Individual heterogeneity in the functional network topography of medial prefrontal cortex

Claudio Toro-Serey[1], Sean M. Tobyne[2], & Joseph T. McGuire[1]

1. Department of Psychological and Brain Sciences, Boston University, Boston, USA

2. Graduate Program for Neuroscience, Boston University, Boston, USA

Corresponding author: Claudio Toro-Serey (ctoro@bu.edu)

Abstract

Human medial prefrontal cortex (mPFC) and posterior cingulate cortex (PCC) are implicated in multiple cognitive functions, including subjective valuation processes and processes linked to the default network (DN). Our ability to interpret these seemingly co-localized effects is constrained by a limited understanding of the individual-level heterogeneity in mPFC/PCC functional organization. Here we used cortical surface-based meta-analysis to identify a parcel in human PCC that was preferentially implicated in DN effects relative to valuation. We then used resting-state fMRI data and a data-driven network analysis algorithm, spectral partitioning, to partition mPFC and PCC into "DN" and "non-DN" subdivisions in individual participants (n = 100 from the Human Connectome Project). The spectral partitioning algorithm efficiently identified individual-level cortical subdivisions that were reliable across test/retest sessions and varied markedly across individuals, especially in mPFC. Our results point toward a new generation of strategies for assessing whether distinct cognitive functions engage common or distinct mPFC subregions.

Keywords: brain networks; DN; individual differences

Introduction

Human medial prefrontal cortex (mPFC) and posterior cingulate cortex (PCC) are jointly associated with a large and diverse set of cognitive functions (Hiser & Koenigs, 2018; Kragel et al., 2018). One such function is subjective valuation during economic decision making. Both mPFC and PCC (together with ventral striatum) consistently show greater BOLD activity in response to more highly valued choice prospects and outcomes, relative to prospects and outcomes that are less highly valued (Bartra, McGuire, & Kable, 2013; Clithero & Rangel, 2014; Hiser & Koenigs, 2018; Kable & Glimcher, 2007; Levy, Lazzaro, Rutledge, & Glimcher, 2011). Both mPFC and PCC are also among the regions implicated in a set of functions associated with the default network (DN). DN effects include a decrease in BOLD activity during cognitively demanding tasks compared with less-demanding task conditions or periods of rest (Laird et al., 2009; McKiernan, Kaufman, Kucera-Thompson, & Binder, 2003), as well as a distinctive pattern of inter-region correlations in resting-state fMRI data (Fox et al., 2005; Greicius, Krasnow, Reiss, & Menon, 2003; Yeo et al., 2011).

Coordinate-based meta-analyses of the fMRI literature show that valuation and DN effects occur in overlapping and largely indistinguishable regions of ventral mPFC (Acikalin, Gorgolewski, & Poldrack, 2017). The two sets of effects diverge in other brain areas; a region of the striatum is preferentially associated with valuation, whereas bilateral temporoparietal junction is linked with the DN. Within PCC and adjacent precuneus, a more posterior subregion appears DN-specific, whereas a more anterior subregion is implicated in both effects. One possible interpretation of these findings is that mPFC and anterior PCC operate as common nodes within partially overlapping distributed brain systems that subserve valuation and DN effects. This colocalization has prompted theories that valuation and DN-related processes share common elements at the psychological level. For example, the two sets of processes might share elements of self-referential cognition, episodic memory, mental simulation, or monitoring and regulation of internal states (Acikalin et al., 2017; Clithero & Rangel, 2014; Northoff & Hayes, 2011; Reddan, Wager, & Schiller, 2018).

However, strong conclusions about functional colocalization require consideration of individual-level heterogeneity in topographic patterns of brain activity. A recognized limitation of group averaging and meta-analysis is that the functional topography of individual brains can be misaligned and blurred (Fedorenko, Duncan, & Kanwisher, 2012; Guntupalli, Feilong, & Haxby, 2018; Michalka, Kong, Rosen, Shinn-Cunningham, & Somers, 2015; Tobyne et al., 2018; Wang et al., 2015; Woo et al., 2014), exaggerating the apparent overlap across domains. This concern is especially pronounced in mPFC, which is subject to considerable idiosyncratic cortical folding (Lopez-Persem, Verhagen, Amiez, Petrides, & Sallet, 2019; Mackey & Petrides, 2014; Zilles, Palomero-Gallagher, & Amunts, 2013) and inter-subject functional variability (Mueller et al., 2013). An

alternative approach is to focus on analyses at the individual-participant level. Individual-level analyses of fMRI data have identified idiosyncratic, reliable, and valid patterns of functional organization that would be blurred in aggregative estimates (Gordon et al., 2017; Gratton et al., 2018; Laumann et al., 2015; Tobyne et al., 2018), and subject-specific network arrangements have been found to predict behavioral characteristics (Kong et al., 2018). Recent work has uncovered fine-grained subdivisions within the DN using both data-driven clustering and individually customized seed-based connectivity analysis (Braga & Buckner, 2017; Braga, Van Dijk, Polimeni, Eldaief, & Buckner, 2019). It is therefore possible that the previously reported overlap of DN and valuation effects can be attributed to low effective spatial resolution, and that the organization of mPFC and PCC would be better understood at the individual level. An important first step, and the goal of the present paper, is to quantify the degree of topographic heterogeneity of DN effects within mPFC and PCC in a large sample of individuals.

A useful way to characterize individual-specific brain organization is to examine patterns of resting-state functional connectivity. Connectome-based analyses of resting-state functional connectivity have been fruitful in identifying individualized functional subregions that correspond well to task-induced activity patterns (Gordon et al., 2017; Laumann et al., 2015; Smith et al., 2009; Tobyne et al., 2018). A functional connectome can be represented in the form of a network, and graph theoretic methods can be applied to analyze the network's structure (Bassett & Sporns, 2017; Rubinov & Sporns, 2010). In the context of network analysis, community detection algorithms subdivide brain networks into sets of nodes that share more connections with each other than with the rest of the network (Fortunato & Hric, 2016; Garcia, Ashourvan, Muldoon, Vettel, & Bassett, 2018). Here we use the technique of spectral partitioning (SP), an efficient community detection algorithm that deterministically subdivides a network into two communities (Belkin & Niyogi, 2003; Chung, 1997; Fiedler, 1975). SP has previously been used to characterize the posterior-anterior functional gradient of the insula using resting-state fMRI data (Tian & Zalesky, 2018), and was shown to robustly and reliably separate both simulated and actual primate ECoG networks (Toker & Sommer, 2019). We use SP here to identify subsets of nodes within mPFC and PCC that share spontaneously covarying temporal activation patterns.

In this study, we aimed to subdivide mPFC and PCC into individual-specific DN and non-DN communities, and to quantify the degree of topographic heterogeneity in the resulting community structure over time and across individuals. We did this by capitalizing on the respective strengths of meta-analysis and subject-specific analyses of brain networks. First, we established a search space by selecting parcels from an established brain atlas (Glasser et al., 2016) that corresponded to previously defined DN and limbic networks on the medial cortical wall (Yeo et al., 2011). A cortical surface-based meta-analysis of the DN and valuation

literatures identified a parcel in PCC that was DN-specific at the aggregate level. We then derived a functional connectivity network of all the surface vertices within the search space for each of 100 individual resting-state fMRI data sets from the Human Connectome Project (HCP; Van Essen et al., 2012), and used the SP algorithm to subdivide each individual's network into DN and non-DN communities (labeled according to which community included the meta-analytically identified DN-specific parcel in PCC). Focusing on individual vertices in the search space rather than the parcels (as is typical in brain network analyses) allowed us to finely delineate the topographic extent of each community. The resulting communities varied topographically across individuals, while also appearing to follow common organizational principles. Sliding window and test-retest analyses showed that these partitionings were highly similar across scanning days within (but not between) individuals, and that individual-level idiosyncrasy was greater in mPFC. Partitionings obtained from the SP algorithm had higher test-retest reliability than did analogous results from seed-based functional connectivity. Lastly, we describe how the structure of the resulting automatically defined DN and non-DN communities both aligns with and differs from recently proposed scheme for identifying subdivisions within the DN (Braga & Buckner, 2017; Braga et al., 2019). Our work highlights the usefulness of estimating brain effects at the individual level in mPFC and PCC, and provides a new framework and tool set for future investigations of overlap across cognitive domains.

Methods

All code used in this study is openly available at https://github.com/ctoroserey/mPFC_partitioning

Search space

For all analyses, we defined our search space based on the 17-network parcellation proposed by Yeo et al. (2011). First, we selected vertices on the medial cortical surface that were contained by the DN and limbic networks in HCP's 32,000 vertex surface space (fs_LR_32k). Next, we overlaid those networks on a parcellated atlas of the human cortical surface (360 regions; Glasser et al., 2016), and retained a set of parcels that covered approximately the same brain regions. This resulted in a search space that consisted of 40 parcels across hemispheres (Supplementary Table 1). The search space in each hemisphere was naturally divided into two spatially non-contiguous clusters in PCC and mPFC, facilitating the examination of each region separately.

Meta-analysis

We used a novel approach to cortical surface parcel-based meta-analysis to assess whether individual parcels within the search space were preferentially associated with valuation or DN effects. For subjective valuation, we gathered peak activation coordinates from 200 studies that were associated with positive effects in contrasts of higher-value minus lower-value outcomes or prospects (Bartra et al., 2013). For DN, we acquired coordinates from 80 studies that were related to task-deactivation (Laird et al., 2009). The coordinates represent areas that exceeded the statistical significance threshold in each original study. For each study, we created an indicator map in standard volumetric space (MNI152, 1 mm resolution) which contained values of 1 in a 10 mm radius sphere around each reported activation peak, and values of 0 elsewhere (Wager, Lindquist, Nichols, Kober, & Van Snellenberg, 2009). The indicator map for each study was then projected to a standard cortical mesh (fsaverage, 160,000 vertices, projfrac-max from 0 to 1 by 0.25, registered using mni152.register.dat) using FreeSurfer's mri_vol2surf (Dale, Fischl, & Sereno, 1999; Fischl, Sereno, & Dale, 1999) (http://surfer.nmr.mgh.harvard.edu/). We then resampled the Glasser et al. (2016) parcellation to fsaverage, and tallied how many studies had positive indicator values intersecting with each cortical parcel (the details of the resampling procedure are described in https://wiki.humanconnectome.org/display/PublicData/HCP+Users+FAQ#HCPUsersFAQ-9.HowdoImapdatabetweenFreeSurferandHCP, and were implemented using a custom script available at https://github.com/stobyne/Spherical-Surface-Swapper). Two studies from the subjective valuation corpus were removed because they did not contain activation peaks that overlapped with cortex, leaving a final number of 198 studies.

To test for parcels that were significantly more strongly associated with one domain than the other, we performed per-parcel chi-squared tests comparing the proportion of studies with activation in that parcel between the two domains. We permuted the study domain labels (DN or valuation) 5000 times while preserving the total number of studies in each domain, and on each iteration stored the maximum resulting chi-squared statistic across all parcels. This gave us a null distribution of 5000 maximum chi-squared values. The 95th percentile of this distribution served as an FWE-corrected significance threshold to evaluate unpermuted chi-squared values.

