



Published in final edited form as:

Neuroimage. 2019 November 15; 202: 115990. doi:10.1016/j.neuroimage.2019.07.003.

The community structure of functional brain networks exhibits scale-specific patterns of inter- and intra-subject variability

Richard F. Betzel^{1,2,3}, Maxwell A. Bertolero¹, Evan M. Gordon^{4,5}, Caterina Gratton^{6,7}, Nico U.F. Dosenbach^{6,8,9}, Danielle S. Bassett^{1,10,11,12,*}

¹Department of Bioengineering, School of Engineering and Applied Sciences, University of Pennsylvania, Philadelphia, PA, 19104 USA

²Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN, 47401 USA

³Cognitive Science Program, Indiana University, Bloomington, IN, 47401 USA

⁴VISN 17 Center of Excellence for Research on Returning War Veterans, Waco, TX, 76711 USA

⁵Center for Vital Longevity, School of Behavioral and Brain Sciences, University of Texas at Dallas, Dallas, TX, 75235, USA

⁶Department of Neurology, Washington University School of Medicine, St. Louis, MO, 63110, USA

⁷Department of Psychology, Northwestern University, Evanston, IL, 60208, USA

⁸Program in Occupational Therapy, Washington University School of Medicine, St. Louis, MO, 63110, USA

⁹Department of Pediatrics, Washington University School of Medicine, St. Louis, MO, 63110, USA

¹⁰Department of Electrical & Systems Engineering, School of Engineering and Applied Sciences, University of Pennsylvania, Philadelphia, PA, 19104 USA

¹¹Department of Neurology, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, 19104 USA

¹²Department of Physics & Astronomy, School of Arts and Sciences, University of Pennsylvania, Philadelphia, PA, 19104 USA

Abstract

The network organization of the human brain varies across individuals, changes with development and aging, and differs in disease. Discovering the major dimensions along which this variability is

* dsb@seas.upenn.edu.

AUTHOR CONTRIBUTIONS

This study was designed, carried out, and written by RFB. DSB secured funding for the work. MAB processed and provided HCP MRI and behavioral data. EMG, CG, NUFD processed and provided MSC MRI data. All authors contributed to the direction of the research and edited the paper.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

displayed remains a central goal of both neuroscience and clinical medicine. Such efforts can be usefully framed within the context of the brain's modular network organization, which can be assessed quantitatively using powerful computational techniques and extended for the purposes of multi-scale analysis, dimensionality reduction, and biomarker generation. Though the concept of modularity and its utility in describing brain network organization is clear, principled methods for comparing multi-scale communities across individuals and time are surprisingly lacking. Here, we present a method that uses multi-layer networks to simultaneously discover the modular structure of many subjects at once. This method builds upon the well-known *multi-layer modularity maximization* technique, and provides a viable and principled tool for studying differences in network communities across individuals and within individuals across time. We test this method on two datasets and identify consistent patterns of inter-subject community variability, demonstrating that this variability – which would be undetectable using past approaches – is associated with measures of cognitive performance. In general, the multi-layer, multi-subject framework proposed here represents an advancement over current approaches by straightforwardly mapping community assignments across subjects and holds promise for future investigations of inter-subject community variation in clinical populations or as a result of task constraints.

INTRODUCTION

The human brain is a complex network of functionally interconnected brain areas. Its architecture is strikingly non-random across a spectrum of scales, ranging from the local scale of individual brain areas to the global scale of the entire brain [1, 2]. Situated between these two extremes is the meso-scale comprised of sub-networks of topologically-related neural elements referred to as “communities” or modules [3, 4]. The brain's community structure reflects regularities in its wiring diagram, delineating groups of brain areas with shared functionality [5, 6]. Critically, the brain's community structure spans multiple organizational scales, ranging from small communities associated with functionally-specialized areas (the scale measurable with MRI) to larger communities associated with more general brain and cognitive functions [2].

Increasingly, the brain's community structure has become the focus of many investigations. By characterizing the variability of community structure across individuals or between clinical populations, recent studies seek a deeper understanding of neuropsychiatric disorders [7, 8], development and aging [9–11], and diverse cognitive processes [12]. Despite such broad interest, there remains a paucity of principled methods for detecting and comparing communities across individuals, and little consensus on which approach maximizes advantages and minimizes disadvantages in the context of neuroscientific inquiry. Virtually all extant community detection methods rely on heuristics or make strong assumptions about the number and size of communities, the consistency of communities across individuals, and the nature of community identity as one maps community structure from one subject to another. Although these assumptions are made in order to facilitate further analyses, they can also entail unwanted biases, thereby limiting inference.

Broadly, existing approaches for comparing community structure across individuals exhibit both striking benefits and marked limitations. Consider, as an example, the popular

consensus approach, in which group-level communities are imposed uniformly across all subjects¹. This approach confers two notable advantages. First, the uniformity of communities permits straightforward comparisons across individuals. Second, because the communities are defined at the group level, they are likely less susceptible to overfitting. Yet, a notable disadvantage of the approach is that it precludes the possibility that communities vary across individuals. While this assumption of conservation of community structure across individuals might be reasonable in certain analyses of neurotypical cohorts (although even this assumption may be too strong; see [15]), it becomes increasingly problematic in clinical populations where heterogeneity in pathology leads to patient-specific disruptions in neuroanatomy and physiology.

Other approaches can overcome these specific issues, but are limited in different ways. Data-driven community detection methods [16, 17] naturally accommodate inter-individual community variability, but can overfit noisy network data or result in situations where the mapping of communities across individuals is ambiguous or even impossible. This latter issue can be partially mitigated using template-matching techniques, such as those that register detected communities to a common set of “canonical” communities [15]. However, template-matching requires users to specify a set of template communities, restricting investigation to a single topological scale, and precluding hierarchical or multi-scale community analysis [2, 18, 19]. Moreover, template communities are oftentimes established based on resting-state data whose relationship with task-evoked community structure remains unclear [12].

Here, we propose a framework that flexibly accommodates multi-scale community analysis while unambiguously mapping communities from one subject to another. This framework builds on the well-known technique of modularity maximization [16], which algorithmically decomposes networks into internally cohesive communities and has recently been extended to be compatible with multi-layer networks [20]. Past applications of multi-layer modularity maximization have been largely restricted to so-called time-varying networks, where each layer represents a snapshot of a functional brain network localized to a particular window in time [21–23]. We present a modification in which network layers represent connectivity data from single subjects that are then made interdependent upon one another through the addition of inter-layer couplings.

Here, we apply multi-layer, multi-subject modularity maximization to functional connectivity data acquired as part of the Human Connectome Project (HCP) [24] and the “Midnight Scan Club” (MSC) [25]. First, we show that this approach naturally resolves ambiguities related to the mapping of communities across subjects while simultaneously recapitulating the topography of known resting-state and intrinsic connectivity networks. Next, we show that community structure varies across subjects along “modes” that are aligned with distinct organizational scales. Importantly, we show that the association of cognitive performance measures with community variability is also scale-dependent, emphasizing the necessity of detecting and analyzing community structure at multiple

¹Group-level community labels can be obtained in a number of ways, including data-driven methods such as community detection [13] or system labels taken from canonical brain atlases [5, 14]

topological scales. Finally, using MSC data, we replicate modes of inter-subject community variability, but show that these modes differ from those associated with *intra*-subject community variability, reaffirming recent findings that variability of network organization within individuals is unique and may be a powerful source of behavioral variation. Our findings showcase the relevance of multi-scale community analysis and present methodology that will reduce the ambiguity typically associated with mapping and comparing communities across cohorts of individuals.

RESULTS

The brain's community structure reflects cohesive groups of functionally related brain areas and spans multiple organizational scales. inter-subject variability in the community assignments of particular brain areas has been associated with an individual's disease, developmental, and cognitive state [7–11]. Current methods for studying this variability suffer from performance issues and tradeoffs that limit their utility. Here, we present an extension of the multi-layer modularity maximization framework to accommodate multi-subject datasets. This extension addresses many of the existing shortcomings associated with current methods and seamlessly maps communities across individuals and scales. In this section, we present the results of applying the multi-layer, multi-subject modularity maximization framework to functional connectivity data taken from the Human Connectome Project (HCP) [24] and the “Midnight Scan Club” (MSC) [25].

Detecting multi-layer, multi-subject community structure in the HCP dataset

Basic analysis—Characterizing inter-subject variability in community structure can provide valuable behavioral and clinical insight. Here, we examine patterns of inter-subject community variability using a novel community detection approach, which we apply to functional connectivity (FC) data made available as part of the HCP dataset. Specifically, we analyze “discovery” and “replication” cohorts each composed of $T = 80$ subjects (see Materials and Methods for preprocessing and cohort definition details). Our proposed community detection approach is based on a multi-layer variant [20] of the well-known modularity maximization framework [16]. In this approach, we treat subject's connectivity matrices as “layers” that are consolidated in a unified multi-layer network, which serves as input to the community detection algorithm. By applying modularity maximization to a single multi-layer network object, we can detect communities in all layers (i.e. subjects) simultaneously. Critically, this allows us to preserve community identity across subjects. That is, two brain areas with the same community label are treated as members of the same community, irrespective of whether they correspond to different parts of the brain or appear in unique subjects. This feature of multi-layer modularity maximization allows us to trivially map community assignments from one subject to another (Fig. 1).

Multi-layer modularity maximization depends upon two free parameters. The structural resolution parameter, γ , determines the size of communities: smaller or larger values of γ result in correspondingly larger or smaller communities. The inter-subject coupling parameter, ω , determines the consistency of communities across layers, which in our case represent subjects: smaller or larger values of ω emphasize community organization that is

either unique to individual subjects or shared by the entire cohort. Usually, applications using multi-layer modularity maximization focus on a restricted subset of the $\{\gamma, \omega\}$ parameter space, resulting in communities of characteristic size and variability. Here, however, we develop a procedure to efficiently sample a much larger region of $\{\gamma, \omega\}$ parameter space (see Materials and Methods for details). In addition, this procedure estimates the boundaries of a subspace of $\{\gamma, \omega\}$ parameter space in which detected partitions result in more than 1 community but fewer than N communities and in which intersubject variability is non-extreme, meaning that there exists at least one detected community that is shared across subjects. As a result, the effective parameter space sampled by this procedure is non-rectangular. See, for example, Fig. 2a.

Using this procedure, we generated 40000 samples of multi-subject community structure, i.e. simultaneous estimates of each brain area's community assignment for each subject in the cohort. We first aimed to benchmark the sampling procedure to confirm that it was capable of detecting meaningful community structure at different topological scales. To this end, we examined subdivisions of parameter space where the algorithm detected consensus partitions of the brain into 2, 5, 8, 11, 14, and 17 communities (note: these numbers excluded singleton communities; Fig. 2a). For each community within a consensus partition, we calculated its homogeneity with respect to the 13 systems identified by [26]. That is, a community comprised only of nodes assigned to the *Auditory* system would have greater homogeneity than a community comprised of nodes belonging to both the *Auditory* and *Saliency* systems. In Fig. 2b and c we show examples of detected communities. Each community and its constituent nodes are colored as the weighted sum of those nodes' systems' RGB triplets. As expected, the homogeneity increased with the number of communities, and the detected communities approximated the systems from [26].

