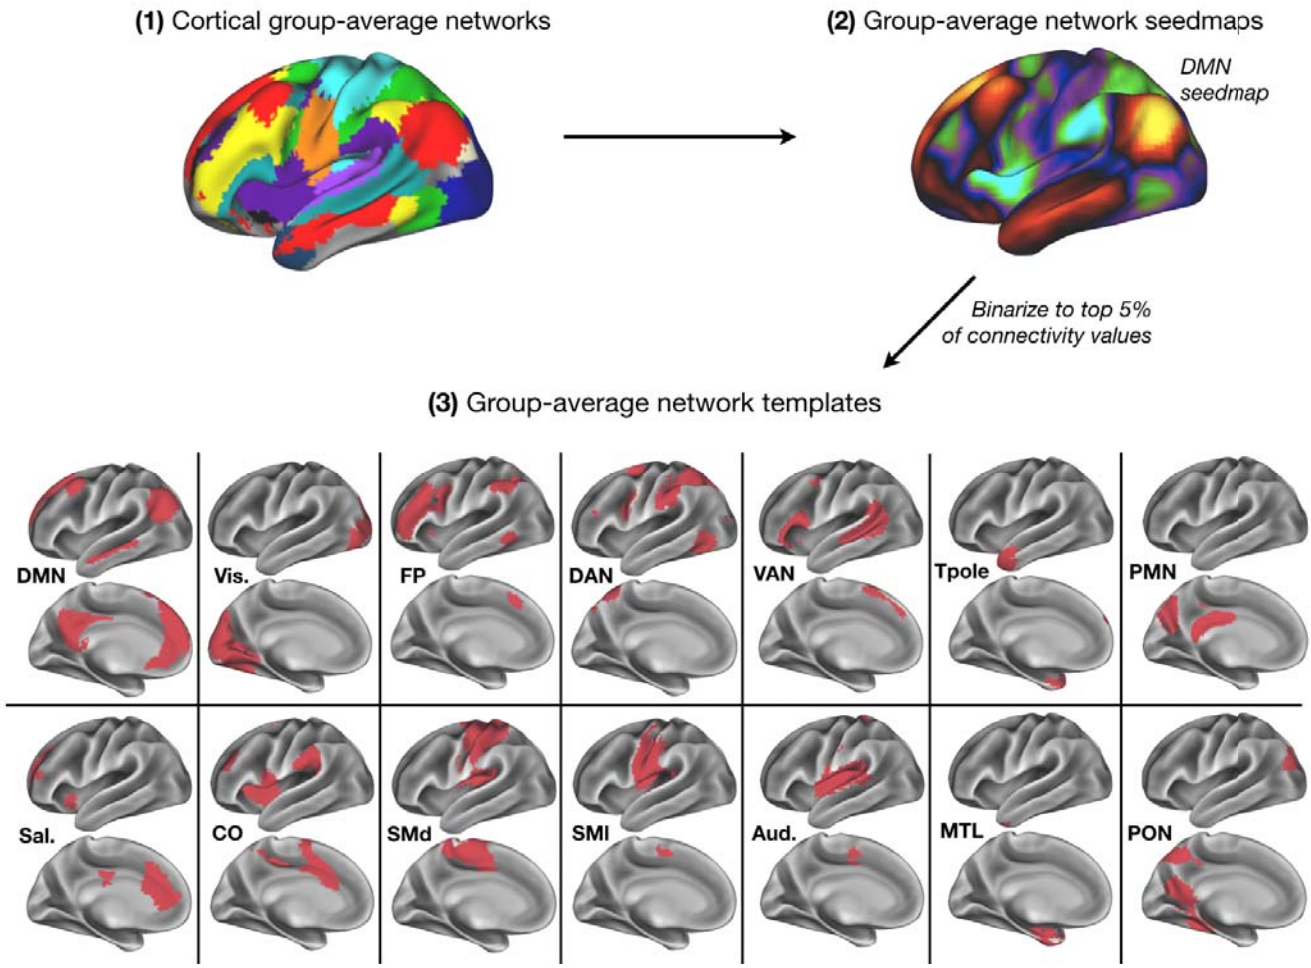

Supp. Fig. 1: Creation of network templates and binarized versions of network templates for each of the 14 networks analyzed in this study.