Resting-state fMRI Data

Our fMRI analyses used resting-state fMRI data from the Human Connectome Project (Van Essen et al., 2012) Q6 release (N = 100, randomly sampled from the total pool of 469 available subjects). Each subject's data was acquired over two days at Washington University in St. Louis on a Siemens CONNECTOM Skyra

MRI scanner (Siemens, Erlangen, Germany). Four resting state runs (repetition time = 0.720 s, echo time = 33.1 ms, flip angle = 52°, multiband factor = 8, 72 slices, 2 mm isotropic voxels) each comprised 1200 time points (14 min 24 s) for a total of 4800 time points. Two runs were acquired on each day, with the phase encoding direction set to left-right for one run and right-left for the other. Only subjects with both left-right and right-left phase encoding for each day were included (i.e. subjects with four resting-state fMRI sessions). In addition, only datasets with low motion levels (under 1.5 mm) and less than 5% of points over 0.5 mm framewise displacement (FD) were used. See (Van Essen et al., 2012) for more details about the data acquisition protocol.

Data initially underwent the HCP minimal preprocessing pipeline (Glasser et al., 2013), which included gradient nonlinearity correction, motion correction, EPI distortion correction, high-pass filtering (0.0005 Hz threshold), MNI152-based normalization, surface reconstruction, and mapping of functional data to a standardized cortical surface model (details can be found in Glasser et al., 2013). In addition, data underwent temporal denoising based on independent components (FMRIB's ICA-based X-noiseifier, FIX; Griffanti et al., 2014; Salimi-Khorshidi et al., 2014). Data were further preprocessed using an in-house pipeline described previously (Tobyne, Osher, Michalka, & Somers, 2017). Steps (in order) included linear interpolation across high motion timepoints with over 0.5 mm of framewise displacement, band-pass filtering (allowed frequencies ranged from 0.009 and 0.08 Hz), and temporal denoising via mean grayordinate signal regression (Burgess et al., 2016). After filtering and denoising, the interpolated high-motion time points were censored by deletion and each run was temporally de-meaned. The processed time series had a median of 4799 time points (minimum = 4661) across participants. Each subject's brain was comprised of 32k standard grayordinates per hemisphere (combined in a CIFTI file). We retained only the cortical surfaces, which resulted in 59,412 total surface vertices per subject.

Network Definition

All network analyses were performed using the igraph package (v. 1.1.2; https://igraph.org/r/; Csardi & Nepusz, 2006) in R (v. 3.4.1; https://www.r-project.org/; R Core Computing Team, 2017). To establish each subject's network, we selected all the vertices contained within the mPFC/PCC search space (n = 4,801 per subject; mPFC = 2854, PCC = 1947) and computed the Pearson correlation of the time series for every pair of vertices. This produced a weighted network for each subject, in which the nodes were surface vertices and the edges were the correlations among them. Next, we applied Fisher's r to z transformation, and performed a two-sided significance test to identify significant connections. The resulting p-values were corrected for multiple comparisons (FDR < 0.05), and edges with non-significant correlations were set to 0 (all other edges

retained their respective z-values). The proportion of remaining edges (i.e. network density) was high for all individuals (mean = 0.78, SD = 0.024), and the retained edges mostly consisted of positive correlations (mean proportion positive = 0.69, sd = 0.04). Next, we took the exponential of these values so that all weights were positive while maintaining the ordinal ranks of the original correlation distribution within the set of retained edges. Non-retained edges were reset to 0 to exclude them from the network analyses. We generated and analyzed network weight matrices at four levels: (1) for each subject's full concatenated dataset (up to 4800 TRs); (2) on each step of a sliding window analysis (see Partition Evaluation for more details); (3) for the concatenated time series for the two runs on each day (up to 2400 TRs); and (4) for each run separately (up to 1200 TRs).

Community Detection

Communities (i.e. clusters) were identified using the SP algorithm (Belkin & Niyogi, 2003; Chung, 1997; Fiedler, 1975; Higham, Kalna, & Kibble, 2007). First, each network was represented as an n x n network weight matrix W as described above (where n equals the number of vertices in the search space, 4,801). The matrix was then transformed into its symmetric normalized Laplacian form

$$L = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$$

Where I is an identity matrix of size n, and D is a diagonal matrix containing the strength of each vertex (i.e. the sum of its edge weights with all other vertices). This resulted in a matrix wherein each entry was the negative normalized value of the connection (from 0 to 1) between any two vertices relative to their combined connectivity strength, and with ones along the diagonal. The transformation ensures that every row sums to zero. We then computed the eigenvalues and eigenvectors of the symmetric normalized Laplacian matrix, and used the eigenvector associated with the second-to-lowest eigenvalue (traditionally called the 'Fiedler vector') to divide the network into two. The Fiedler vector consists of a set of positive and negative values and is binarized by sign to partition the network into two similarly-sized communities (Fiedler, 1975). In this way, SP avoids producing communities that are too small to be physiologically meaningful (for example, small sets of vertices that are spuriously correlated due to measurement noise). Given that this data-driven method does not label the two communities or establish correspondence across participants, we defined each individual's "DN" community as that which contained the DN-specific PCC parcel identified in our meta-analysis. The high density of the graphs ensured that SP did not face the issues associated with its use in sparse networks (Fortunato & Hric, 2016).

In order to evaluate the validity of the resulting partitionings across community-detection methods, we also estimated network communities using the more traditional approach of modularity maximization (Garcia et al., 2018), based on the algorithm from Clauset et al. (2004). The method heuristically iterates through many possible combinations of vertices, and selects the partitioning that maximizes the within-community edge weights, relative to a random network containing the same number of edges and communities. Unlike SP, modularity can fractionate a network into more than two communities. Agreement between the partitions provided by the bounded and unbounded community detection methods would suggest the results are not distorted by the restriction of SP to binary partitionings.

Partition Evaluation

We used the Adjusted Rand index (ARI) to evaluate the stability and topographical heterogeneity of the communities within and across individuals (Hubert & Arabie, 1985), which was calculated using the "mcclust" package in R (Fritsch, 2012). The ARI is a metric that quantifies the similarity between two alternative clusterings of the same data. The base of the ARI is computed by the formula

$$\frac{a+b}{a+b+c+d}$$

where a is the number of pairs of nodes that were grouped together in both partitionings, b is the number that were grouped separately, and c and d denote the number of nodes grouped together (separately) in one partitioning, but separately (together) in the other. Therefore, the ARI estimates the fraction of all possible node pairs that had the same status (connected or not) in both partitionings (with the denominator equal to $n(n-1)/2$). The resulting ratio is adjusted against a baseline given by the expectation assuming independent partitionings to yield an index that ranges from 0 to 1, where 0 denotes the value expected by chance. This means that even though differences are heavily penalized, positive ARI values compare favorably against chance clustering (and the index can take negative values if the ratio given by the formula above falls below the chance level). In short, the ARI quantifies the chance-corrected agreement between any two partitions while being agnostic to the labeling scheme.

We performed a number of comparisons among partitions. First, we computed the degree of agreement between SP and modularity per subject. SP and Modularity have been previously found to show underfitting and overfitting tendencies, respectively, in their community detection performance in a diverse set of network types (Ghasemian, Hosseinmardi, & Clauset, 2019), so alignment between the two algorithms would increase our confidence in the validity of the resulting partitionings. Next, we compared the subject-level SP partitionings

across individuals, and calculated the mean pairwise ARI for the group. We then performed the same evaluation for PCC and mPFC separately, and examined whether there were mean differences in overall agreement within these regions by performing a paired permutation analysis (5000 iterations) on all pairwise ARI comparisons across subjects.

To estimate the degree of variability of partitionings per individual we performed a sliding window analysis (20 min windows, 1 min increments, median number of windows per subject = 37, range = 35 - 37), comparing each window's resulting partitioning against the partitioning derived from the subject's whole data set. A 20-min window has previously been found adequate for identifying stable community features in brain networks (Gordon et al., 2017). We assessed whether the magnitude of the Fiedler vector value for a given vertex (for the full subject-level data set) was associated with the stability of that vertex's sub-network assignment across time windows. To do this, we fit a mixed effects logistic regression model, in which the dependent variable was the proportion of times each vertex participated in the DN community across windows, and the explanatory variables included a random effect of subject and a fixed effect of the Fiedler vector value for that vertex (derived from their full time series). Based on this relationship, we identified a threshold Fiedler vector value for each subject, such that above-threshold vertices were persistently associated with either DN or non-DN more than 99% of the time.

We then estimated the level of agreement between network partitions estimated using data across individual scan days (with 2 days per participant). If the functional organization estimated by SP is indeed individual-specific, we should see higher agreement within individual (test/re-test across days) than across individuals. We tested this idea by computing the ratio of the mean ARI within and between individuals. Ratios close to one would denote similar within-participant and across-participant alignment, whereas ratios considerably higher than one would suggest that partitions were more similar within-participant than across participants. We then extended this idea by computing the agreement across individual runs (4 per subject). Similar to the day-based analysis, we assessed whether run-level data showed higher agreement within-subject than between subjects.

Seed-based Resting-state Functional Connectivity versus Community Detection

We evaluated the performance of the SP algorithm in comparison to a simpler partitioning approach based on seed-based functional connectivity. Independently for each day (2 per individual), we estimated each subject's DN partition in mPFC based on its vertex-wise functional correlations (Pearson) with the spatially averaged activity across all vertices in PCC. We compared these seed-based maps with the unthresholded

Fiedler vectors produced by SP, with the sign of the Fiedler vector oriented so the DN community was marked by positive values in every subject. We calculated three sets of across-day similarity values for each individual: 1) between the two seed-based maps; 2) between the two SP-based maps; and 3) between seed- and SP-based maps. Because the values in the maps were continuous-valued (and not categorical labels, which would be amenable to ARI), we quantified the similarity between maps in terms of the spatial Spearman correlation across vertices. These spatial correlations were meant to determine the test/re-test reliability of each approach, as well as the overall level of agreement between them. For 7 subjects, the communities produced with one of the days' data sets had split coverage of area 7m, and so our community labeling scheme for the Fiedler vector produced a sign mismatch across days. ARI is robust to such labeling issues, but the inconsistency produced strong negative correlations of the Fiedler vector across days for these individuals. Visual inspection showed that the community layout was well aligned across days, and so we matched the labeling of their partitionings based on the day that sufficiently covered area 7m.

The two methods were expected to produce somewhat similar results, but the one displaying greater within-subject agreement across days should be preferred (for a discussion on the stability of functional networks see (Gratton et al., 2018; Kong et al., 2018)). We therefore compared the within-subject spatial correlation coefficients produced by each method through a paired permutation analysis. For each of the 100 individuals, we computed the difference in inter-day correlations between methods, randomized the sign of these values 5000 times, and computed the mean of these differences on each iteration. The empirical difference in means was then evaluated against this permuted distribution.