We further characterized each sampled partition by measuring the variability of brain area i 's community assignment across subjects as the normalized entropy, h_i . Intuitively, the value of h_i is equal to zero when i has the same community assignment across individuals and is close to 1 when i 's assignment is less consistent. The $N \times 1$ vector $H = \{h_i\}$, therefore, encodes the pattern of community variability across the entire brain. Accordingly, the average normalized entropy, $\bar{H} = \frac{1}{N} \sum_{i=1}^N h_i$, served as an index of community variability across brain areas and individuals.

As expected, we found that the number of communities varied monotonically with γ , the structural resolution parameter that shifts the scale at which communities are detected (Fig. 3a). Smaller values of γ generally resulted in larger communities while larger values of γ resulted in smaller communities. We also found that the number of non-singleton communities peaked at an intermediate value of $\gamma \approx 1.1$ (Fig. 3b), so that the number of singleton communities increased with γ . In addition, we observed that the mean normalized entropy varied considerably over the full parameter space, and was greatest when both the inter-subject coupling and structural resolution parameters, ω and γ , were small (Fig. 3c).

Our community detection approach was designed to uncover communities of different sizes and with varying inter-subject consistency. However, some hypotheses may be easier to test by focusing on communities with a more restricted set of characteristics and statistics, such

as a given number or size or a given level of inter-subject variability. Our approach naturally allows us to test these more focused hypotheses. As an example, we could specify the set of all partitions resulting in six communities (Fig. 3d; blue points) or the set of all partitions resulting in a particular level of inter-subject variability (average normalized entropy between 0.2 and 0.25) (Fig. 3d; red points). These partitions, or even their intersection (Fig. 3d; yellow points), could be extracted for additional analyses, allowing for a more detailed and nuanced exploration of community structure across subjects. As an example, we show detected communities and community entropies in Fig. 3e,f corresponding to one of the yellow points in Fig. 3d.

Collectively, these observations illustrate the mechanisms by which modularity maximization can be used to generate estimates of multi-subject community structure. Due to the two free parameters in the optimization, the method has marked utility in detecting communities at different organizational scales and with varying levels of consistency across subjects, motivating further characterization of community structure across the γ, ω plane.

Principal component analysis and modes of inter-subject variation—The sampling procedure described in the previous section generated 40000 estimates of multi-scale, multi-subject community structure. For each sample, we characterized the inter-subject variability of communities as a normalized entropy vector, H . An important practical question is whether these patterns of variability are themselves variable, and if so, whether that variability is structured in some meaningful way. If, for example, community structure varies across subjects differently depending upon the size and the number of detected communities, such variability could have profound implications for any study of community-level correlates of behavioral measures or clinical scores. To address this possibility, we stored the full set of entropy vectors in a 333×40000 matrix (each column corresponds to a single sample). We column-normalized this matrix and then performed a principal component analysis (PCA), generating 332 orthonormal vectors (principal component scores) and their relative contribution to each of the 40000 entropy estimates (principal component coefficient). That is, PC scores generate brain maps while PC coefficients are defined in the parameter space. Intuitively, these can be thought of as “modes” by which communities varied across subjects.

To test whether the PC scores generated here were meaningful, both practically and statistically, we performed two confirmatory analyses. First, we compared PCs calculated from discovery and replication datasets. In general, we found excellent correspondence across cohorts, with strong one-to-one PC score correlations persisting over, at least, the first twenty PCs (See Fig. S1a). Second, we compared the cumulative variance explained by PCs in the discovery cohort with the cumulative variance explained under a null model in which elements of the entropy matrix were permuted randomly and independently within columns. We found that the variance explained by the observed data was greater than that of the null model for up to thirteen components ($p \approx 0$ based on 100 repetitions of a null model in which the elements of entropy vectors were uniformly and randomly permuted; See Fig. S1b,c). We also note that, in general, the first few principal components explain a large fraction of the variance in community structure, suggesting that inter-subject variability, although taking place along multiple modes, is nonetheless relatively low-dimensional.

Collectively, these findings indicate that PC scores were largely replicable across two non-overlapping groups of human participants and that the total variance explained by the first few PCs exceeded that of a chance model.

Next, we characterized PC properties including their localization in parameter space and cortical topography. First, we observed that the PCs were highly localized, both in terms of their location in parameter space as well as their cortical topography. Focusing on the first four components collectively explaining $\approx 78\%$ variance, we found that PC coefficients were largely non-overlapping, tiling the parameter space and varying as a function of the structural resolution parameter, γ (Fig. 4a,d,g,j). This tiling phenomenon indicates that the “modes” by which communities vary across subjects are scale-dependent, varying with the number and size of detected communities.

Next, we focused on the topographic distribution of PCs across the cortical surface. We found that the first component, PC_1 , which corresponded to relatively small values of γ where communities were large and few in number, implicated areas that make up the default mode (DMN), cingulo-parietal (CP), and salience systems (Fig. 4b,c). In particular, the PC scores of regions within the DMN were low (high consistency across individuals) while the PC scores of regions within the CP and SAL were high (low consistency and high variability across individuals). Across the first four principal components, different patterns of inter-subject variability were apparent. For example, in the case of PC_3 , we observed that the community assignments of the retrosplenial-temporal system (RT) were highly variable across individuals, while the dorsal attention (DAN), fronto-parietal (FP), and somatomotor (SMhand; SMmouth) systems were particularly stable (Fig. 4e,f). In the case of PC_4 , on the other hand, variability in virtually all primary sensory systems tended to be low, including auditory (AUD), somatomotor (SMhand; SMmouth), and visual (VIS), while the variability in community assignments across higher-order systems tended to be high (Fig. 4k,l). Details concerning additional components can be found in Fig. S2. Notably, we find similar patterns of variability using a different parcellation of the brain into $N=200$ regions [19] (Fig. S3).

To better understand the origins of each PC, we performed additional characterizations. Whereas the previous analysis described PCs in terms of canonical systems from [26], this additional analysis contextualized PCs with respect to detected communities themselves. Briefly, we focused only on the sub-region of parameter space for which $\omega < 10^{-2}$, as virtually all inter-subject variation occurred as a function of γ within this regime. We then partitioned this space into four segments according to the “dominant component,” i.e. the PC with the greatest z-scored average coefficient (Fig. 5a). Separately and within each regime, we calculated the co-assignment probability for every pair of nodes and derived consensus communities (Fig. 5b-i).

Each segment corresponds to a different PC. Our aim was to describe, in greater detail, the typical community structure within each segment and to understand the differences between segments. The first segment, which corresponded to the smallest values of γ and PC_1 , was composed of four communities: the first two were large and divided the cerebral cortex broadly along the association-sensory dimension. The first community was comprised of default (DMN) and fronto-parietal (FP) networks, while the second community consisted of

auditory (AUD), motor (SMhand; SMmouth), and visual plus cingulo-opercular (CO) and dorsal attention (DAN) systems (Fig. 5j). The smaller communities consisted of cingulo-parietal (CP) and salience (SAL) networks, which were also among the most variable across subjects (Fig. 5n), suggesting that inter-subject variability at this coarse scale may be driven by differences in aligning this small community within the two larger communities across subjects.

The next three segments involved progressive refinement of these two larger communities as they sub-divide into smaller clusters. In the second segment, which corresponds to PC3, the sensory network fragments, with the dorsal attention (DAN) and cingulo-opercular (CO) networks merging with the cingulo-parietal (CP) network, leaving the auditory (AUD), somatomotor (SMhand; SMmouth), and visual (VIS) networks as a community of their own (Fig. 5k). This sensory network is further refined within the third segment, corresponding to PC2, in which the visual network breaks away to form its own community. In the same segment, the default (DMN) and fronto-parietal (FP) community, which had to this point remained largely intact since the first segment, divided into separate communities (Fig. 5l). In the final segment, which corresponded to PC4, the cingulo-opercular (CO) and dorsal attention (DAN) systems sub-divide into distinct clusters (Fig. 5m). Each step in this progression is also accompanied by changes in inter-subject community variability, beginning with a stable default mode (DMN), (Fig. 5n) and terminating with stable sensory systems (Fig. 5q), suggesting that γ -dependent misalignment of subjects to specific communities within each consensus partition may drive distinct modes of variability.

Taken together, these findings suggest that inter-subject community variability is not well-summarized by a single spatial pattern nor is it localized to a particular cognitive or functional system. Rather, the variability of community assignments across subjects depends on topological scale, as operationalized by the number and size of identified communities. This point is important, as the number and size of detected communities is usually determined by a user-defined resolution parameter [27, 28], implying that the pattern of inter-subject variability in communities can effectively be tuned by the user. This fact has important implications for applications in which one wishes to understand the relation between some aspect of network community structure and a clinical or cognitive outcome. In particular, these findings suggest that patterns of inter-subject variability can be modulated by a resolution parameter, which in principle could result in different patterns of community-behavior correlations. We explore this possibility in greater detail in the next section.

Brain-behavior correlations are scale dependent—One approach for relating brain network architecture to behavior involves associating measures of community structure with a behavioral, cognitive, or clinical measure of interest. In the previous section, we demonstrated that community structure varies across subjects along distinct scale-dependent modes, suggesting that brain-behavior associations could exhibit similar dependencies. This observation has important practical implications; if brain-behavior correlations are multi-scale, then any study focusing on a single organizational scale may fail to fully characterize all relevant patterns of brain-behavior correlations.

In principle, the very fact that we observe scale-specific modes of community variation indicates that any comparison of community structure with behavioral or psychometric data would also exhibit scale specificity. To show that this is indeed the case, we calculated a subject-level analog of the normalized entropy measure, $h_{i,r}$. This score measured for each node, i , and for each subject, r , the fraction of all other subjects in which i 's community assignment differed at a given point in parameter space (see Materials and Methods). Intuitively, this measure assesses the similarity of that node's community assignment to its assignment in other subjects. We calculated this measure for each subject and assessed whether its variability across subjects was correlated with any of four behavioral indices related to social (SOC) or relational cognition (REL), language (LANG), or working memory (WM) task performance. Note that the derivation of these indices has been described elsewhere [29, 30] and are also briefly summarized in Materials and Methods.

In this section we aim to assess whether there exists evidence that brain-behavior correlations are scale-dependent. That is, we aim to address the question of whether our choice of community detection parameters γ and ω influences the brain-wide pattern of correlations between community structure and behavioral measures. Whereas the goal of many neuroimaging analyses is to identify strong associations between community structure and a particular behavioral variable of interest, we aim simply to show that the correlation pattern varies as a function of γ and ω . In this particular example, we are not necessarily seeking strong correlations, nor are we emphasizing a particular behavioral variable more than any other. In fact, three of the four measures that we study are strongly correlated with one another, implying that if we were to observe a scale-dependent effect in one of these measures, it is likely that we would find a similar effect in the remaining two (see Fig. S8). Our aim is simply to assess whether correlation patterns vary across the parameter space.