Supplemental Figure 2

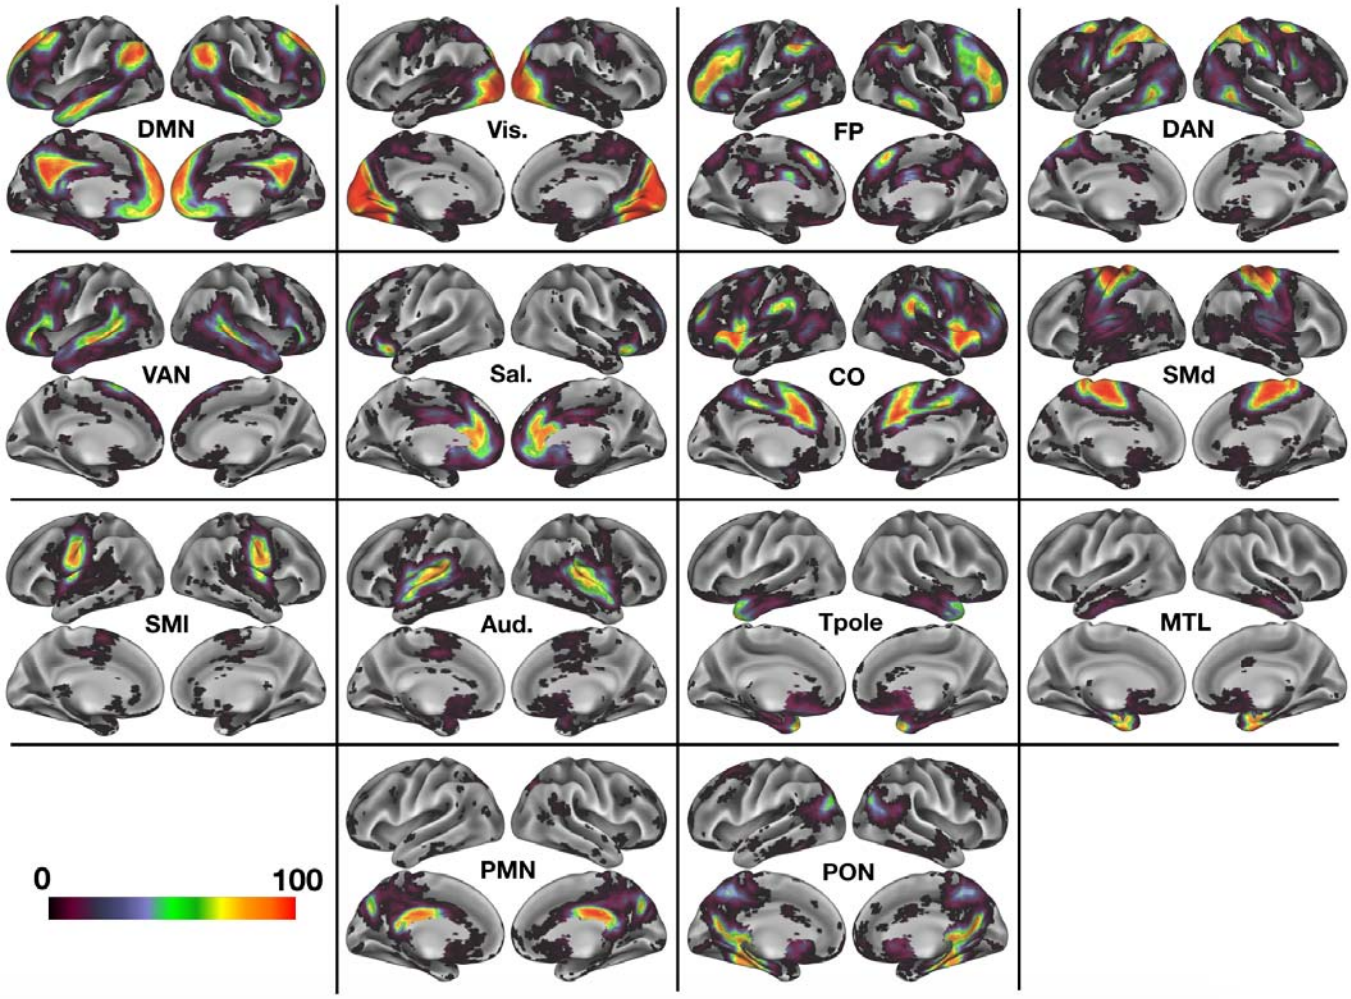

Supp. Fig. 2: Probabilistic overlap maps for all 14 networks.

References

Fair, D. A., Miranda-Dominguez, O., Snyder, A. Z., Perrone, A., Earl, E. A., Van, A. N., Koller, J. M., Feczko, E., Tisdall, M. D., van der Kouwe, A., Klein, R. L., Mirro, A. E., Hampton, J. M., Adeyemo, B., Laumann, T. O., Gratton, C., Greene, D. J., Schlaggar, B. L., Hagler, D. J., … Dosenbach, N. U. F. (2020). Correction of respiratory artifacts in MRI head motion estimates. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2019.116400

Friston, K. J., Williams, S., Howard, R., Frackowiak, R. S. J., & Turner, R. (1996). Movement-related effects in fMRI time-series. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.1910350312

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Van Essen, D. C., & Jenkinson, M. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2013.04.127

Lancaster, J. L., Glass, T. G., Lankipalli, B. R., Downs, H., Mayberg, H., & Fox, P. T. (1995). A modality-independent approach to spatial normalization of tomographic images of the human brain. *Human Brain Mapping*. https://doi.org/10.1002/hbm.460030305

Miezin, F. M., Maccotta, L., Ollinger, J. M., Petersen, S. E., & Buckner, R. L. (2000). Characterizing the hemodynamic response: Effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *NeuroImage*. https://doi.org/10.1006/nimg.2000.0568

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2011.10.018

Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artifact in resting state fMRI. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2013.08.048

Siegel, J. S., Mitra, A., Laumann, T. O., Seitzman, B. A., Raichle, M., Corbetta, M., & Snyder, A. Z. (2017). Data quality influences observed links between functional connectivity and behavior. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhw253

Talairach, J., & Tournoux, P. (1988). Co-Planar Stereotaxis Atlas of the Human Brain: 3-D Proportional System. In *Thieme Medical Publisher*.