Results

Meta-analysis

We performed a coordinate-based meta-analysis to identify cortical surface parcels within mPFC and PCC that were preferentially associated with the DN or valuation literature. Volumetric coordinates from 80 studies with task deactivation contrasts and 198 studies with valuation contrasts were projected onto a cortical surface, and mapped to discrete parcels from a multimodal cortical parcellation (Glasser et al., 2016) to produce a list of brain areas reported per study. The 40 parcels considered were limited to the medial portion of the default and limbic networks defined by the Yeo et al. (2011) 17-network parcellation. Domain-specificity was tested by first permuting the domain labels across studies (DN or valuation) to create a null distribution for the maximum chi-squared statistic in the search space (see Methods for details). The

null distribution was used to identify regions that were reported significantly more often in one literature or the other.

Figure 1 shows the proportion of times each parcel was reported for each domain, as well as the significant differences between domains. The 95th percentile of the permuted chi-squared distribution was 8.87. Based on this threshold, area 7m in PCC/precuneus was the only parcel to show a preferential association with the DN bilaterally (Left: observed $\chi^2 = 10.07$, $p = 0.029$; Right: observed $\chi^2 = 18.89$, $p < 0.001$). The adjacent area v23 exhibited a similar effect, albeit only unilaterally (Right: observed $\chi^2 = 11.51$, $p = 0.011$; Left: observed $\chi^2 = 8.25$, $p = 0.067$). There appeared to be a bilateral preference toward valuation effects in mPFC area 25 (Left: observed $\chi^2 = 12.91$, $p = 0.005$; Right: observed $\chi^2 = 12.83$, $p = 0.005$); however, closer inspection suggested this effect was driven by subcortical foci centered in adjacent ventral striatum. No other parcels were preferentially implicated in valuation relative to DN. We therefore selected area 7m as an interpretable, bilateral reference point for labeling DN and non-DN communities in the analyses that follow. We note that the area labeled 7m in the parcellation used here (Glasser et al., 2016) is different from (and located inferiorly on the medial surface to) the non-DN area 7m discussed in previous work (Andrews-Hanna, Reidler, Huang, & Buckner, 2010).

Individual-level DN and non-DN communities

Within the mPFC/PCC search space, we estimated the topography of the DN for each individual. Using each individual's full time series (approximately 4800 total TRs from four 14-min scanning runs acquired over two days), we calculated the full vertex-to-vertex correlation matrix for the 4801 surface vertices in the search space. We represented each individual's correlation matrix in the form of a network, with cortical surface vertices as nodes and thresholded/transformed correlation values as edge weights. We then applied the SP community detection algorithm to partition the network into two cohesive functional communities.

Figure 2 shows a representative partitioning of the search space for a single participant (100307; additional examples are presented in Supplemental Figure 1). The SP algorithm subdivides a network according to the positive versus negative values in the Fiedler vector (the eigenvector related to the second-to-lowest eigenvalue of the network's Laplacian matrix, see Methods). Since this is a data-driven approach, there is no a priori labeling for the two communities. We assigned the DN label to the community that contained the DN-specific PCC parcel from the meta-analysis (7m). We oriented each individual's Fiedler vector so positive values corresponded to the DN community, and were assigned a value of 1 in the binarized partitionings (with 0 denoting non-DN). In qualitative terms, the resulting patterns contained substantial DN coverage in posterior

PCC (as dictated by our labeling strategy), with non-DN vertices in anterior PCC. The mPFC region tended to include DN vertices in its ventral-anterior and dorsal-anterior areas, with a persistent non-DN pocket between them. This non-DN section extended posteriorly into pregenual cinglate cortex (area a24).

Before evaluating the degree of generazibility of this topographic pattern across individuals, we examined the validity of the partitionings by comparing them to results from an alternative community detection algorithm, modularity maximization (Clauset et al., 2004). Modularity seeks to find the set of communities that maximizes within-community connection weights relative to a null model. Since modularity is not constrained to a predetermined number of communities, it was capable of finding more than two in our data set. We quantified the cross-method agreement in terms of the Adjusted Rand Index (ARI; see Methods), which measures the proportion of node pairs in a network that were either clustered together or separately in both partitionings, while being agnostic to labeling schemes and controlling for chance clustering. The ARI normally takes values ranging from 0 to 1, with 0 indicating chance agreement (but can take negative values if the similarity falls below chance).

The two clustering methods had high agreement (mean ARI = 0.87, SD = 0.11). Modularity showed a tendency to produce additional communities (median = 3, range = 2, 4). However, the additional communities encompassed a small number of vertices (median = 14.5, IQR = 4 - 31.75) compared to the principal two (median = 4786.5, IQR = 4769.25 - 4797), suggesting that a binary partitioning provided a reasonable approximation of the network's true community structure.

Next, we examined the similarity of SP-based partitionings across individuals by computing the ARI between every pair of subjects, and found modestly above-chance agreement overall (mean = 0.12, SD = 0.05). Qualitative inspection of the community organization showed good alignment for PCC, whereas the pattern in mPFC was consistent but shifted topographically across subjects. To quantify this heterogeneity in mPFC, we calculated the between-subject ARI for each region separately (Figure 3). The functional topography of PCC was better aligned across individuals (mean = 0.18, SD = 0.09) than mPFC (mean = 0.09, SD = 0.04; paired permutation, $p < 0.001$; Cohen's D = 1.34), although the mean ARI in mPFC still exceeded the chance value of zero (Wilcoxon signed rank test, $p < 0.001$; Cohen's D = NA).

Pattern variability over time

We next sought to identify a set of temporally stable core nodes in each community. In order to estimate the stability of partitions over time, we performed a sliding window analysis on each subject's full time series (20 min windows shifting by 1 min). We compared the partitioning derived from each window with the

partitioning computed using the entire time series (Figure 4).

The mean ARI along each subject's time series was high for both PCC (mean = 0.59; SD = 0.14) and mPFC (mean = 0.49; SD = 0.12), with significantly higher ARI for PCC (paired permutation, $p < 0.001$; Cohen's D = 0.71). A subset of nodes showed exceptionally high stability, in that they were assigned to the same community in every time window. The percentage of stable nodes ranged from 0 to 73% across individuals (median = 50%, IQR = 32% - 60%).

We next tested whether the continuous-valued Fiedler vector (before binarization into discrete communities) carried information about the stability of individual nodes. There is precedent in the literature for the idea that the magnitude (and not just the sign) of the Fiedler vector values conveys important information about the role of each node in the network (Gkantsidis, Mihail, & Zegura, 2003; Tian & Zalesky, 2018). Therefore, we tested whether the magnitude of the eigenvector values was associated with the stability of nodes over time. Specifically, we estimated the proportion of DN affiliations per node as a function of Fiedler vector values, using a logistic mixed effects model (Figure 4). The model identified a positive significant relationship between these features ($\beta = 211.54$, SE = 0.51, $p < 0.001$), signifying that vertices with higher absolute Fiedler vector values were more persistent in their relationship with their corresponding community over time.

These analyses suggest that there is potential value in thresholding the Fiedler vector as a means to identify reliable DN and non-DN vertices on an individual subject basis. We therefore thresholded each subject's Fiedler vector to produce these refined maps. For each individual, we estimated the threshold by selecting the smallest absolute Fiedler vector value that yielded an average stability across suprathreshold nodes of 99%, for positive (mean = 0.0132, SD = 0.0051) and negative (mean = -0.0144, SD = 0.0056) values separately. Individuals without such stable nodes (n = 17) were not thresholded, and were included in the subsequent analyses in unthresholded form. The median proportion of retained vertices per individual was 0.48 (IQR = 0.29 - 0.63). Sub-threshold vertices were set to zero in Fiedler vector maps and 0.5 in the binarized maps (so that they would not bias the calculation of averages). Figure 5A shows the thresholded partitioning for the same individual shown in Figure 2. The maps used in all subsequent analyses were thresholded by this individualized criterion.

With these thresholded partitions, we recomputed the overall similarity across participants. Compared to before, there was lower topographic agreement across individuals (mean ARI = 0.07, SD = 0.04). The same was true for both PCC (mean = 0.11, SD = 0.07) and mPFC (mean = 0.04, SD = 0.03) separately, although the significance of the differences between areas was preserved (paired permutation, $p < 0.001$; Cohen's

D = 1.2). Figure 5B shows the average of the thresholded partitions across all participants, denoting the proportion of times a vertex was affiliated with the DN community. This summary illustrates the common organizational layout of both communities, but also highlights the considerable variability across individuals.

To test the possibility that the higher inter-subject variability in mPFC was driven merely by lower signal quality in the retained vertices, we quantified the temporal signal to noise ratio (tSNR) for each region, both before and after thresholding. We calculated tSNR using time series that were not demeaned, but were otherwise equivalent to the data originally used. In terms of tSNR variability across vertices within each region, mPFC had overall greater spatial standard deviation both before and after thresholding (mPFC: pre-threshold mean spatial SD = 33.96, post-threshold mean spatial SD = 30.15; PCC: pre-threshold mean spatial SD = 15.28, post-threshold mean spatial SD = 14.59). However, tSNR after thresholding was significantly higher for mPFC than PCC (mPFC: mean = 77.72, SD = 14.54; PCC: mean = 64.78, SD = 10.32; permutation p-value < 0.001, Cohen's D = 1.03). This reflected a significant increase in mean tSNR in mPFC as a result of the thresholding step (pre-threshold mean = 66.5, SD = 7.87; paired permutation p-value < 0.001, Cohen's D = 0.96), whereas the mean signal quality in PCC was unchanged (pre-threshold mean = 64.56, SD = 10.02; paired permutation p-value = 0.2388, Cohen's D = 0.02). In short, mPFC had higher overall tSNR, albeit with greater variability across nodes. Applying the thresholding step allowed us to focus the analysis on vertices with high signal quality.

Test/re-test reliability across days

The relatively high inter-individual variability seen in the aggregate map could reflect at least three factors: (1) measurement noise, (2) dynamic variation in mPFC network organization, and (3) stable patterns of functional organization that differ across individuals. To arbritrate among these possibilities, we examined the test/re-test reliability of mPFC/PCC community structure across separate days of testing. Insofar as the observed variability reflects individual-specific brain organization, across-day ARI values should be consistently higher within-individual than between individuals (an example comparison for two individuals is provided in Supplemental Figure 2). Figure 6 shows pairwise comparisons among ten example subjects for PCC and mPFC separately (left).

Once again we found low alignment across individuals for PCC (mean = 0.08, SD = 0.06) and mPFC (mean = 0.04, SD = 0.03), but both areas showed comparatively high levels of within-individual agreement (PCC: mean = 0.36, SD = 0.14; mPFC: mean = 0.26, SD = 0.1). We calculated an index of relative specificity by computing the ratio of each individual's across-day (within-participant) ARI to the mean of all

between-participant ARI values involving that individual. The index is expected to take on a value near 1 if partitionings are well aligned across individuals and/or are subject to a common level of measurement noise. It is expected to exceed 1 insofar as functional network organization is reliable and individual-specific. This index is intended to factor out the potential contributions of measurement noise or dynamic instability, which would introduce variability both across individuals and across days.