To assess whether this was the case, we identified the points in parameter space where PC coefficients were greatest for the first four PCs (Fig. 6a). Separately for each PC, we computed the average correlation of inter-individual community variability with behavioral indices for each node. We illustrate this procedure more clearly in Fig. S4. The result of this procedure was a series of sixteen vectors of length $[333 \times 1]$. The i th element of a given vector represented the correlation of brain region i 's inter-individual community variability with one of four behavioral measures, assayed separately for each of four PCs. To assess the spatial similarity of these brain-wide correlation patterns across PCs and behavioral measures, we computed their spatial similarity as a 16×16 correlation matrix (Fig. 6b). In general, we observed that the cortical topography of these correlation patterns varied both across behavioral indices and across the different PCs (Fig. 6b). To further illustrate this observation, we show examples of brain-wide correlation patterns of subject-level normalized entropy with the working memory (WM) behavioral index (Fig. 6c-j; the results for other behavioral indices are included in Supplementary Figs. S5 and S7). In general, the correlation pattern varied across parameter space, although most variation occurred as a function of γ . Depending upon where in the (γ, ω) parameter space the correlations were computed, the correlation magnitude was largest within different brain systems. For example, examining correlation patterns at parameter values with strong loadings onto PC_1 , we find evidence of system-level correlations within cingulo-opercular (CO) and visual (VIS) systems, and anti-correlations in default mode (DMN) and both dorsal/ventral

attention (DAN/VAN) systems (Fig. 6e). These correlation patterns are contrasted with those observed at PC_4 , which are dominated by a strong positive correlation in the somatomotor (SMhand) system. Together, these findings demonstrate that not only does the brain's modular architecture vary across individuals in a scale specific manner, but that associations of this variability with behavior also vary in a scale-specific manner. This observation motivates the re-exploration of previously analyzed data, wherein multi-scale associations might have been overlooked, and should spur the collection and analysis of datasets designed to tease apart scale-specific effects.

Finally, we note that, here, we deliberately resist detailed interpretation of reported brain-behavior correlations. Further, we report unthresholded correlation maps, retaining all correlations irrespective of their statistical significance. In general, valid statistical testing is essential in any study attempting to link neuroimaging and behavioral data (and we show the results of statistical tests in Fig. S6). However, we note that in the present study, our aim was not to identify particular brain areas whose modular structures were robustly and significantly associated with behavior, but to demonstrate the level of subjectivity that enters into module detection and how it propagates to later analyses where those modules are linked to behavior.

Detecting multi-layer, multi-subject community structure in the Midnight Scan Club dataset

In the previous sections we used multi-layer modularity maximization to detect and characterize patterns of inter-subject community variability in multi-subject cohorts. While network organization certainly varies across individuals, networks also vary within an individual over successive scans separated by hours, days, or weeks [25, 31–33]. Within-subject variability along different dimensions of network organization has proven useful for explaining variation in task state [34], level of attention [35], and affective state [23].

Here, we repeated our previous analyses using the recently published “Midnight Scan Club” (MSC) dataset in which ten participants underwent repeated fMRI scans (10 times per subject). Using these data we aimed to compare patterns of inter- and intra-subject variability in the modular organization of functional brain networks. To accomplish this aim, we performed two rounds of multi-layer modularity maximization. For each subject, we first generated an average connectivity matrix by aggregating usable fMRI BOLD time series data from across all resting scan sessions, and then we calculated the inter-areal correlation matrix. This procedure resulted in ten connectivity matrices (one per subject) that were treated as the layers of a multi-layer network model. As part of our second analysis, we focused on individual scan sessions from the six subjects (MSC01, MSC02, MSC03, MSC05, MSC06, and MSC07) with at least 300 low-motion frames in each of their ten scan sessions. We then estimated session-specific connectivity matrices for each subject and constructed their respective multi-layer network model whose t -th layer represented a subject's connectivity pattern on the t -th scan session. As with the HCP dataset, we used our multi-scale sampling procedure to sample values for the structural and inter-layer coupling parameters detect communities at those points in parameter space. As before, we calculated normalized entropy measures to characterize patterns of variability across layers (either subjects or scan sessions) and to decompose these patterns into principal component scores

and coefficients using SVD. By comparing the PCs generated from the inter-subject analysis with those generated by the intra-subject analysis, we could effectively identify differences in patterns of inter- *versus* intra-subject community variability.

To do so, we first assessed patterns of inter-subject variability by examining the multi-layer network where each layer corresponded to the session-averaged connectivity matrix for a different subject. Interestingly, despite the fact that these data were independently acquired and processed using a different pipeline, we found similar patterns of inter-subject variability in the MSC dataset as in the HCP dataset. We focus, in particular, on the first three PCs, which we denote MSC PC_1 , MSC PC_2 , and MSC PC_3 . As with HCP PC_1 , the brain areas in MSC PC_1 associated with the lowest inter-subject variability were concentrated in the DMN (Fig. 7a). This similarity is further illustrated by computing the system-averaged correlation of MSC PC_1 scores and HCP PC_1 scores; we observed that the two variables exhibit a strong positive correlation ($r = 0.79$, $p < 0.05$; Fig. 7d). We find analogous relations by pairing MSC PC_2 with HCP PC_2 (Fig. 7b,e), and by pairing MSC PC_3 with HCP PC_4 (Fig. 7c,f). Note that we obtain similar results using Spearman rank correlations: $\rho_{HCPPC1, MSCPC1} = 0.89$, $\rho_{HCPPC2, MSCPC2} = 0.47$, and $\rho_{HCPPC4, MSCPC3} = 0.86$. Also similar to the HCP dataset, we observed a clear dependence of the PCs on the γ parameter. In Fig. S9 we show this dependence by projecting the PCs back into the $\{\gamma, \omega\}$ parameter space.

Next, we performed a series of analyses to characterize *intra*-subject variability in community structure. As noted earlier, this process entailed constructing multi-layer networks for the six MSC subjects with greater than 300 frames of low-motion data in each of their ten scan sessions. In these multi-layer networks, layers corresponded to functional connectivity estimated for single sessions. We aggregated these patterns of intra-subject variability with the previously obtained patterns of inter-subject variability and performed a joint PCA, the results of which we visualized in two-dimensions in a low-dimensional space defined using the first two principal components. If the patterns of community variability observed within subjects were similar to those observed between subjects, then we would expect these projected intra- and inter-subject variability patterns to overlap in this low-dimensional space. However, we found that this was not the case. In Fig. 7g and Fig. 7h, we show two-dimensional histograms for patterns of inter-subject and intra-subject community variability, respectively. Though similar, we find that when we perform an element-wise subtraction of the distributions, there remain distinct portions of parameter space occupied by one or the other modality of variability, suggesting that inter- and intra-subject variability patterns are distinct from one another. One possible explanation is that the day-to-day variability in subject-level communities is driven by personalized factors such as sleep duration, sleep quality, affective state, or level of arousal [23, 35], whereas inter-subject community variability may be driven by temporally stable anatomical factors.

These findings consider the modes of inter-subject variability reported in the HCP data and reproduce them in an independently acquired dataset. Our results also build upon recent work characterizing intra-subject variability in functional network organization [25, 34], confirming that the patterns by which community structure varies within individuals is largely distinct from inter-subject patterns of variation. These observations suggest that the

detected modes are not easily dismissed as artifacts of a particular acquisition strategy, scanner, or processing pipeline. Rather, these results suggest that communities reconfigure across individuals and *datasets* in a robust, general, and multi-scale manner [2], motivating the exploration of techniques designed to detect communities at different organizational levels.

DISCUSSION

In this paper we adapted the popular multi-layer modularity maximization framework so that each layer represented the functional connectivity network of a single subject or scan session. This conceptual alteration facilitated unambiguous comparisons of community structure across subjects and time, permitted the straightforward calculation of consensus communities, and helped us to localize patterns of inter- and intra-subject variability. These advances made it possible for us to detect robust “modes” or (brain patterns) of inter-subject variation in community structure, which we showed were distinct from intra-subject variability. In summary, our findings offer new insight into the multi-scale community structure of functional brain networks and how that structure varies across different putative cognitive systems. The methodological advances presented here enable future studies to investigate variation in community structure across subjects from different clinical cohorts, behavioral states, and (as has already been investigated) time points.

Advantages over current methodology

Here we propose an extension of multi-layer modularity maximization for studying how community structure varies across subjects. Our approach, which has been suggested before but never realized [36, 37], offers distinct advantages over existing methods. First, because community assignments are determined simultaneously for all subjects and because community labels are preserved across layers, we avoid the use of heuristics for mapping community assignments from one subject to another. This facilitates straightforward comparisons of community structure across individuals and allows us to easily obtain consensus communities [38]. Whereas recent work has focused on generating subject-specific community or system assignments by matching to predefined systems [15, 33] using statistical models based on shared covariance structure [39], our approach is grounded in methodology from network science, beginning with the representation of FC as a fully-weighted and signed graph, and ending with the application of modularity maximization to detect community structure [3, 13, 16].

Our use of the multi-layer modularity maximization approach for community detection is not without precedent. Past studies have leveraged this same technique to uncover the evolving community structure in time-varying FC [21–23, 40–43], across tasks [12], and motor behaviors [44]. In recasting single-subject FC matrices as individual layers and moving the problem of community detection from the level of single subjects to that of the cohort, we highlight the flexibility and generalizability of the multi-layer modularity maximization framework. It is not difficult to imagine further extensions of these methods in which network layers correspond to different connection modalities, such as structural connectivity or gene coexpression, for example [45, 46].

Multi-scale community structure and modes of variability

In this study we characterized how communities varied across individuals and time. We found that patterns of variability could be decomposed into a series of “modes,” stereotypical and orthogonal patterns of community variability whose expression was correlated with different numbers and sizes of communities. This observation is in line with the hypothesis that brain networks, both structural and functional, exhibit community structure that spans multiple organizational scales, ranging from coarse divisions into a few large communities (e.g. the division of the resting brain into task-positive and task-negative networks [47]) to much finer divisions reflecting increased functional specialization [48]. The multi-scale and hierarchical structure of communities has a strong theoretical basis: this type of organization facilitates separation of dynamic timescales [49], efficient spatial embedding [50], evolutionary adaptability [51], and robustness to perturbations [52]. There is also considerable empirical evidence demonstrating that brain networks exhibit multi-scale modular structure [53–55] that reconfigures adaptively to support goal-directed cognitive processes [43] and during passive viewing of naturalistic movies [56].

Here, we interpret these modes of variation as evidence of multi-scale topology. An alternative interpretation is that variation in community size and number reflect distinct dynamic regimes. Due to the mechanics of modularity maximization [16] and the use of Pearson correlations to define connection weights [57], the smallest communities reflect collections of nodes that are extremely strongly connected to one another across time (by definition). Large communities, on the other hand, are more diffuse and composed of areas whose activity is correlated with one another, albeit more weakly [58]. Whereas the strongest correlations are also least variable across time [59–61], weaker time-averaged correlations can occur for multiple reasons, e.g. transient periods of strong coupling interspersed with periods of decoupling. Thus, the variation in community structure across topological scales is inextricably linked to variation in network structure across time [62]. Future work could be fruitfully directed towards an investigation into the relationship of time and topology in greater detail.

Implications for analysis of inter-subject community variability

The observation that communities vary across subjects along distinct modes has important implications for future work. Generally, these types of studies calculate community structure at a single scale. That is, single-layer modularity maximization is performed with the structural resolution parameter set uniformly across subjects to a particular value (although this is not always the case [63, 64]). Then, variability in the detected communities is associated with some clinical score or behavioral measure. While this approach has oftentimes proven fruitful, our findings suggest that it is also limiting. Studies that focus exclusively on inter-subject variability at a single scale will fail to characterize variability in communities at other scales. Here, we demonstrated that patterns of brain-behavior correlations depend on the scale at which communities are detected. It follows, then, that any single-scale analysis of the association between community structure and clinical or behavioral measures will result in only one correlation pattern, failing to fully characterize brain-behavior associations at other scales.

Notably, our observations do not detract from these past efforts. Rather, they suggest the possibility that past studies have only scratched the surface in terms of characterizing inter-subject variability in community structure, and that there likely exists a wealth of unexplored brain-behavior associations. This notion dovetails nicely with the missions of recent data-collection and data-sharing projects in the MRI community that have made large-scale datasets and statistical maps freely available to any researcher [65–67]. These datasets, some of which have already been studied through the lens of community structure [68, 69], in principle could be reanalyzed in future studies with an increased focus on characterizing multi-scale patterns of variability.

Our findings suggest that it is important to consider a range of topological scales in any analysis of brain network community structure. From a practical standpoint, however, this can prove challenging. Repeatedly optimizing modularity while varying γ and ω incurs a massive computational burden, especially as the number of nodes in a network or subjects in a cohort increases [13]. In turn, this computational burden limits the utility of our approach. Fortunately, under some contexts this burden can be lessened. Our findings suggest that much of the inter-subject variation in community structure occurs as a function of γ and that while ω may change the baseline magnitude of intersubject variability, the overall pattern, i.e. the brain regions that vary in community assignments across subjects, tends to stay the same. This observation suggests that it may be more important to explore variation as a function of γ rather than variation as a function of ω . Additionally, new software packages such as CHAMP (<https://github.com/wweir827/champ>) [70] show promise in reducing the size of the parameter space to be explored. In summary, although the approach used here has some computational limits, these limits can be partially circumvented, making multi-subject modularity maximization a viable tool for studying inter-subject variation in community structure even in large cohorts.

Here we use the multi-scale, multi-subject modularity maximization approach to detect different patterns of community variability and to show that those patterns, which are sensitive to community detection parameters, correspond to different associations of community structure with measures of behavior. In doing so we focused on associations at the level of brain regions. In general, however, there are other possible strategies for linking community structure to behavioral measures. Another possible approach involves assessing the relationship between inter-subject similarity matrices – one that measures the similarity of subjects' partitions to one another and another that measures the similarity of their behavioral profiles. We could then ask whether the elements of these matrices are correlated with one another, which would suggest that the geometry of behavior is represented in subjects' community structures (analogous, in some ways, to the popular representational similarity analysis framework [71]). We note that the behavioral similarity matrix could be replaced by matrices that encode inter-subject similarity along other dimensions, including genes [72, 73], structural connectivity [74], white-matter microstructure [75], or morphometric data [76]. We also note that here we did not pursue this alternative strategy, but that it could be explored further in future work.

Finally, while our findings show that brain-behavior correlations are sensitive to user-defined parameters, we intentionally do not interpret the correlation patterns in detail. Rather, our

aim was simply to show that these patterns can be manipulated to some extent by a user's decision to select particular combinations of community detection parameters. In future work, more focused analysis should investigate the behavioral and cognitive consequences of the findings reported here.

Limitations

Here, we extended the modularity maximization approach to be compatible with multi-subject cohorts. Though this approach has clear advantages, it also suffers from some limitations. First, as with single-layer and other multi-layer formulations of modularity maximization, the composition of detected communities depends on the free parameters, γ and ω . In our application we aimed to explore the space defined by these parameters. In fact, our findings suggest that this exploration may be necessary, as brain-behavior correlations vary across parameter space. Nonetheless, it may be advantageous in some applications to focus on a particular region of parameter space. In general, choosing the "correct" values of γ and ω is difficult, although many heuristics exist. This includes, for example, selecting the parameter values that result in consistently similar partitions [69] or identifying partitions that maximally differ from an appropriate null model [77, 78]. Finally, it is also worth noting that the modularity maximization framework could be extended much further than we did here. For instance, we fixed the values of γ and ω to be uniform across all layers. It would be interesting and potentially informative to devise heuristics for determining the values of these parameters in a meaningful and subject-specific way, allowing for finer control over the character of detected communities.

Though modularity maximization facilitates our multi-subject analysis, it also serves to limit the scope of our findings. Modularity maximization implicitly assumes that a network's community structure is uniformly assortative. This means that the communities detected using this method will be internally dense and externally sparse. While assortative communities are well-suited for segregated information processing, networks can exhibit more general classes of community structure, including core-periphery and disassortative organization [3, 79]. These communities emphasize integrative information processing and cross-community interactions. Modularity maximization, however, will fail to detect communities of this type. Other methods, such as stochastic blockmodels [80], can detect more general classes of communities. Future work could build on recent applications of blockmodeling to brain network data [81, 82] while taking advantage of multi-layer formulations to study multi-subject cohorts [83–85]. In addition, the modularity maximization framework is subject to so-called resolution limits [86, 87] that, for a given set of parameters, $\{\gamma, \omega\}$, render it incapable of resolving communities below some characteristic size. This limitation, in addition to the aforementioned assumptions about partition composition, motivates the continued exploration of alternative community detection methods.

Another potential limitation concerns the procedure for associating variability at the community level with cognitive measures. Here, we assessed that relationship using community variability measures derived from many distinct partitions. Although this approach enabled us to uncover brainwide patterns of correlations, it comes with the

statistical cost of multiple comparisons. A potential strategy for mitigating this issue is to generate for each mode of variability (PC) a single representative consensus partition [38]. Rather than computing and averaging multiple correlations, generating a consensus partition would allow for the computation of one correlation pattern. Although appealing, the consensus partition strategy has some downsides. In particular, the effectiveness of consensus partitions is limited when they are not representative of the partitions they are intended to represent. In the present study, the principal components identify partitions with correlated patterns of intersubject variability but whose baseline or mean intersubject variability may vary considerably. This feature of the data exists because each principal component is restricted to a narrow range of γ values, ensuring similar numbers of partitions within a given layer, but spans orders of magnitude in terms of ω . Future work could focus on more principled statistical frameworks for dealing with the large numbers of partitions generated by the multi-layer modularity maximization framework.

An additional limitation concerns the manner in which brain network nodes (parcels) are defined. Here, we imposed an identical parcellation across all individuals, which implicitly assumes that brain areas' locations are consistent across subjects. Recent studies, however, have shown that areas vary in their locations across individuals [88]. Consequently, it remains a possibility that inter-subject community variability could be explained by systematic differences in the locations of brain areas. Here, we use group-defined parcels as this approach remains the field standard. Future work should investigate these issues more directly, for example, by exploring the effect of subject-specific parcellations on inter-subject community variability [25, 31, 33, 34].

Another limitation is related to the use of the *a priori* system labels defined by [26]. Among these systems is a group labeled “None,” which collects brain regions with no clear modular assignment, characterized by poor signal to noise, and located in areas of the brain susceptible to artifact and dropout. Here, because we use a data-driven approach to assign nodes to modules and because we tried to avoid manual interventions, nodes with the “None” label were always grouped into communities, though usually not collectively. Rather, the “None” system was broken up and incorporated into distinct communities across the brain. Here, we did not assess the effect of excluding “None” nodes from our analysis. Future work could investigate their impact in greater detail.

A final consideration concerns the scalability of the multi-layer approach. Here, we studied cohorts that included (at most) 80 subjects and connectivity data from 333 brain areas. The flattened modularity tensor associated with these data had dimensions 26640×26640 (333×80). The Louvain algorithm used to maximize $Q(\gamma, \omega)$ was quite fast, which enabled us to sample community structure from many points in the γ, ω parameter space. However, for much larger matrices, corresponding to greater numbers of subjects or finer parcellations, it becomes more difficult to efficiently sample such large numbers of parameters. In such cases, one may wish to consider principled methods for constructing sparse representations of the full data matrices, or for efficiently clustering large networks [89, 90].

CONCLUSION

Here, we studied inter-subject variability in community structure by extending the well-known modularity maximization framework to multiple subjects. We find, surprisingly, that communities vary across individuals along distinct “modes” and that these modes correspond to different organizational scales, ranging from a few large communities to many small communities. Interestingly, when we compare the variation of community structure with different measures of cognitive performance, we find that the correlation patterns are scale-dependent. This observation demonstrates that analyses that calculate brain-behavior correlations based on community structure detected at a single organizational scale may fail to uncover interesting and behaviorally relevant associations. Finally, using a second dataset, we reproduce our previously observed modes of inter-subject community variability, and also show that communities vary within subjects but along different dimensions.

MATERIALS AND METHODS

Datasets

We analyzed functional connectivity data from two independent datasets processed using different pipelines.

Human Connectome Project—We analyzed data from the Human Connectome Project (HCP), a multi-site consortia that collected extensive MRI, behavioral, and demographic data from a large cohort of subjects (>1000) [24]. As part of the HCP protocol, subjects underwent two separate resting state scans. All functional connectivity data analyzed in this report came from these scans and was part of the HCP S1200 release [24]. Only subjects that completed both resting-state scans were analyzed. We utilized a cortical parcellation that maximizes the similarity of functional connectivity within each parcel ($N = 333$ parcels) [26].

We preprocessed resting-state data using the following pipeline. Our analyses were based on the ICA-FIX resting-state data provided by the Human Connectome Project, which used ICA to remove nuisance and motion signals [91]. We removed the mean global signal and bandpass filtered the time series from 0.009 to 0.08 Hz. To reduce artifacts related to in-scanner head motion, frames with greater than 0.2 millimeters of framewise displacement or a derivative root mean square above 75 were removed [92]. Subjects whose scans resulted in fewer than 50% of the total frames left were not analyzed further; a total of 827 subjects met this criteria for all resting-state scans.

For all scans, the MSMAll registration was used, and the mean time series of vertices on the cortical surface (fsL32K) in each of the $N = 333$ parcels was calculated [26]. We used this particular parcellation as it has high functional connectivity homogeneity within each parcel and the number of nodes is consistent with other analyses. The functional connectivity matrix for each subject was calculated as the pairwise Pearson correlation coefficient (subsequently Fisher z -transformed) between time series of all nodes. Both left-right and right-left phase encoding directions scans were used, and the mean functional connectivity matrix across the four resting-state scans was calculated.

In-scanner head motion is known to induce spurious correlations in resting state FC [92]. To reduce the impact of head motion on detected communities, we focused our analyses onto smaller, more exclusive subsets of subjects. Specifically, we divided the full HCP dataset into smaller discovery and validation datasets comprising 80 subjects each. We chose this number of subjects to ensure that dataset size was comparable to what is reported in the typical fMRI study. The test cohort included the 1st, 3rd, 5th ...and 159th subjects, in ascending order of mean framewise displacement . The validation cohort was defined as the 2nd, 4th, 6th, ...160th subjects, ordered according to the same criterion. Note, because this procedure generated datasets with low head motion, it is possible that related subjects appear together in the same cohort.

We also analyzed HCP behavioral data. For the Working Memory task, our measure of task performance was given by the mean accuracy across all conditions (WM_Task_Acc). For the Relational task, our measure of task performance was given by the mean accuracy across all conditions (Relational_Task_Acc). For the Language task, our measure of task performance was given by the highest level reached in either the Language or Math conditions (maximum of Language_Task_Story_Avg_Difficulty_Level and Language_Task_Math_Avg_Difficulty_Level). For the Social task, our measure of task performance was given by the mean across the random and theory of mind conditions (mean of Social_Task_TOM_Perc_TOM and Social_Task_Random_Perc_Random).

Midnight Scan Club—Data were collected from ten healthy, right-handed, young adult subjects (5 females; age: 24–34). One of the subjects is an author (NUFD), and the remaining subjects were recruited from the Washington University community. Informed consent was obtained from all participants. The study was approved by the Washington University School of Medicine Human Studies Committee and Institutional Review Board. This dataset was previously reported in [15, 34] and is publicly available at <https://openneuro.org/datasets/ds000224/versions/00002>. Imaging for each subject was performed on a Siemens TRIO 3T MRI scanner over the course of 12 sessions conducted on separate days, each beginning at midnight. In total, four T1-weighted images, four T2-weighted images, and 5 hours of resting-state BOLD fMRI were collected from each subject. For further details regarding data acquisition parameters, see [15].