Figure 6 shows ARI ratios for PCC and mPFC. A signed-rank test showed evidence for specificity (i.e. ratios > 1) in both mPFC (median = 7.38, IQR = 5.93 - 8.77, V = 5029, $p < 0.001$) and PCC (median = 4.3, IQR = 3.41 - 5.13, V = 5034, $p < 0.001$). Moreover, the ratios for mPFC were significantly greater than those for PCC when compared with a paired permutation test ($p < 0.001$; Cohen's D = 0.33). These test/retest results suggest that the topographic variability seen in mPFC arises at least in part from stable and subject-specific organizational patterns.

Test/re-test reliability across runs

We extended the analysis of per-day data by examining whether the organization of the DN could be extracted using per-run data only. The duration of each run (approximately 14 min) falls below a previously suggested stability threshold for fMRI-based modularity estimations (Gordon et al., 2017). Nonetheless, high ARI ratios would indicate that the SP algorithm can reliably estimate individual-specific patterns of DN organization from a single run of data.

Run-specific SP results still captured unique organizational patterns to some degree, even though the overall levels of agreement decreased (PCC between subjects: mean = 0.04, SD = 0.05; mPFC between subjects: mean = 0.02, SD = 0.02; PCC within subjects: mean = 0.17, SD = 0.14; mPFC within subjects: mean = 0.1, SD = 0.08). We again computed each subject's ARI ratio in order to quantify the specificity of the partitions, this time using the mean of 6 across-run (within-participant) ARI values in the numerator of the ratio (Figure 6, right).

As before, a signed rank test showed that both regions had ARI ratios significantly greater than 1 (mPFC: median = 5.85, IQR = 4.29 - 7.84, V = 5025, $p < 0.001$; PCC: median = 3.51, IQR = 2.36 - 4.4, V = 4961, $p < 0.001$), and ratios for mPFC were higher than those of PCC (permutation $p = < 0.001$; Cohen's D = 1.1). This result further confirms that the intrinsic functional organization of mPFC is uniquely arranged per individual, and provides evidence that information about such patterns can be extracted from relatively small amounts of data.

Correlation vs community detection in mPFC

We next explored the possible advantage of community detection relative to a more conventional seed-based functional connectivity analysis for estimating the individual-specific functional topography of mPFC. We examined whether maps generated with SP were more similar per participant across days than those computed from seed-based correlations. We generated a seed time-series by averaging all vertices in the PCC region of our search space, and calculated its correlation with the activity of each vertex in mPFC. We compared the map of correlation values in mPFC to the map of unthresholded Fiedler vector values using Spearman correlations across vertices. Pairwise spatial correlations were calculated among maps computed for each day and method from all individuals. Figure 7A shows that these pairwise comparisons resemble those from the across-day comparisons above, and suggests good alignment between methods, but particularly high agreement within subject and method.

Figure 7B shows the test/re-test reliability across days for patterns derived using community detection, seed-based correlation, and across methods (e.g. Day 1 community detection versus Day 2 seed-based correlation). While both approaches were reliable, community detection displayed a significantly higher mean correlation coefficient across days than seed-based correlation (Community: mean = 0.7, SD = 0.26; Seed-based: mean = 0.63, SD = 0.12; paired permutation $p = 0.007$; Cohen's D = 0.32). Agreement across methods was fair (mean = 0.45, SD = 0.24), signifying that the two approaches identified similar topographic features but also had systematic differences. These findings suggest that graph-theoretic community detection algorithms are advantageous for detecting stable functional topologies, in addition to their other advantages of being data-driven, unbiased and observer agnostic.

Alignment of mPFC community structure with a proposed DN sub-network organization

The thresholded partitions we identified had conceptual and topographic similarities to DN sub-networks A and B proposed by Braga and Buckner (2017). We explored the relationship between the two sets of sub-regions by reproducing their seed-based connectivity approach in two of our subjects. In their original work, Braga and Buckner (2017) manually selected individual vertices in areas including the dorsolateral prefrontal cortex (DLPFC), temporo-parietal junction (TPJ), and parahippocampal cortex (PHC) that produced two spatially anticorrelated networks (Braga & Buckner, 2017; Braga et al., 2019). We hypothesized that if the SP communities corresponded to one or both of the previously proposed sub-networks, our partitionings should match networks A and B generated by seed-based functional connectivity in these diagnostic areas. For whole-brain functional connectomes from two individuals (100307 and 101006), we selected seeds for

networks A and B in TPJ, and confirmed their placement based on functional connectivity patterns observed in TPJ, PCC, DLPFC, and PHC (correlation coefficients thesholded at 0.2), as stipulated by Braga et al. (2019). The whole-brain seed-based functional connectivity maps for two individuals are juxtaposed with the corresponding community detection results in Figure 8. Visual inspection of these networks shows high similarity between our DN community and the previously reported sub-network A. However, the non-DN community filled areas not covered by either DN-A or DN-B. This result supports the idea that the two approaches serve complementary purposes. Whereas Braga and colleagues (2017; 2019) identified subdivisions within the DN, the present community detection approach might be better understood as partitioning DN from non-DN cortex.

Discussion

A considerable amount of meta-analytic work has been dedicated to characterizing the brain activity patterns associated with psychological processes in medial prefrontal cortex (mPFC), revealing both dissociable and overlapping activation across domains (De La Vega, Chang, Banich, Wager, & Yarkoni, 2016; Hiser & Koenigs, 2018; Kragel et al., 2018). In particular, activation patterns associated with subjective valuation and with the default network (DN) have been suggested to be inseparable in this area, with overlap also partially extending to posterior cingulate cortex (PCC) (Acikalin et al., 2017; Bartra et al., 2013; Clithero & Rangel, 2014; Laird et al., 2009). This apparent neural overlap has important implications, as it has motivated theoretical proposals about ways in which these superficially dissimilar domains might involve a shared set of core cognitive processes (Acikalin et al., 2017; Clithero & Rangel, 2014; Northoff & Hayes, 2011).

However, the interpretation of overlap in group-level data depends on the degree to which functional organization is heterogeneous across individuals. Recent studies have shown that heteromodal brain regions have considerable variability in functional connectivity across individuals (Mueller et al., 2013), individual-specific functional topography can be occluded in aggregative estimations (Braga & Buckner, 2017; Gordon et al., 2017; Michalka et al., 2015; Tobyne et al., 2018), and overlap in functional activation can vanish with increases in spatial precision (Woo et al., 2014). These findings suggest that group-level and meta-analysis-level overlap does not necessarily imply overlap in individual brains. However, our understanding of the individual-level heterogeneity in the functional topography of mPFC has been mostly descriptive so far (Braga & Buckner, 2017; Braga et al., 2019; Gordon et al., 2017). To affirm that DN and valuation share neural substrates in this area requires a method to reliably and precisely capture the functional topography of mPFC in isolated individuals, as well as a quantitative estimate of the degree of topographic heterogeneity

across a large group of individuals.

Here we address these challenges by using spectral partitioning (SP), a graph-theoretic community detection algorithm that efficiently separates a network into two (Fiedler, 1975; Higham et al., 2007; Toker & Sommer, 2019). For 100 individuals, we subdivided brain regions that typically show overlap between DN and valuation effects into DN and non-DN communities. Restricting our analyses to a general mPFC/PCC search space made it appropriate to use a technique that identified a vertex-wise, binary partitioning that was sensitive to the complex topography of the brain. This contrasts with whole-brain network analyses, which need to allow for multiple sub-networks and which often use parcels that are several orders of magnitude larger than vertices as the units of analysis. Partitioning an individual's brain network through SP has a number of advantages, including identifying communities deterministically, constraining communities to contain a similar number of vertices (i.e. preventing the allocation of most vertices to a single community), providing continuous values that relate to the strength of a node's community affiliation, and the ability to diagnose the connectedness of a network through examination of its resulting eigenvalues (Chung, 1997; Higham et al., 2007). Comparisons with partitionings formed by modularity maximization, which heuristically determines the ideal number of communities (Garcia et al., 2018), suggested the binary partitioning was appropriate.

We found a generalizable pattern across individual partitionings, in which the DN community covered ventral/dorsal mPFC and posterior PCC, with the non-DN community concentrated in pregenual ACC and anterior PCC. The precise spatial positioning of this general community structure was highly heterogenous across individuals, yet stable across test/re-test evaluations within-individual. The idiosyncrasy in functional topography was particularly pronounced in mPFC, and was identified in both run-based and day-based analyses. Individual-specificity could theoretically arise from a variety of sources. For example, the functional topography of mPFC could be governed by its underlying sulcal and gyral organization, which has been shown to vary systematically across individuals (Lopez-Persem et al., 2019; Mackey & Petrides, 2014). Individual variability could also be due to shifts in functional organization that are independent of structural features (Conroy, Singer, Guntupalli, Ramadge, & Haxby, 2013), or be characterized by the pattern of functional connections with the rest of the brain (Mars, Passingham, & Jbabdi, 2018; Passingham, Stephan, & Kötter, 2002; Tobyne et al., 2018). An important goal for future work will be to assess whether the network layout in this region can be predicted on the basis of aspects of brain structure, such as sulcal morphology, myeloarchitecture (Glasser et al., 2016), or structural connectivity (Osher et al., 2016; Saygin et al., 2011, 2016).

Network-partitioning methods such as SP are data-driven, and therefore provide no labeling information about the resulting communities. We circumvented this issue by independently identifying the DN community

based on its coverage of area 7m, a region in PCC that was preferentially associated with the DN relative to subjective valuation in our meta-analysis. We were able to apply labels derived from this group-level approach on the basis of the topography in PCC, where functional organization was more consistent across individuals. Because each community spanned both mPFC and PCC, the labels extended to mPFC where topography was more heterogeneous.

Our results extend previous work that described individual-specific brain organization. Several recent investigations have identified topographic heterogeneity using a different data aspect ratio than we used here (a small number of individuals and a large number of scanning sessions per individual; Braga & Buckner, 2017; Braga et al., 2019; Gordon et al., 2017). Previous work has also shown that functional correlations among pre-defined cortical parcels are highly stable within an individual (Gratton et al., 2018; Kong et al., 2018). Here we were able to quantify the variability and stability of functional topography in a large sample at a fine, vertex-level spatial granularity, using moderately low amounts of data (down to a single 14 minute scan). The motivation to sub-divide DN also stems from recent work by Kernbach et al. (2018), which identified specialized communication of parcels within DN with the rest of the brain in a large pool of individuals.

In addition to the technical advantages noted above, the SP algorithm offers analytical advantages specific to neuroscience. We found that SP outperformed a traditional seed-based correlation approach in capturing idiosyncratic functional topography. Community detection methods such as SP are stabilized by relying on all pairwise correlations among cortical vertices (rather than correlations with an individual seed). In addition, we found we could increase the temporal stability of SP results by thresholding the underlying Fiedler vector. The magnitude of Fiedler vector values has been recently used to characterize the continuous connectivity profile of the insula with the rest of the brain, challenging the notion of discrete parcellations in that region (Tian & Zalesky, 2018). The combination of discrete classification and graded information yielded by SP provides additional flexibility and richness relative to some other clustering algorithms.