MRI data were preprocessed and sampled to the surface as described in [15] and with shared code available at <https://github.com/MidnightScanClub>. The steps are summarized briefly below. The high-resolution structural MRI data were averaged together, and the average T1 images were used to generate hand-edited cortical surfaces using Freesurfer [93]. The resulting surfaces were registered into fs_LR_32k surface space as described in [91]. Separately, an average native T1-to-Talairach [94] volumetric atlas transform was calculated. That transform was applied to the fs_LR_32k surfaces to put them into Talairach volumetric space.

All fMRI data first underwent pre-processing (in the volume) to correct for artifacts and align data, including slice-timing correction, frame-to-frame alignment to correct for motion, and intensity normalization to mode 1000. Functional data were then registered to the T2 image, which was registered to the high-resolution T1 anatomical image, which in turn had

been previously registered to the template space. Finally, functional data underwent distortion correction [15]. Registration, atlas transformation, resampling to 3 mm isotropic resolution, and distortion correction were all combined and applied in a single transformation step [95]. Subsequent steps were all completed on the atlas transformed and resampled data.

Processing steps specific to functional connectivity were undertaken to reduce the influence of artifacts on functional network data. These steps are described in detail in [96], and include (1) demeaning and de-trending of the data, (2) nuisance regression of signals from white matter, cerebrospinal fluid, and the global signal, (3) removal of high motion frames (with framewise displacement (FD) > 0.2 mm; see [15]) and their interpolation using power-spectral matched data, and (4) bandpass filtering (0.009 Hz to 0.08 Hz). After this volumetric pre-processing, functional data were sampled to the cortical surface and combined with volumetric subcortical and cerebellar data into the CIFTI data format using the Connectome Workbench [97]. Finally, data were smoothed (Gaussian kernel, $\sigma = 2.55$ mm) with 2-D geodesic smoothing on the surface and 3-D Euclidean smoothing for subcortical volumetric data.

Modularity maximization

A critical step in any modularity analysis is to determine the community assignments of brain areas. Because real-world networks are too complex to identify communities from visual inspection alone, their nodes' community assignments must be determined algorithmically through a process known as community detection [13]. There are many different ways to define a network's communities and equally many algorithms for detecting them. Of these, modularity maximization is likely the most popular [16].

Intuitively, modularity maximization operates according to a simple principle: communities correspond to groups of nodes that are more strongly connected to one another than would be expected by chance alone. The detection of these communities, in practice, is accomplished by optimizing a modularity objective function:

$$Q = \sum_{ij} B_{ij} \delta(\sigma_i \sigma_j). \quad (1)$$

In this expression, $B_{ij} = W_{ij} - P_{ij}$, where W_{ij} and P_{ij} are the observed and expected weight of the connection between nodes i and j . The full matrix, $B = \{B_{ij}\}$, is referred to as the modularity matrix whose elements encode the difference between observed and expected connection weights. The variable $\sigma_i \in \{1, \dots, K\}$ indicates to which of the K communities node i is assigned. The Kronecker delta function, $\delta(x, y)$, takes on a value of 1 when $x = y$ and 0 when $x \neq y$. Finally, Q is the modularity quality function to be maximized. Effectively, Q is the sum over within-community elements of the modularity matrix. In general, larger values of Q are taken to indicate higher quality partitions and correspond to internally dense and externally sparse communities.

The modularity function suffers from a "resolution limit," meaning that communities below a characteristic scale are undetectable [87]. To circumvent this issue, some applications

introduce a tunable structural resolution parameter, γ , which scales the relative importance of the null connectivity model:

$$Q(\gamma) = \sum_{ij} [A_{ij} - \gamma P_{ij}] \delta(\sigma_i \sigma_j). \quad (2)$$

The effect of this parameterization is that optimizing $Q(\gamma)$ for large values of γ results in the detection of many small communities, whereas optimizing $Q(\gamma)$ for small values of γ results in a few large communities.

The modularity function has been expanded further so that it is compatible with multi-layer networks composed of distinct layers corresponding to different connection modalities (e.g., structural and functional connectivity) or estimates of network structure at different time points [20]. In the multi-layer analog of the modularity function, nodes are linked to themselves across layers by an inter-layer coupling parameter, ω . This parameter determines the similarity of communities detected across layers, with larger values of ω resulting in increased homogeneity of communities across layers. The multi-layer expression for modularity is given by:

$$Q(\gamma, \omega) = \sum_{ijsr} [(W_{ijs} - \gamma P_{ijs}) \delta_{sr} + \omega \delta(ij)] \delta(\sigma_{is} \sigma_{jr}). \quad (3)$$

The multi-layer, multi-scale modularity, $Q(\gamma, \omega)$, operates on the same principle as the single-layer version. Communities are defined by placing stronger-than-expected connections within communities. The main difference between the single- and multi-layer modularity functions is that, by adding interlayer connections of weight ω , it can be advantageous from the perspective of optimizing $Q(\gamma, \omega)$ to assign nodes in different layers to the same community.

Multi-layer, multi-subject modularity

Our approach builds directly upon the canonical multi-layer modularity framework. It differs in two important ways, the first is conceptual while the second is methodological.

Conceptually, rather than letting layers represent different connection modalities [45, 46], frequency bands [98, 99], or network estimates at different time points [21], we let layers represent functional connectivity matrices corresponding to single subjects. Effectively, this choice allows us to obtain estimates of community assignments simultaneously for all subjects. Importantly, this approach also preserves community labels across subjects (provided $\omega > 0$), making the process of comparing communities across subjects trivially easy. Methodologically, our approach differs from the general case in that we let $P_{ijs} = 1$ for all i, j, s . This choice is in line with best practices for modularity maximization when the networks correspond to correlation matrices [58, 100, 101]. In our past work we have shown that this decision results in communities with broadly recognizable topographic features [60, 102].

As with the canonical version of multi-layer modularity maximization, the multi-subject version depends upon two free parameters, γ and ω . The value of γ determines the scale of detected communities, as reflected in their number and size. As before, smaller values of γ

result in a few large communities while larger values result in many small communities. Similarly, the value of ω emphasizes the consistency of communities across individuals. Larger values of ω emphasize network features and communities that are common across subjects, whereas small values of ω emphasize the uniqueness of individuals' community structures. The parameter space defined by γ and ω contains partitions of the network into communities ranging from single nodes to the whole network, that are entirely unique to individuals or perfectly consistent across individuals. Note, that while ω influences the consistency of communities across subjects, it is implemented globally and its influence is exerted over all subjects equally.

Sampling multi-scale, multi-consistency community structure—Though many studies fix the values of γ and ω in order to focus on partitions of networks into roughly the same number of communities across subjects, we aimed to explore the full range of partitions, both in terms of community size but also in terms of the consistency of community structure across individuals. In general, however, the relevant ranges of γ and ω are unknown ahead of time. In our application, we are interested in characterizing variability in community structure across individuals, and therefore wish to avoid extreme partitions; that is, combinations of γ and ω that result in:

1. singleton communities (each node is its own community),
2. whole-network partitions (all nodes are assigned to the same community),
3. partitions that are identical across all subjects, or
4. partitions that are maximally dissimilar across all subjects.

To sample community structure, we employed a novel two-stage procedure. The first stage allowed us to bipartition the parameter space defined by γ and ω into two sub-spaces, one where the above-defined criteria were true and another where the above-defined criteria were false. In the second stage we sampled parameter pairs from within this sub-space to generate a distribution of possible partitions that spanned all organizational scales and levels of consistency.

The first stage is initialized with the user defining the boundaries of a rectangular parameter space, forcing γ and ω to fall within $[\gamma_{min}, \gamma_{max}]$ and $[\omega_{min}, \omega_{max}]$, respectively. Here, we set $\gamma_{min} = \min_{ijs} W_{ijs}$, $\gamma_{max} = \max_{ijs} W_{ijs}$, $\omega_{min} = 0$ and $\omega_{max} = 1$. Next, we sampled parameter pairs uniformly and at random from within this sub-space. To identify regions of this sub-space in which the above-defined four criteria hold, we optimized $Q(\gamma, \omega)$ for each pair of parameter samples using a generalization of the so-called Louvain algorithm [103, 104]. This optimization resulted in a multi-layer partition for which we computed the mean number of communities per layer and the mean consistency of communities across layers (see the next section). Intuitively, if the average number of communities was fewer than 2 or greater than N , or the mean consistency of communities was equal to 0 (maximal inconsistency) or 1 (maximal consistency), we considered the corresponding point in parameter space to be “bad,” as it would have failed to satisfy at least one of the four criteria. Parameter values that satisfy these criteria, on the other hand, are considered “good”.

Next, we calculate the local homogeneity of parameter space by calculating the entropy of each sampled point's 25 nearest neighbors. Points with non-zero values of entropy exhibit inhomogeneities and correspond to regions of parameter space where at least one of the four criteria is sometimes not satisfied. We identified all such points and constructed around them the unique non-convex polygon. Intuitively, partitions corresponding to parameter pairs that fall within this polygon are likely to satisfy all four criteria; points located outside this polygon are unlikely to satisfy all four criteria. We repeated this procedure five times (a total of 5000 samples), each time refining our definition of the boundary between “good” and “bad” regions of parameter space.

At this point our procedure moved onto the second stage. In this stage, we sampled 40000 partitions from within the “good” region. Values of γ were sampled uniformly, while values of ω were sampled from an exponential distribution, so that small values of ω are sampled more frequently than large values. For each pair of parameter values we optimized $Q(\gamma, \omega)$, resulting in 40000 multi-layer partitions. These partitions formed the basis for all subsequent analyses.

Consensus communities and community entropy

Optimizing the multi-layer, multi-scale modularity expression returns $\sigma = \{\sigma_{is}\}$, whose elements encode the community assignment of node i in subject s . From this two-dimensional matrix we can calculate several useful statistics that would not have been accessible using the traditional single-layer approach. First, we can obtain a consensus partition, or a group-representative set of communities. Whereas past studies have relied on iterative clustering with an arbitrary threshold to obtain consensus communities [38], the multi-layer ensemble makes their estimation straightforward. The consensus assignment of node i is given as:

$$\sigma_i^c = \text{mode}_s(\sigma_{is}). \quad (4)$$

Next, we can also calculate the variability of community assignments across subjects for a given node. Let $p_i(k)$ equal the fraction of all subjects whose node i is assigned to community k . The variability of this node's assignment can be characterized with the entropy:

$$h_i = - \sum_{k=1}^K p_i(k) \log_2(p_i(k)). \quad (5)$$

We can standardize this entropy to the range [0, 1] by dividing by $\log_2(K)$.

The measure h_i reveals the extent to which a node's community assignment was consistent or variable across a cohort of subjects. To extend this measure to the level of single subjects, we calculated a new measure, h_{ir} . Consider subject r and node i that is assigned to community σ_{ir} . We can quantify the average dissimilarity of its community assignment to that of all other subjects by calculating:

$$h_{ir} = \frac{1}{T-1} \sum_{s \neq r} [1 - \delta(\sigma_{ir}\sigma_{is})] \quad (6)$$

where T is the total number of subjects in the cohort and $\delta(x, y)$ is the Kronecker delta function, which is equal to 1 when $x = y$ and 0 otherwise. Effectively, this summation reveals the dissimilarity of the community assignment of subject r with respect to that of the rest of the cohort. Whereas h_i measures the variability of node i 's assignment across the entire cohort of subjects, h_{ir} measures the variability of that node's assignment from the perspective of subject r and with respect to the rest of the cohort.

Principal component analysis

The entropy vector describes the variability of each brain areas' community assignment across subjects. These vectors, in general, may depend on where in the $\{\gamma, \omega\}$ parameter space those communities were sampled. We wished to assess whether certain patterns of inter-individual community variability appeared more frequently than others and whether these patterns were localized to specific regions of parameter space.

To address this question, we subjected the column-normalized (zero mean, unit variance) matrix of entropy scores to a principal component analysis via singular value decomposition (SVD). Let $H \in [N \times N_{reps}]$ be the matrix of nodes' entropy scores across all repetitions of the Louvain algorithm. In the case of HCP and MSC data with the Gordon atlas, this matrix has dimensions $[333 \times 40000]$. Singular value decomposition factorizes H into left and right singular vectors, $U \in [N \times N]$ and $V \in [N_{reps} \times N]$ and a set of singular values, $\Sigma \in [N \times N_{reps}]$, satisfying the relationship:

$$H = U\Sigma V^T. \quad (7)$$

The principal component scores and coefficients are given by the columns of U , which are by definition orthogonal to one another, and the rows of V , respectively. The squared diagonal elements of Σ give the percent variance accounted for by principal components. We interpret the columns of U as modes of inter-individual community variation, or dominant brain-wide patterns of entropy scores.

In the case of the MSC dataset, we performed SVD on the inter-individual entropy matrix, which was calculated given communities estimated using a multi-layer network whose layers represented session-averaged FC for each of the 10 subjects. We then extracted the first three principal component scores, $U_{inter} = [U_{(:,1)}, U_{(:,2)}, U_{(:,3)}]$, which defined a low-dimensional space of inter-subject community variability.

In addition to calculating inter-individual communities, we also applied modularity maximization to six different multi-layer networks where each layer represented FC for a given subject on a given scan session. Analogously to our calculation of inter-subject community variability, we computed inter-session community variability for each subject, and subsequently calculated each subject's matrix of column normalized entropy scores,

H_{subject} . These, along with the patterns of *inter*-subject variability, we included in a joint PCA and used to generate panels *g*, *h*, and *i* of Fig. 7.

DATA AVAILABILITY STATEMENT

Data were provided in part 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.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGEMENTS

RFB, MB, and DSB would like to acknowledge support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the ISI Foundation, the Paul Allen Foundation, the Army Research Laboratory (W911NF-10-2-0022), the Army Research Office (Bassett-W911NF-14-1-0679, Grafton-W911NF-16-1-0474, DCIST-W911NF-17-2-0181), the Office of Naval Research, the National Institute of Mental Health (2-R01-DC-009209-11, R01-MH112847, R01-MH107235, R21-MH-106799), the National Institute of Child Health and Human Development (1R01HD086888-01), National Institute of Neurological Disorders and Stroke (R01 NS099348), and the National Science Foundation (BCS-1441502, BCS-1430087, NSF PHY-1554488 and BCS-1631550). The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

References

- [1]. Bullmore Ed and Sporns Olaf, "Complex brain networks: graph theoretical analysis of structural and functional systems," *Nature Reviews Neuroscience* 10, 186 (2009). [PubMed: 19190637]
- [2]. Betzel Richard F and Bassett Danielle S, "Multi-scale brain networks," *Neuroimage* 160, 73–83 (2017). [PubMed: 27845257]
- [3]. Newman Mark EJ, "Communities, modules and large-scale structure in networks," *Nature physics* 8, 25 (2012).
- [4]. Sporns Olaf and Betzel Richard F, "Modular brain networks," *Annual review of psychology* 67, 613–640 (2016).
- [5]. Power Jonathan D, Cohen Alexander L, Nelson Steven M, Wig Gagan S, Barnes Kelly Anne, Church Jessica A, Vogel Alecia C, Laumann Timothy O, Miezin Fran M, Schlaggar Bradley L, et al., "Functional network organization of the human brain," *Neuron* 72, 665–678 (2011). [PubMed: 22099467]
- [6]. Crossley Nicolas A, Mechelli Andrea, Vértes Petra E, Winton-Brown Toby T, Patel Ameera X, Ginestet Cedric E, McGuire Philip, and Bullmore Edward T, "Cognitive relevance of the community structure of the human brain functional coactivation network," *Proceedings of the National Academy of Sciences* 110, 11583–11588 (2013).
- [7]. de Haan Willem, van der Flier Wiesje M, Koene T, Smits Lieke L, Scheltens Philip, and Stam Cornelis J, "Disrupted modular brain dynamics reflect cognitive dysfunction in alzheimer's disease," *Neuroimage* 59, 3085–3093 (2012). [PubMed: 22154957]
- [8]. Alexander-Bloch Aaron F, Gogtay Nitin, Meunier David, Birn Rasmus, Clasen Liv, Lalonde Francois, Lenroot Rhoshel, Giedd Jay, and Bullmore Edward T, "Disrupted modularity and local connectivity of brain functional networks in childhood-onset schizophrenia," *Frontiers in systems neuroscience* 4, 147 (2010). [PubMed: 21031030]

- [9]. Gu Shi, Satterthwaite Theodore D, Medaglia John D, Yang Muzhi, Gur Raquel E, Gur Ruben C, and Bassett Danielle S, "Emergence of system roles in normative neurodevelopment," *Proceedings of the National Academy of Sciences* 112, 13681–13686 (2015).
- [10]. Chan Micaela Y, Park Denise C, Savalia Neil K, Petersen Steven E, and Wig Gagan S, "Decreased segregation of brain systems across the healthy adult lifespan," *Proceedings of the National Academy of Sciences* 111, E4997–E5006 (2014).
- [11]. Geerligns Linda, Renken Remco J, Saliassi Emi, Maurits Natasha M, and Lorist Monicque M, "A brainwide study of age-related changes in functional connectivity," *Cerebral Cortex* 25, 1987–1999 (2014). [PubMed: 24532319]
- [12]. Cole Michael W, Bassett Danielle S, Power Jonathan D, Braver Todd S, and Petersen Steven E, "Intrinsic and task-evoked network architectures of the human brain," *Neuron* 83, 238–251 (2014). [PubMed: 24991964]
- [13]. Fortunato Santo, "Community detection in graphs," *Physics reports* 486, 75–174 (2010).
- [14]. Yeo BT Thomas, Krienen Fenna M, Sepulcre Jorge, Sabuncu Mert R, Lashkari Danial, Hollinshead Marisa, Roffman Joshua L, Smoller Jordan W, Zöllei Lilla, Polimeni Jonathan R, et al., "The organization of the human cerebral cortex estimated by intrinsic functional connectivity," *Journal of neurophysiology* 106, 1125–1165 (2011). [PubMed: 21653723]
- [15]. Gordon Evan M, Laumann Timothy O, Adeyemo Babatunde, and Petersen Steven E, "Individual variability of the system-level organization of the human brain," *Cerebral Cortex* 27, 386–399 (2017). [PubMed: 26464473]
- [16]. Newman Mark EJ and Girvan Michelle, "Finding and evaluating community structure in networks," *Physical review E* 69, 026113 (2004).
- [17]. Rosvall Martin and Bergstrom Carl T, "Maps of random walks on complex networks reveal community structure," *Proceedings of the National Academy of Sciences* 105, 1118–1123 (2008).
- [18]. Bellec Pierre, Rosa-Neto Pedro, Lyttelton Oliver C, Benali Habib, and Evans Alan C, "Multi-level bootstrap analysis of stable clusters in resting-state fmri," *Neuroimage* 51, 1126–1139 (2010). [PubMed: 20226257]
- [19]. Schaefer Alexander, Kong Ru, Gordon Evan M, Laumann Timothy O, Zuo Xi-Nian, Holmes Avram J, Eickhoff Simon B, and Yeo BT, "Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri," *Cerebral Cortex*, 1–20 (2017). [PubMed: 28365777]
- [20]. Mucha Peter J, Richardson Thomas, Macon Kevin, Porter Mason A, and Onnela Jukka-Pekka, "Community structure in time-dependent, multiscale, and multiplex networks," *science* 328, 876–878 (2010). [PubMed: 20466926]
- [21]. Bassett Danielle S, Wymbs Nicholas F, Porter Mason A, Mucha Peter J, Carlson Jean M, and Grafton Scott T, "Dynamic reconfiguration of human brain networks during learning," *Proceedings of the National Academy of Sciences* 108, 7641–7646 (2011).
- [22]. Braun Urs, Schäfer Axel, Walter Henrik, Erk Susanne, Romanczuk-Seiferth Nina, Haddad Leila, Schweiger Janina I, Grimm Oliver, Heinz Andreas, Tost Heike, et al., "Dynamic reconfiguration of frontal brain networks during executive cognition in humans," *Proceedings of the National Academy of Sciences* 112, 11678–11683 (2015).
- [23]. Betzel Richard F, Satterthwaite Theodore D, Gold Joshua I, and Bassett Danielle S, "Positive affect, surprise, and fatigue are correlates of network flexibility," *Scientific Reports* 7, 520 (2017). [PubMed: 28364117]
- [24]. Van Essen David C, Smith Stephen M, Barch Deanna M, Behrens Timothy EJ, Yacoub Essa, Ugurbil Kamil, Wu-Minn HCP Consortium, et al., "The wu-minn human connectome project: an overview," *Neuroimage* 80, 62–79 (2013). [PubMed: 23684880]
- [25]. Gordon Evan M, Laumann Timothy O, Gilmore Adrian W, Newbold Dillan J, Greene Deanna J, Berg Jeffrey J, Ortega Mario, Hoyt-Drazen Catherine, Gratton Caterina, Sun Haoxin, et al., "Precision functional mapping of individual human brains," *Neuron* 95, 791–807 (2017). [PubMed: 28757305]
- [26]. Gordon Evan M, Laumann Timothy O, Adeyemo Babatunde, Huckins Jeremy F, Kelley William M, and Petersen Steven E, "Generation and evaluation of a cortical area parcellation from resting-state correlations," *Cerebral cortex* 26, 288–303 (2014). [PubMed: 25316338]