The community organization of PCC and mPFC was congruent with DN sub-networks A and B proposed by Braga & Buckner (2017; Braga et al., 2019). The topography of our thresholded DN community closely matched network A, whereas our non-DN community included cortical territory that was not part of either DN network. Subthreshold vertices from the DN community overlaped with DN-B vertices, suggesting that this sub-network could act as a dynamic intermediary between DN-A and other networks. Our findings therefore complement the initial identification of DN sub-networks by quantifying the systematic variability of their underlying topography in a larger group of people. Understanding the interaction of networks A (DN), B, and non-DN is an important goal for future research. These two sets of results collectively suggest that canonical DN regions can be topographically partitioned into DN and non-DN communities, and that

the DN community can in turn be further divided into sub-networks A and B.

Our findings show that the functional topography of mPFC is variable across a large pool of individuals, and that the SP algorithm is a useful tool for identifying individualized topography in a data-driven way. The ability to capture an individual's functional topography without the need of group priors is clinically relevant, as it could help characterize heterogenous changes in mPFC activity in disorders such as depression and schizophrenia (Hiser & Koenigs, 2018). It will be beneficial for future task-based fMRI experiments to be able to characterize where task-evoked activity is situated relative to an individual's overall mPFC organization. Our work was originally motivated by the apparent spatial overlap between subjective valuation and the default network (Acikalin et al., 2017), but additional distinctions of interest include valuation versus episodic memory functions (Euston, Gruber, & McNaughton, 2012), self-referential thought and theory of mind (Mitchell, Banaji, & Macrae, 2005), and potential topographic distinctions among multiple forms of valuation (Clithero & Rangel, 2014; Shenhav & Karmarkar, 2019; Smith et al., 2010). An individualized frame of reference will enhance the ability of future studies to gauge similarities and differences among brain activity patterns associated with diverse psychological domains.

Acknowledgements

We thank Lauren DiNicola, Daniel Reznik, Xavier Guell, and David Somers for helpful discussions, and Daniel Sussman for initial guidance on community detection and evaluation methods. This work was supported by grants from the National Science Foundation (BCS-1755757 and BCS-1625552) and the Office of Naval Research (MURI awards N00014-16-1-2832 and N00014-17-1-2304). 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.

582 # Figures

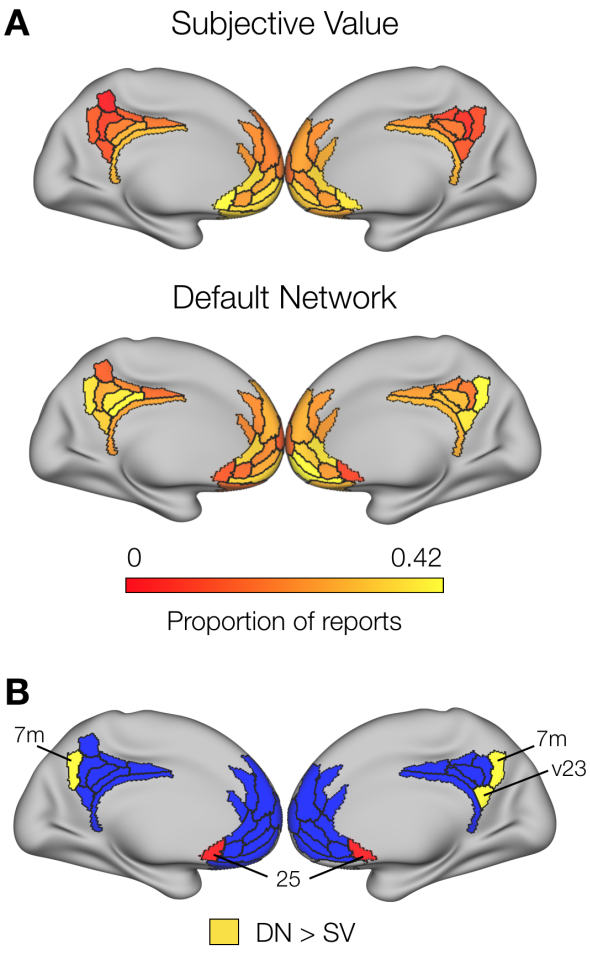

Figure 1: Meta-analysis results. A: Proportion of times each ROI was reported in the valuation and DN literatures. B: Regions identified in permutation-based chi-squared tests contrasting the two literatures (see text for details). Areas in blue represent the remainder of the search space.

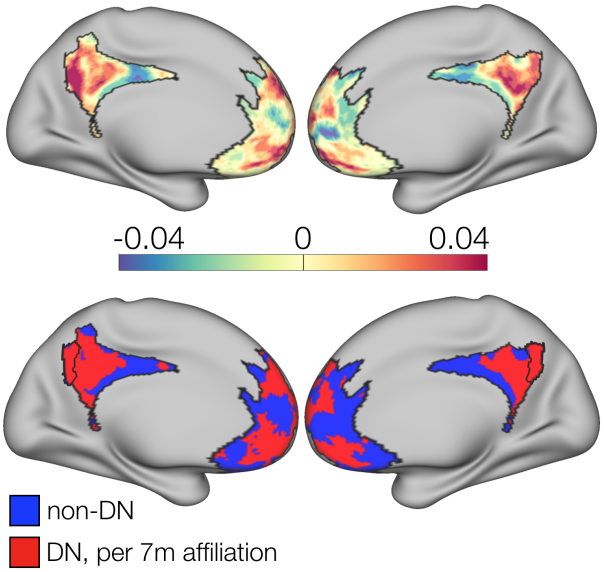

Figure 2: Brain partition for an example subject (100307). Fiedler vector values (top) are mapped onto the brain surface, dividing it into positive and negative communities. The bottom brain shows the binarized Fiedler vector, with red areas denoting the DN community (as indicated by coverage of area 7m, bordered).

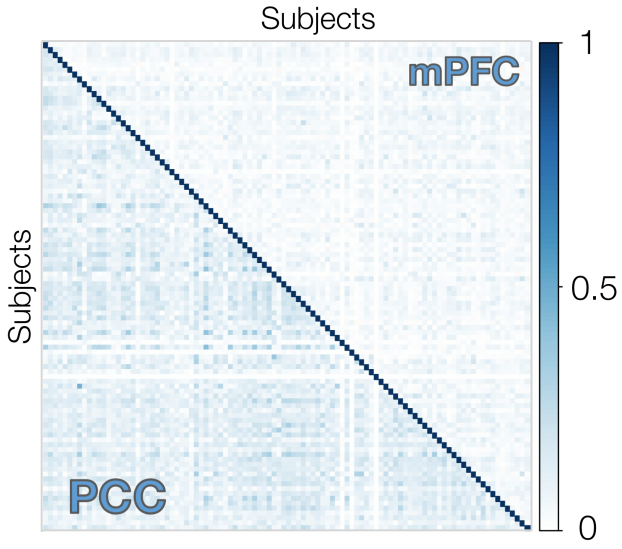

Figure 3: Similarity matrix showing ARI values among all subjects for PCC (lower triangle) and mPFC (upper triangle) separately. Functional topographic patterns were better aligned across individuals in PCC than mPFC.

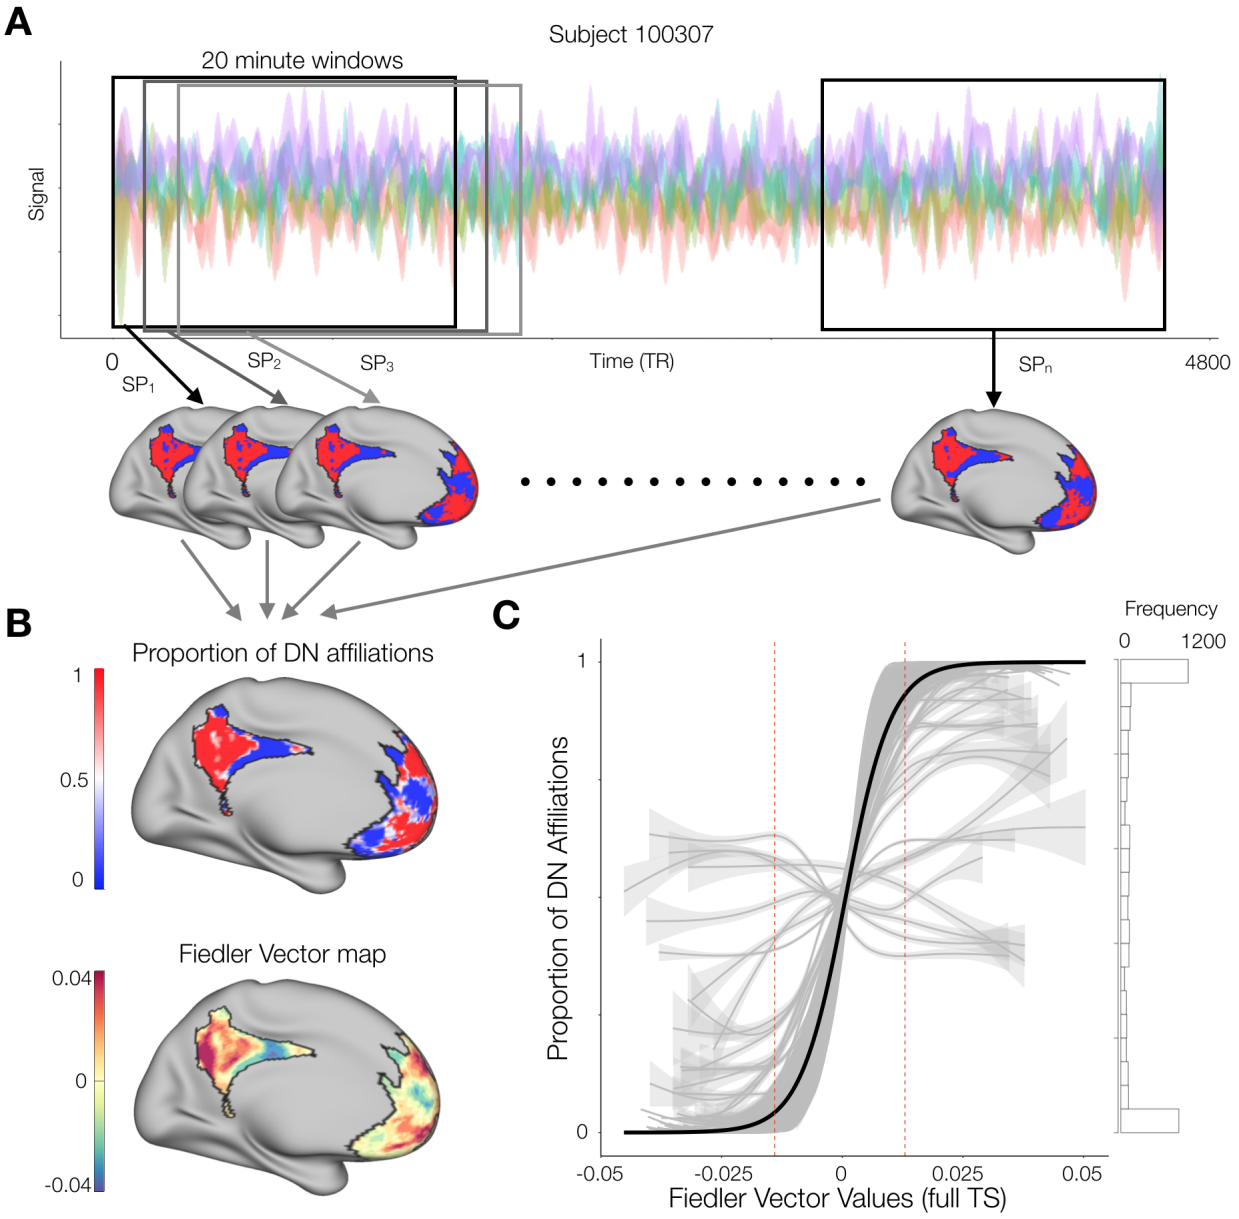

Figure 4: A: For each individual, we produced partitions for each 20 minute sliding window (84 TRs). B: Proportion of times each vertex was affiliated with the DN community across windows (upper), and the continuous Fiedler vector map for the current subject using their full time series (lower). The maps share considerable qualitative similarities in their gradients along the cortical surface. C: Relationship between the magnitude of Fiedler vector values and the proportion of DN affiliations. Grey lines display data for each subject, and the black line shows the fit from a mixed-effects logistic regression. Dashed red lines indicate the mean FV value at which maps were thresholded. The histogram displays the mean frequency distribution of y-axis values.