- [27]. Reichardt Jörg and Bornholdt Stefan, “Statistical mechanics of community detection,” *Physical Review E* 74, 016110 (2006).
- [28]. Kheirkhazadeh Masoumeh, Lancichinetti Andrea, and Rosvall Martin, “Efficient community detection of network flows for varying markov times and bipartite networks,” *Physical Review E* 93, 032309 (2016). [PubMed: 27078368]
- [29]. Bertolero Maxwell A, Yeo BT Thomas, and Mark D’Esposito, “The modular and integrative functional architecture of the human brain,” *Proceedings of the National Academy of Sciences* 112, E6798–E6807 (2015).
- [30]. Bertolero MA, Yeo BTT, and M D’Esposito, “The diverse club,” *Nature communications* 8, 1277 (2017).
- [31]. Laumann Timothy O, Gordon Evan M, Adeyemo Babatunde, Snyder Abraham Z, Joo Sung Jun, Chen Mei-Yen, Gilmore Adrian W, McDermott Kathleen B, Nelson Steven M, Dosenbach Nico UF, et al., “Functional system and areal organization of a highly sampled individual human brain,” *Neuron* 87, 657–670 (2015). [PubMed: 26212711]
- [32]. Poldrack Russell A, Laumann Timothy O, Koyejo Oluwasanmi, Gregory Brenda, Hover Ashleigh, Chen Mei-Yen, Gorgolewski Krzysztof J, Luci Jeffrey, Joo Sung Jun, Boyd Ryan L, et al., “Long-term neural and physiological phenotyping of a single human,” *Nature communications* 6, 8885 (2015).
- [33]. Kong Ru, Li Jingwei, Orban Csaba, Mert R Sabuncu Hesheng Liu, Schaefer Alexander, Sun Nanbo, Zuo Xi-Nian, Holmes Avram J, Eickhoff Simon B, et al., “Spatial topography of individual-specific cortical networks predicts human cognition, personality, and emotion,” *Cerebral Cortex* (2018).
- [34]. Gratton Caterina, Laumann Timothy O, Nielsen Ashley N, Greene Deanna J, Gordon Evan M, Gilmore Adrian W, Nelson Steven M, Coalson Rebecca S, Snyder Abraham Z, Schlaggar Bradley L, et al., “Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation,” *Neuron* 98, 439–452 (2018). [PubMed: 29673485]
- [35]. Shine James M, Koyejo Oluwasanmi, and Poldrack Russell A, “Temporal metastates are associated with differential patterns of time-resolved connectivity, network topology, and attention,” *Proceedings of the National Academy of Sciences* 113, 9888–9891 (2016).
- [36]. Khambhati Ankit N, Sizemore Ann E, Betzel Richard F, and Bassett Danielle S, “Modeling and interpreting mesoscale network dynamics,” *Neuroimage* (2017).
- [37]. Vaiana Michael and Muldoon Sarah Feldt, “Multilayer brain networks,” *Journal of Nonlinear Science*, 1–23 (2018).
- [38]. Lancichinetti Andrea and Fortunato Santo, “Consensus clustering in complex networks,” *Scientific reports* 2, 336 (2012). [PubMed: 22468223]
- [39]. Sripada Chandra, Angststadt Mike, Rutherford Saige, Kessler Daniel, Kim Yura, Yee Mike, and Levina Liza, “Fundamental units of inter-individual variation in resting state connectomes,” *bioRxiv*, 326082 (2018).
- [40]. Shine James M, Bissett Patrick G, Bell Peter T, Koyejo Oluwasanmi, Balsters Joshua H, Gorgolewski Krzysztof J, Moodie Craig A, and Poldrack Russell A, “The dynamics of functional brain networks: integrated network states during cognitive task performance,” *Neuron* 92, 544–554 (2016). [PubMed: 27693256]
- [41]. Gerraty Raphael T, Davidow Juliet Y, Foerde Karin, Galvan Adriana, Bassett Danielle S, and Shohamy Daphna, “Dynamic flexibility in striatal-cortical circuits supports reinforcement learning,” *Journal of Neuroscience*, 2084–17 (2018).
- [42]. Vatansever Deniz, Menon David K, Manktelow Anne E, Sahakian Barbara J, and Stamatakis Emmanuel A, “Default mode dynamics for global functional integration,” *Journal of Neuroscience* 35, 15254–15262 (2015). [PubMed: 26586814]
- [43]. Bola Michal and Sabel Bernhard A, “Dynamic reorganization of brain functional networks during cognition,” *Neuroimage* 114, 398–413 (2015). [PubMed: 25828884]
- [44]. Wymbs Nicholas F, Bassett Danielle S, Mucha Peter J, Porter Mason A, and Grafton Scott T, “Differential recruitment of the sensorimotor putamen and frontoparietal cortex during motor chunking in humans,” *Neuron* 74, 936–946 (2012). [PubMed: 22681696]

- [45]. Battiston Federico, Guillon Jeremy, Chavez Mario, Latora Vito, and De Vico Fallani Fabrizio, "Multiplex core-periphery organization of the human connectome," arXiv preprint arXiv:1801.01913 (2017).
- [46]. Bentley Barry, Branicky Robyn, Barnes Christopher L, Lian Chew Yee, Yemini Eviatar, Bullmore Edward T, Vértés Petra E, and Schafer William R, "The multilayer connectome of *Caenorhabditis elegans*," PLoS computational biology 12, e1005283 (2016). [PubMed: 27984591]
- [47]. Golland Yulia, Golland Polina, Bentin Shlomo, and Malach Rafael, "Data-driven clustering reveals a fundamental subdivision of the human cortex into two global systems," Neuropsychologia 46, 540–553 (2008). [PubMed: 18037453]
- [48]. Rosenthal Gideon, Sporns Olaf, and Avidan Galia, "Stimulus dependent dynamic reorganization of the human face processing network," Cerebral Cortex 27, 4823–4834 (2016).
- [49]. Pan Raj Kumar and Sinha Sitabhra, "Modularity produces small-world networks with dynamical timescale separation," EPL (Europhysics Letters) 85, 68006 (2009).
- [50]. Bassett Danielle S, Greenfield Daniel L, Meyer-Lindenberg Andreas, Weinberger Daniel R, Moore Simon W, and Bullmore Edward T, "Efficient physical embedding of topologically complex information processing networks in brains and computer circuits," PLoS computational biology 6, e1000748 (2010). [PubMed: 20421990]
- [51]. Kirschner Marc and Gerhart John, "Evolvability," Proceedings of the National Academy of Sciences 95, 8420–8427 (1998).
- [52]. Simon Herbert A, "The architecture of complexity," in Facets of systems science (Springer, 1991) pp. 457–476.
- [53]. Meunier David, Lambiotte Renaud, and Bullmore Edward T, "Modular and hierarchically modular organization of brain networks," Frontiers in neuroscience 4, 200 (2010). [PubMed: 21151783]
- [54]. Doucet Gaëlle, Naveau Mikaël, Petit Laurent, Delcroix Nicolas, Zago Laure, Crivello Fabrice, Jobard Gael, Tzourio-Mazoyer Nathalie, Mazoyer Bernard, Mellet Emmanuel, et al., "Brain activity at rest: a multiscale hierarchical functional organization," Journal of neurophysiology 105, 2753–2763 (2011). [PubMed: 21430278]
- [55]. Ferrarini Luca, Veer Ilya M, Baerends Evelinda, van Tol Marie-José, Renken Remco J, van der Wee Nic JA, Veltman Dirk J, Aleman Andre, Zitman Frans G, Penninx Brenda WJH, et al., "Hierarchical functional modularity in the resting-state human brain," Human brain mapping 30, 2220–2231 (2009). [PubMed: 18830955]
- [56]. Kim DoHyun, Kay Kendrick, Shulman Gordon L, and Corbetta Maurizio, "A new modular brain organization of the bold signal during natural vision," Cerebral Cortex 28, 3065–3081 (2017).
- [57]. Smith Stephen M, Miller Karla L, Salimi-Khorshidi Gholamreza, Webster Matthew, Beckmann Christian F, Nichols Thomas E, Ramsey Joseph D, and Woolrich Mark W, "Network modelling methods for fmri," Neuroimage 54, 875–891 (2011). [PubMed: 20817103]
- [58]. Traag Vincent A, Van Dooren Paul, and Nesterov Yurii, "Narrow scope for resolution-limit-free community detection," Physical Review E 84, 016114 (2011).
- [59]. Shen Kelly, Matthew Hutchison R, Bezgin Gleb, Everling Stefan, and McIntosh Anthony R, "Network structure shapes spontaneous functional connectivity dynamics," Journal of Neuroscience 35, 5579–5588 (2015). [PubMed: 25855174]
- [60]. Betzel Richard F, Fukushima Makoto, He Ye, Zuo Xi-Nian, and Sporns Olaf, "Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks," NeuroImage 127, 287–297 (2016). [PubMed: 26687667]
- [61]. Thompson William H and Fransson Peter, "The mean-variance relationship reveals two possible strategies for dynamic brain connectivity analysis in fmri," Frontiers in human neuroscience 9, 398 (2015). [PubMed: 26236216]
- [62]. Lurie Daniel, Kessler Daniel, Bassett Danielle, Betzel Richard F, Breakspear Michael, Keilholz Shella, Kucyi Aaron, Liégeois Raphaël, Lindquist Martin A, McIntosh Anthony Randal, et al., "On the nature of resting fmri and time-varying functional connectivity," (2018).
- [63]. Akiki Teddy J and Abdallah Chadi G, "Determining the hierarchical architecture of the human brain using subject-level clustering of functional networks," bioRxiv, 350462 (2018).

- [64]. Akiki Teddy J, Averill Christopher L, Wrocklage Kristen M, Scott J Cobb, Averill Lynnette A, Schweinsburg Brian, Alexander-Bloch Aaron, Martini Brenda, Southwick Steven M, Krystal John H, et al., "Default mode network abnormalities in posttraumatic stress disorder: A novel network-restricted topology approach," *NeuroImage* 176, 489–498 (2018). [PubMed: 29730491]
- [65]. Poldrack Russell A, Barch Deanna M, Mitchell Jason, Wager Tor, Wagner Anthony D, Devlin Joseph T, Cumba Chad, Koyejo Oluwasanmi, and Milham Michael, "Toward open sharing of task-based fmri data: the openfmri project," *Frontiers in neuroinformatics* 7, 12 (2013). [PubMed: 23847528]
- [66]. Gorgolewski Krzysztof J, Varoquaux Gael, Rivera Gabriel, Schwarz Yannick, Ghosh Satrajit S, Maumet Camille, Sochat Vanessa V, Nichols Thomas E, Poldrack Russell A, Poline Jean-Baptiste, et al., "Neurovault. org: a web-based repository for collecting and sharing unthresholded statistical maps of the human brain," *Frontiers in neuroinformatics* 9, 8 (2015). [PubMed: 25914639]
- [67]. Satterthwaite Theodore D, Elliott Mark A, Ruparel Kosha, Loughhead James, Prabhakaran Karthik, Calkins Monica E, Hopson Ryan, Jackson Chad, Keefe Jack, Riley Marisa, et al., "Neuroimaging of the philadelphia neurodevelopmental cohort," *Neuroimage* 86, 544–553 (2014). [PubMed: 23921101]
- [68]. Baum Graham L, Ciric Rastko, Roalf David R, Betzel Richard F, Moore Tyler M, Shinohara Russell T, Kahn Ari E, Vandekar Simon N, Rupert Petra E, Quarmley Megan, et al., "Modular segregation of structural brain networks supports the development of executive function in youth," *Current Biology* 27, 1561–1572 (2017). [PubMed: 28552358]
- [69]. Betzel Richard F, Medaglia John D, Papadopoulos Lia, Baum Graham L, Gur Ruben, Gur Raquel, Roalf David, Satterthwaite Theodore D, and Bassett Danielle S, "The modular organization of human anatomical brain networks: Accounting for the cost of wiring," *Network Neuroscience* 1, 42–68 (2017). [PubMed: 30793069]
- [70]. Weir William, Emmons Scott, Gibson Ryan, Taylor Dane, and Mucha Peter, "Post-processing partitions to identify domains of modularity optimization," *Algorithms* 10, 93 (2017). [PubMed: 29046743]
- [71]. Kriegeskorte Nikolaus, Mur Marieke, and Bandettini Peter A, "Representational similarity analysis connecting the branches of systems neuroscience," *Frontiers in systems neuroscience* 2, 4 (2008). [PubMed: 19104670]
- [72]. Betzel Richard F, Medaglia John D, Kahn Ari E, Soffer Jonathan, Schonhaut Daniel R, and Bassett Danielle S, "Inter-regional ecog correlations predicted by communication dynamics, geometry, and correlated gene expression," *arXiv preprint arXiv:1706.06088* (2017).
- [73]. Hawrylycz Michael J, Lein Ed S, Guillozet-Bongaarts Angela L, Shen Elaine H, Ng Lydia, Miller Jeremy A, Van De Lagemat Louie N, Smith Kimberly A, Ebbert Amanda, Riley Zackery L, et al., "An anatomically comprehensive atlas of the adult human brain transcriptome," *Nature* 489, 391 (2012). [PubMed: 22996553]
- [74]. Hagmann Patric, Cammoun Leila, Gigandet Xavier, Meuli Reto, Honey Christopher J, Van Wedeen J, and Sporns Olaf, "Mapping the structural core of human cerebral cortex," *PLoS biology* 6, e159 (2008). [PubMed: 18597554]
- [75]. Yeatman Jason D, Dougherty Robert F, Myall Nathaniel J, Wandell Brian A, and Feldman Heidi M, "Tract profiles of white matter properties: automating fiber-tract quantification," *PloS one* 7, e49790 (2012). [PubMed: 23166771]
- [76]. He Yong, Chen Zhang J, and Evans Alan C, "Small-world anatomical networks in the human brain revealed by cortical thickness from mri," *Cerebral cortex* 17, 2407–2419 (2007). [PubMed: 17204824]
- [77]. Jeub Lucas GS, Sporns Olaf, and Fortunato Santo, "Multiresolution consensus clustering in networks," *Scientific reports* 8, 3259 (2018). [PubMed: 29459635]
- [78]. Aldecoa Rodrigo and Marín Ignacio, "Deciphering network community structure by surprise," *PloS one* 6, e24195 (2011). [PubMed: 21909420]
- [79]. Peel Leto, Delvenne Jean-Charles, and Lambiotte Renaud, "Multiscale mixing patterns in networks," *Proceedings of the National Academy of Sciences* 115, 4057–4062 (2018).