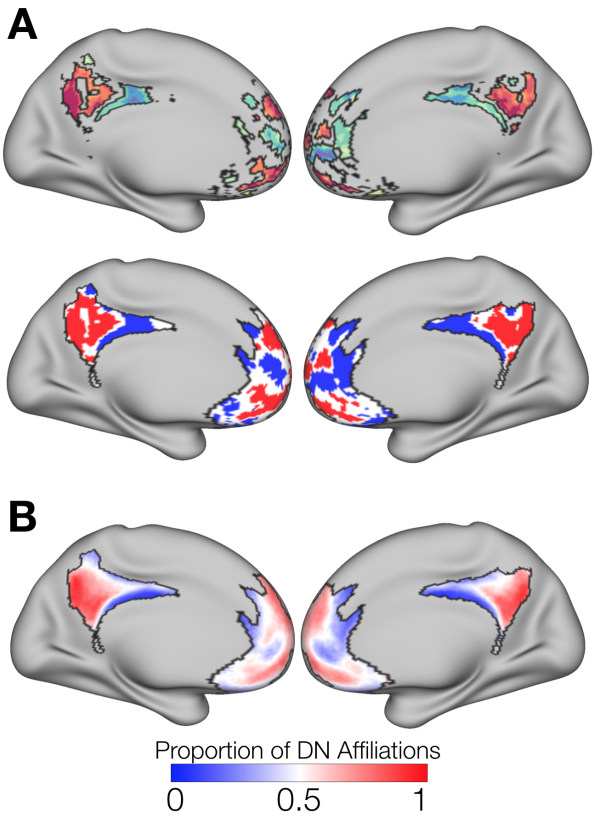

Figure 5: A: Thresholded Fiedler vector maps for subject 100307 (top), and its binarized form (bottom). Subthreshold values effectively formed a third community of high-variability vertices. B: Mean of the binarized maps across all participants, indicating the proportion of DN affiliations per vertex in our sample. This aggregate map shows the common organizational principle of the DN and non-DN communities, while also showing the high level of variability in mPFC.

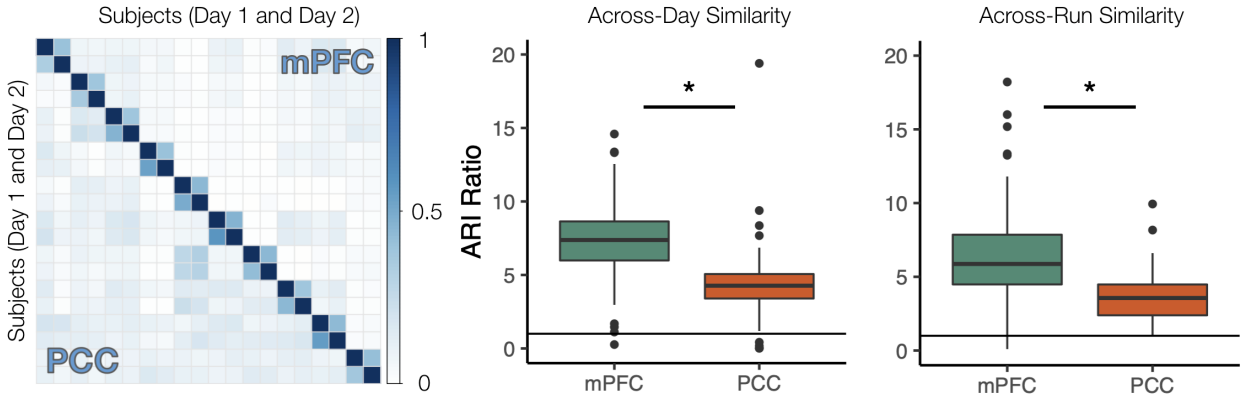

Figure 6: Left: Similarity matrix for 10 example participants (2 scanning days each), showing pattern agreement across days and subjects for PCC and mPFC separately. The block-diagonal structure is indicative of test-retest reliability across days within an individual. Middle: ratio of within-subject ARI to between-subject mean ARI for all individuals across days suggests idiosyncratic community arrangement for both PCC and mPFC (ratios > 1, solid line), with greater subject-specificity in mPFC. Right: within-to-between subject mean ARI ratios for run-specific partitionings again show greater subject-specific organization for mPFC.

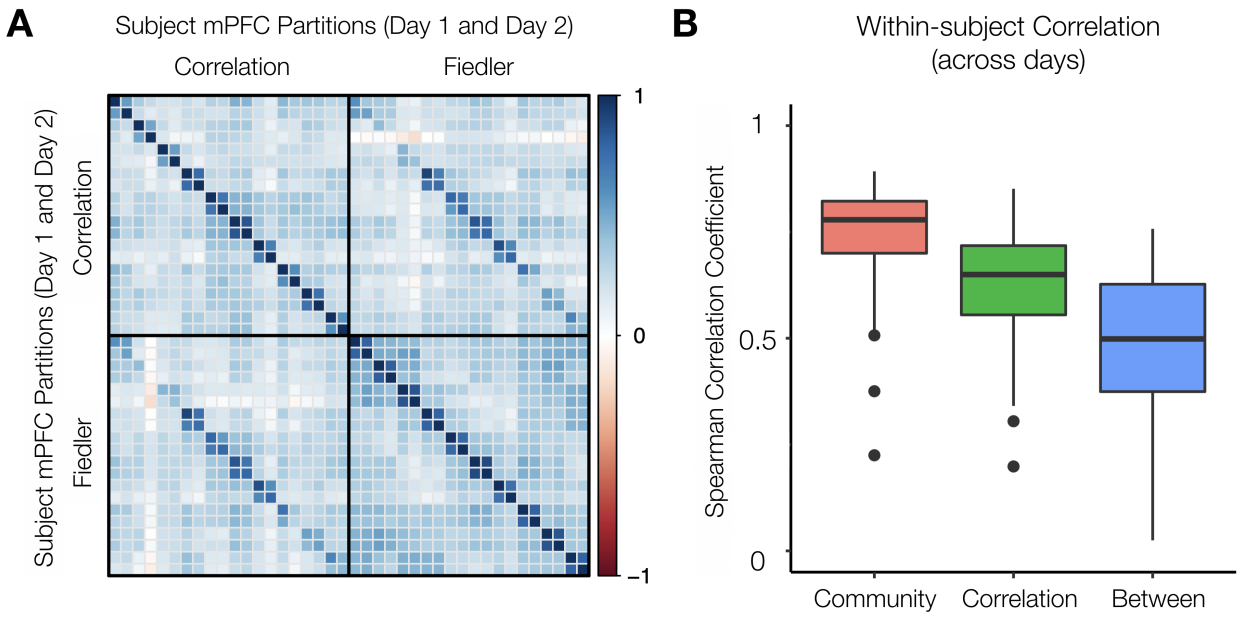

Figure 7: A: Correlation matrix comparing the across-day spatial stability of mPFC maps derived from seed-based functional connectivity (using a PCC seed) and the Fiedler vector for 10 example subjects. The top-left quadrant represents seed-based FC maps, and the bottom-right the Fiedler vector, with two single-day-based maps per individual. The upper-right and lower-left quadrants show across-method agreement. B: Day 1 vs Day 2 within-subject correlation coefficients for each method, as well as between methods. Community detection through spectral partitioning provided more stable estimates, even though both methods showed good levels of agreement.

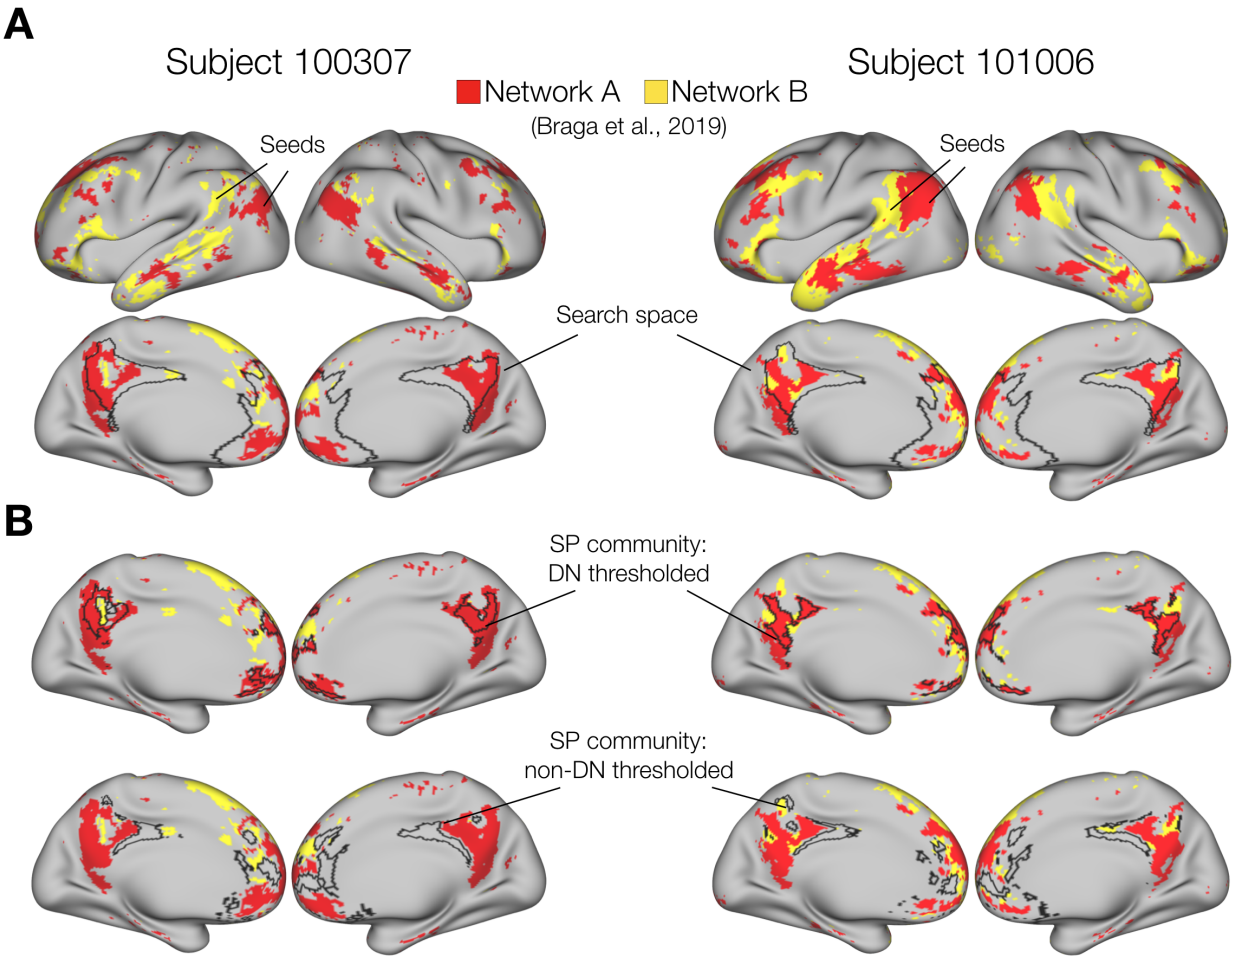

Figure 8: Qualitative comparison between DN sub-networks A and B from Braga et al. (2019) and SP communities for two individuals. Panel A: Whole-brain networks A and B produced by selecting seeds in TPJ, with our community detection search space delineated by black borders. Correlation values are thresholded at 0.2. Panel B: thresholded communities (indicated by borders) show strong resemblance between the DN community and network A. The non-DN community covers sections of cortex not associated with either DN sub-network.

Supplemental Materials

Table 1. Parcels from Glasser et al. (2016) contained in the search space.

Hemisphere	mPFC	PCC
Left	10d, 10r, 10v, 25, 9m, a24, d32, OFC, p24, p32, s32	23d, 31a, 31pd, 31pv, 7m, d23ab, PCV, RSC, v23ab
Right	10d, 10r, 10v, 25, 9a, 9m, a24, d32, OFC, p24, p32, s32	23d, 31a, 31pd, 31pv, 7m, d23ab, RSC, v23ab

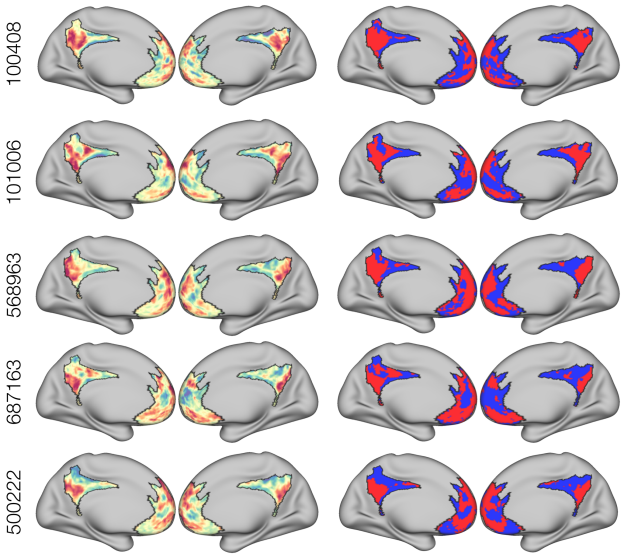

Figure S1: Additional examples of individualized partitionings, showing both Fiedler vector values (left) and binarized communities (right). A common organizational principle is visible, even though it shifts topographically across individuals.

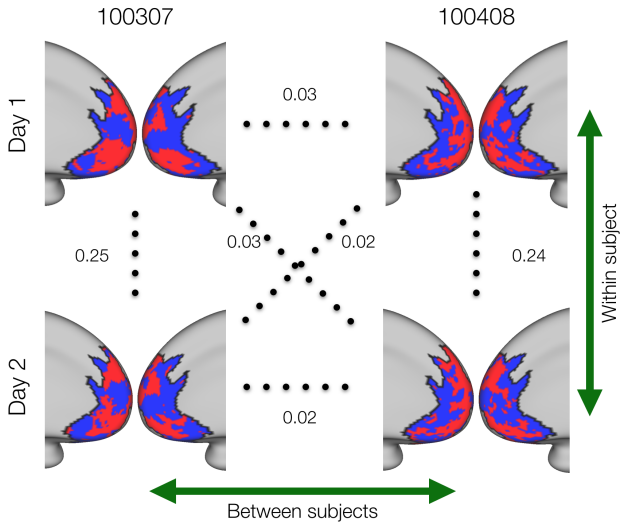

Figure S2: Example of an across-day comparison using ARI for two subjects (100307 and 100408). This reflects how qualitatively similar, within-subject partitionings can have relatively small ARI values (here 0.24-0.25), and how partitionings across individuals are much closer to the chance level of zero.

References

Acikalin, M. Y., Gorgolewski, K. J., & Poldrack, R. A. (2017). A coordinate-based meta-analysis of overlaps in regional specialization and functional connectivity across subjective value and default mode networks. *Frontiers in Neuroscience*, *11*(JAN), 1–11. https://doi.org/10.3389/fnins.2017.00001

Andrews-Hanna, J. R., Reidler, J. S., Huang, C., & Buckner, R. L. (2010). Evidence for the Default Network's Role in Spontaneous Cognition. *Journal of Neurophysiology*, *104*(1), 322–335. https://doi.org/10.1152/jn.00830.2009

Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, *76*, 412–427. https://doi.org/10.1016/j.neuroimage.2013.02.063

Bassett, D. S., & Sporns, O. (2017). Network neuroscience. *Nature Neuroscience*, *20*(3), 353–364. https://doi.org/10.1038/nn.4502

Belkin, M., & Niyogi, P. (2003). Laplacian eigenmaps for dimensionalitu reduction and data representation. *Neural Computation*, *15*, 1373–1396. https://doi.org/10.1162/089976603321780317

Braga, R. M., & Buckner, R. L. (2017). Parallel Interdigitated Distributed Networks within the Individual Estimated by Intrinsic Functional Connectivity. *Neuron*, *95*(2), 457–471.e5. https://doi.org/10.1016/j.neuron.2017.06.038

Braga, R. M., Van Dijk, K. R., Polimeni, J. R., Eldaief, M. C., & Buckner, R. L. (2019). Parallel distributed networks resolved at high resolution reveal close juxtaposition of distinct regions. *Journal of Neurophysiology*, jn.00808.2018. https://doi.org/10.1152/jn.00808.2018

Burgess, G. C., Kandala, S., Nolan, D., Laumann, T. O., Power, J. D., Adeyemo, B., ... & Barch, D. M. (2016). Evaluation of Denoising Strategies to Address Motion-Correlated Artifacts in Resting-State Functional Magnetic Resonance Imaging Data from the Human Connectome Project. *Brain Connectivity*, *6*(9), 669–680.

Chung, F. R. K. (1997). *Spectral Graph Theory* (Vol. 92). American Mathematical Soc.

Clauset, A., Newman, M. E. J., & Moore, C. (2004). *Finding community structure in very large networks*. *066111*, 1–6. https://doi.org/10.1103/PhysRevE.70.066111

Clithero, J. A., & Rangel, A. (2014). Informatic parcellation of the network involved in the computation of subjective value. *Social Cognitive and Affective Neuroscience*, *9*(9), 1289–1302. https://doi.org/10.1093/

scan/nst106

Conroy, B. R., Singer, B. D., Guntupalli, J. S., Ramadge, P. J., & Haxby, J. V. (2013). Inter-subject alignment of human cortical anatomy using functional connectivity. *NeuroImage*, *81*, 400–411. https://doi.org/10.1016/j.neuroimage.2013.05.009

Csardi, G., & Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal Complex Systems*, *1695*.

Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis: I. Segmentation and surface reconstruction. *NeuroImage*, *9*(2), 179–194. https://doi.org/10.1006/nimg.1998.0395

De La Vega, A., Chang, L. J., Banich, M. T., Wager, T. D., & Yarkoni, T. (2016). Large-Scale Meta-Analysis of Human Medial Frontal Cortex Reveals Tripartite Functional Organization. *The Journal of Neuroscience*, *36*(24), 6553–6562. https://doi.org/10.1523/JNEUROSCI.4402-15.2016

Euston, D. R., Gruber, A. J., & McNaughton, B. L. (2012). The role of medial prefrontal cortex in memory and decision making. *Neuron*, *76*(6), 1057–1070. https://doi.org/10.1016/j.neuron.2012.12.002

Fedorenko, E., Duncan, J., & Kanwisher, N. (2012). Language-selective and domain-general regions lie side by side within Broca's area. *Current Biology*, *22*(21), 2059–2062. https://doi.org/10.1016/j.cub.2012.09.011

Fiedler, M. (1975). A Property of Eigenvectors of Nonnegative Symmetric Matrices and its Application to Graph Theory. *Czechoslovak Mathematical Journal*, *25*(100), 619–633.

Fischl, B., Sereno, M. I., & Dale, A. M. (1999). Cortical surface-based analysis: II. Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, *9*(2), 195–207. https://doi.org/10.1006/nimg.1998.0396

Fortunato, S., & Hric, D. (2016). Community detection in networks: A user guide. *Physics Reports*, *659*, 1–44. https://doi.org/10.1016/j.physrep.2016.09.002

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, *102*(27), 9673–9678. https://doi.org/10.1073/pnas.0504136102

Fritsch, A. (2012). *mcclust: Process an MCMC Sample of Clusterings. R package version 1.0.* Retrieved from https://cran.r-project.org/package=mcclust

Garcia, J. O., Ashourvan, A., Muldoon, S., Vettel, J. M., & Bassett, D. S. (2018). Applications of Community Detection Techniques to Brain Graphs: Algorithmic Considerations and Implications for Neural Function.

Proceedings of the IEEE, *106*(5), 846–867. https://doi.org/10.1109/JPROC.2017.2786710

Ghasemian, A., Hosseinmardi, H., & Clauset, A. (2019). Evaluating Overfit and Underfit in Models of Network Community Structure. *IEEE Transactions on Knowledge and Data Engineering*. https://doi.org/10.1109/TKDE.2019.2911585

Gkantsidis, C., Mihail, M., & Zegura, E. (2003). Spectral analysis of Internet topologies. *Proc. IEEE INFOCOM*, *00*(C), 364–374. https://doi.org/10.1109/INFCOM.2003.1208688

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., … Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171–178. https://doi.org/10.1038/nature18933

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., … Jenkinson, M. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, *80*, 105–124. https://doi.org/10.1016/j.neuroimage.2013.04.127

Gordon, E. M., Laumann, T. O., Gilmore, A. W., Newbold, D. J., Greene, D. J., Berg, J. J., … Dosenbach, N. U. (2017). Precision Functional Mapping of Individual Human Brains. *Neuron*, *95*(4), 791–807.e7. https://doi.org/10.1016/j.neuron.2017.07.011

Gratton, C., Laumann, T. O., Nielsen, A. N., Greene, D. J., Gordon, E. M., Gilmore, A. W., … Petersen, S. E. (2018). Functional Brain Networks Are Dominated by Stable Group and Individual Factors, Not Cognitive or Daily Variation. *Neuron*, 439–452. https://doi.org/10.1016/j.neuron.2018.03.035

Greicius, M. D., Krasnow, B., Reiss, A. L., & Menon, V. (2003). *Functional connectivity in the resting brain: A network analysis of the default mode hypothesis*. *100*(1), 253–258. https://doi.org/10.1073/pnas.0135058100

Griffanti, L., Salimi-Khorshidi, G., Beckmann, C. F., Auerbach, E. J., Douaud, G., Sexton, C. E., … Smith, S. M. (2014). ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *NeuroImage*, *95*, 232–247. https://doi.org/10.1016/j.neuroimage.2014.03.034

Guntupalli, J. S., Feilong, M., & Haxby, J. V. (2018). A computational model of shared fine-scale structure in the human connectome. *PLoS Computational Biology*, *14*(4), 1–26. https://doi.org/10.1371/journal.pcbi.1006120

Higham, D. J., Kalna, G., & Kibble, M. (2007). Spectral clustering and its use in bioinformatics. *Journal of Computational and Applied Mathematics*, *204*(1), 25–37. https://doi.org/10.1016/j.cam.2006.04.026

Hiser, J., & Koenigs, M. (2018). The Multifaceted Role of the Ventromedial Prefrontal Cortex in Emotion,

Decision Making, Social Cognition, and Psychopathology. *Biological Psychiatry*, *83*(8), 638–647. https://doi.org/10.1016/j.biopsych.2017.10.030

Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, *2*(1), 193–218. https://doi.org/10.1007/BF01908075

Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625–1633.

Kernbach, J. M., Yeo, B. T. T., Smallwood, J., Margulies, D. S., Thiebaut de Schotten, M., Walter, H., … Bzdok, D. (2018). Subspecialization within default mode nodes characterized in 10,000 UK Biobank participants. *Proceedings of the National Academy of Sciences*, (November), 201804876. https://doi.org/10.1073/pnas.1804876115

Kong, R., Li, J., Orban, C., Sabuncu, M. R., Liu, H., Schaefer, A., … Yeo, B. T. T. (2018). Spatial Topography of Individual-Specific Cortical Networks Predicts Human Cognition, Personality, and Emotion. *Cerebral Cortex*, (May 2018), 2533–2551. https://doi.org/10.1093/cercor/bhy123

Kragel, P. A., Kano, M., Van Oudenhove, L., Ly, H. G., Dupont, P., Rubio, A., … Wager, T. D. (2018). Generalizable representations of pain, cognitive control, and negative emotion in medial frontal cortex. *Nature Neuroscience*, 1. https://doi.org/10.1038/s41593-017-0051-7

Laird, A. R., Eickhoff, S. B., Li, K., Robin, D. A., Glahn, D. C., & Fox, P. T. (2009). Investigating the Functional Heterogeneity of the Default Mode Network Using Coordinate-Based Meta-Analytic Modeling. *Journal of Neuroscience*, *29*(46), 14496–14505. https://doi.org/10.1523/JNEUROSCI.4004-09.2009

Laumann, T. O., Gordon, E. M., Adeyemo, B., Snyder, A. Z., Joo, S. J., Chen, M. Y., … Petersen, S. E. (2015). Functional System and Areal Organization of a Highly Sampled Individual Human Brain. *Neuron*, *87*(3), 658–671. https://doi.org/10.1016/j.neuron.2015.06.037

Levy, I., Lazzaro, S. C., Rutledge, R. B., & Glimcher, P. W. (2011). Choice from Non-Choice: Predicting Consumer Preferences from Blood Oxygenation Level-Dependent Signals Obtained during Passive Viewing. *Journal of Neuroscience*, *31*(1), 118–125. https://doi.org/10.1523/JNEUROSCI.3214-10.2011

Lopez-Persem, A., Verhagen, L., Amiez, C., Petrides, M., & Sallet, J. (2019). The human ventromedial prefrontal cortex: sulcal morphology and its influence on functional organization. *The Journal of Neuroscience*, *39*(19), 2060–2018. https://doi.org/10.1523/JNEUROSCI.2060-18.2019

Mackey, S., & Petrides, M. (2014). Architecture and morphology of the human ventromedial prefrontal cortex.

European Journal of Neuroscience, *40*(5), 2777–2796. https://doi.org/10.1111/ejn.12654

Mars, R. B., Passingham, R. E., & Jbabdi, S. (2018). Connectivity Fingerprints: From Areal Descriptions to Abstract Spaces. *Trends in Cognitive Sciences*, *22*(11), 1026–1037. https://doi.org/10.1016/j.tics.2018.08.009

McKiernan, K. A., Kaufman, J. N., Kucera-Thompson, J., & Binder, J. R. (2003). A Parametric Manipulation of Factors Affecting Task-induced Deactivation in Functional Neuroimaging. *Journal of Cognitive Neuroscience*, *15*, 394–408. https://doi.org/10.1162/089892903321593117

Michalka, S. W., Kong, L., Rosen, M. L., Shinn-Cunningham, B. G., & Somers, D. C. (2015). Short-Term Memory for Space and Time Flexibly Recruit Complementary Sensory-Biased Frontal Lobe Attention Networks. *Neuron*, *87*(4), 882–892. https://doi.org/10.1016/j.neuron.2015.07.028

Mitchell, J. P., Banaji, M. R., & Macrae, C. N. (2005). The link between social cognition and self-referential thought in the medial prefrontal cortex. *Journal of Cognitive Neuroscience*, *17*(8), 1306–1315.

Mueller, S., Wang, D., Fox, M. D., Yeo, B. T. T., Sepulcre, J., Sabuncu, M. R., ... Liu, H. (2013). Individual Variability in Functional Connectivity Architecture of the Human Brain. *Neuron*, *77*(3), 586–595. https://doi.org/10.1016/j.neuron.2012.12.028

Northoff, G., & Hayes, D. J. (2011). Is our self nothing but reward? *Biological Psychiatry*, *69*, 1019–1025. https://doi.org/10.1016/j.biopsych.2010.12.014

Osher, D. E., Saxe, R. R., Koldewyn, K., Gabrieli, J. D., Kanwisher, N., & Saygin, Z. M. (2016). Structural Connectivity Fingerprints Predict Cortical Selectivity for Multiple Visual Categories across Cortex. *Cerebral Cortex*, *26*(4), 1668–1683. https://doi.org/10.1093/cercor/bhu303

Passingham, R. E., Stephan, K. E., & Kötter, R. (2002). The anatomical basis of functional localization in the cortex. *Nature Reviews Neuroscience*, *3*(8), 606–616. https://doi.org/10.1038/nrn893

R Core Computing Team. (2017). *R: A Language and Environment for Statistical Computing*. Retrieved from https://www.r-project.org/

Reddan, M. C., Wager, T. D., & Schiller, D. (2018). Attenuating Neural Threat Expression with Imagination. *Neuron*, *100*(4), 994–1005.e4. https://doi.org/10.1016/j.neuron.2018.10.047

Rubinov, M., & Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage*, *52*(3), 1059–1069. https://doi.org/10.1016/j.neuroimage.2009.10.003

Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., & Smith, S. M. (2014). Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical

fusion of classifiers. *NeuroImage*, *90*, 449–468. https://doi.org/10.1016/j.neuroimage.2013.11.046

Saygin, Z. M., Osher, D. E., Koldewyn, K., Reynolds, G., Gabrieli, J. D., & Saxe, R. R. (2011). Anatomical connectivity patterns predict face selectivity in the fusiform gyrus. *Nature Neuroscience*, *15*(2), 321–327. https://doi.org/10.1038/nn.3001

Saygin, Z. M., Osher, D. E., Norton, E. S., Youssoufian, D. A., Beach, S. D., Feather, J., ... Kanwisher, N. (2016). Connectivity precedes function in the development of the visual word form area. *Nature Neuroscience*, *19*(9), 1250–1255. https://doi.org/10.1038/nn.4354

Shenhav, A., & Karmarkar, U. R. (2019). Dissociable components of the reward circuit are involved in appraisal versus choice. *Scientific Reports*, *9*(1), 172320. https://doi.org/10.1038/s41598-019-38927-7

Smith, D. V., Hayden, B. Y., Truong, T.-K., Song, A. W., Platt, M. L., & Huettel, S. A. (2010). Distinct Value Signals in Anterior and Posterior Ventromedial Prefrontal Cortex. *Journal of Neuroscience*, *30*(7), 2490–2495. https://doi.org/10.1523/JNEUROSCI.3319-09.2010

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., ... Beckmann, C. F. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, *106*(31), 13040 LP–13045. https://doi.org/10.1073/pnas.0905267106

Tian, Y., & Zalesky, A. (2018). Characterizing the functional connectivity diversity of the insula cortex: Subregions, diversity curves and behavior. *NeuroImage*, *183*, 716–733. https://doi.org/10.1016/j.neuroimage.2018.08.055

Tobyne, S. M., Osher, D. E., Michalka, S. W., & Somers, D. C. (2017). Sensory-biased attention networks in human lateral frontal cortex revealed by intrinsic functional connectivity. *NeuroImage*, *162*(August), 362–372. https://doi.org/10.1016/j.neuroimage.2017.08.020

Tobyne, S. M., Somers, D. C., Brissenden, J. A., Michalka, S. W., Noyce, A. L., & Osher, D. E. (2018). Prediction of individualized task activation in sensory modality-selective frontal cortex with 'connectome fingerprinting'. *NeuroImage*, *183*(August), 173–185. https://doi.org/10.1016/j.neuroimage.2018.08.007

Toker, D., & Sommer, F. T. (2019). Information Integration In Large Brain Networks. *PLoS Computational Biology*, *15*(2). https://doi.org/10.1371/journal.pcbi.1006807

Van Essen, D. C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T. E., Bucholz, R., ... Yacoub, E. (2012). The Human Connectome Project: A data acquisition perspective. *NeuroImage*, *62*(4), 2222–2231.

https://doi.org/10.1016/j.neuroimage.2012.02.018

Wager, T. D., Lindquist, M. A., Nichols, T. E., Kober, H., & Van Snellenberg, J. X. (2009). Evaluating the consistency and specificity of neuroimaging data using meta-analysis. *NeuroImage*, *45*(1 Suppl), S210–S221. https://doi.org/10.1016/j.neuroimage.2008.10.061

Wang, D., Buckner, R. L., Fox, M. D., Holt, D. J., Holmes, A. J., Stoecklein, S., ... Liu, H. (2015). Parcellating cortical functional networks in individuals. *Nature Neuroscience*, *18*(12), 1853–1860. https://doi.org/10.1038/nn.4164

Woo, C.-W., Koban, L., Kross, E., Lindquist, M. A., Banich, M. T., Ruzic, L., ... Wager, T. D. (2014). Separate neural representations for physical pain and social rejection. *Nature Communications*, *5*(May), 1–12. https://doi.org/10.1038/ncomms6380

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

Zilles, K., Palomero-Gallagher, N., & Amunts, K. (2013). Development of cortical folding during evolution and ontogeny. *Trends in Neurosciences*, *36*(5), 275–284. https://doi.org/10.1016/j.tins.2013.01.006