- [80]. Karrer Brian and Newman Mark EJ, “Stochastic block-models and community structure in networks,” *Physical review E* 83, 016107 (2011).
- [81]. Betzel Richard F, Medaglia John D, and Bassett Danielle S, “Diversity of meso-scale architecture in human and non-human connectomes,” *Nature Communications* 9, 346 (2018).
- [82]. Betzel Richard F, Bertolero Maxwell A, and Bassett Danielle S, “Non-assortative community structure in resting and task-evoked functional brain networks,” *bioRxiv*, 355016 (2018).
- [83]. Stanley Natalie, Shai Saray, Taylor Dane, and Mucha Peter J, “Clustering network layers with the strata multilayer stochastic block model,” *IEEE transactions on network science and engineering* 3, 95–105 (2016). [PubMed: 28435844]
- [84]. Paul Subhadeep, Chen Yuguo, et al., “Consistent community detection in multi-relational data through restricted multi-layer stochastic blockmodel,” *Electronic Journal of Statistics* 10, 3807–3870 (2016).
- [85]. Han Qiuyi, Xu Kevin, and Airoidi Edoardo, “Consistent estimation of dynamic and multi-layer block models,” in *International Conference on Machine Learning* (2015) pp. 1511–1520.
- [86]. Vaiana Michael and Muldoon Sarah, “Resolution limits for detecting community changes in multilayer networks,” *arXiv preprint arXiv:1803.03597* (2018).
- [87]. Fortunato Santo and Barthelemy Marc, “Resolution limit in community detection,” *Proceedings of the National Academy of Sciences* 104, 36–41 (2007).
- [88]. Bijsterbosch Janine Diane, Woolrich Mark W, Glasser Matthew F, Robinson Emma C, Beckmann Christian F, Van Essen David C, Harrison Samuel J, and Smith Stephen M, “The relationship between spatial configuration and functional connectivity of brain regions,” *Elife* 7, e32992 (2018). [PubMed: 29451491]
- [89]. Clauset Aaron, Newman Mark EJ, and Moore Cristopher, “Finding community structure in very large networks,” *Physical review E* 70, 066111 (2004).
- [90]. Capocci Andrea, Servedio Vito DP, Caldarelli Guido, and Colaiori Francesca, “Detecting communities in large networks,” *Physica A: Statistical Mechanics and its Applications* 352, 669–676 (2005).
- [91]. Glasser Matthew F, Sotiropoulos Stamatios N, Wilson J Anthony, Coalson Timothy S, Fischl Bruce, Andersson Jesper L, Xu Junqian, Jbabdi Saad, Webster Matthew, Polimeni Jonathan R, et al., “The minimal preprocessing pipelines for the human connectome project,” *Neuroimage* 80, 105–124 (2013). [PubMed: 23668970]
- [92]. Power Jonathan D, Barnes Kelly A, Snyder Abraham Z, Schlaggar Bradley L, and Petersen Steven E, “Spurious but systematic correlations in functional connectivity mri networks arise from subject motion,” *Neuroimage* 59, 2142–2154 (2012). [PubMed: 22019881]
- [93]. Dale Anders M, Fischl Bruce, and Sereno Martin I, “Cortical surface-based analysis: I. segmentation and surface reconstruction,” *Neuroimage* 9, 179–194 (1999). [PubMed: 9931268]
- [94]. Talairach Jean and Tournoux Pierre, “Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system: an approach to cerebral imaging,” (1988).
- [95]. Smith Stephen M, Jenkinson Mark, Woolrich Mark W, Beckmann Christian F, Behrens Timothy EJ, Johansen-Berg Heidi, Bannister Peter R, De Luca Marilena, Drobnyak Ivana, Flitney David E, et al., “Advances in functional and structural mr image analysis and implementation as fsl,” *Neuroimage* 23, S208–S219 (2004). [PubMed: 15501092]
- [96]. Power Jonathan D, Mitra Anish, Laumann Timothy O, Snyder Abraham Z, Schlaggar Bradley L, and Petersen Steven E, “Methods to detect, characterize, and remove motion artifact in resting state fmri,” *Neuroimage* 84, 320–341 (2014). [PubMed: 23994314]
- [97]. Marcus Daniel, Harwell John, Olsen Timothy, Hodge Michael, Glasser Matthew, Prior Fred, Jenkinson Mark, Laumann Timothy, Curtiss Sandra, and Van Essen David, “Informatics and data mining tools and strategies for the human connectome project,” *Frontiers in neuroinformatics* 5, 4 (2011). [PubMed: 21743807]
- [98]. De Domenico Manlio, Sasai Shuntaro, and Arenas Alex, “Mapping multiplex hubs in human functional brain networks,” *Frontiers in neuroscience* 10, 326 (2016). [PubMed: 27471443]
- [99]. Brookes Matthew J, Tewarie Prejaas K, Hunt Benjamin AE, Robson Sian E, Gascoyne Lauren E, Liddle Elizabeth B, Liddle Peter F, and Morris Peter G, “A multi-layer network approach to meg connectivity analysis,” *Neuroimage* 132, 425–438 (2016). [PubMed: 26908313]

- [100]. Bazzi Marya, Porter Mason A, Williams Stacy, McDonald Mark, Fenn Daniel J, and Howison Sam D, "Community detection in temporal multilayer networks, with an application to correlation networks," *Multiscale Modeling & Simulation* 14, 1–41 (2016).
- [101]. MacMahon Mel and Garlaschelli Diego, "Community detection for correlation matrices," arXiv preprint arXiv:1311.1924 (2013).
- [102]. Betzel Richard F and Bassett Danielle S, "Specificity and robustness of long-distance connections in weighted, interareal connectomes," *Proceedings of the National Academy of Sciences*, 201720186 (2018).
- [103]. Jutla Inderjit S, Jeub Lucas GS, and Mucha Peter J, "A generalized louvain method for community detection implemented in matlab," URL <http://netwiki.amath.unc.edu/GenLouvain> (2011).
- [104]. Blondel Vincent D, Guillaume Jean-Loup, Lambiotte Renaud, and Lefebvre Etienne, "Fast unfolding of communities in large networks," *Journal of statistical mechanics: theory and experiment* 2008, P10008 (2008).

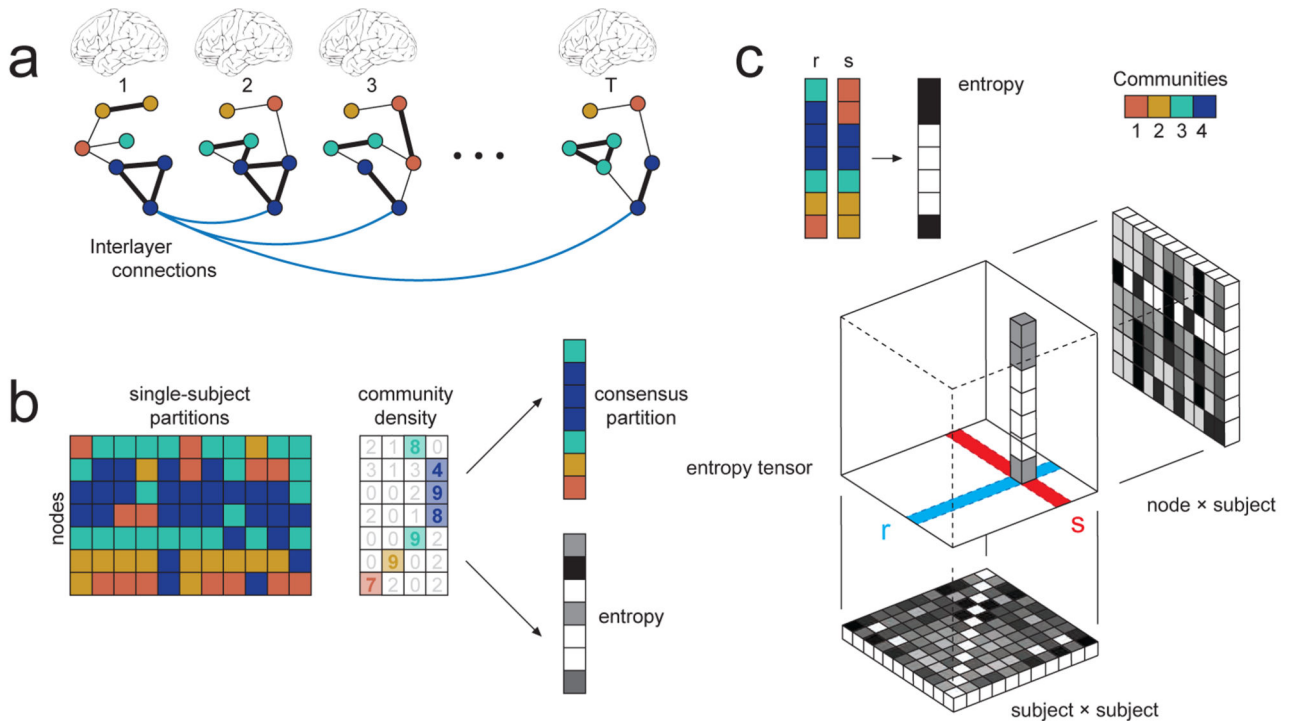


FIG. 1. Multi-subject modularity, communities, and areal entropy.

(a) Single-subject networks are represented as layers in a multi-layer network ensemble. Each node is linked to itself across layers, here illustrated by interlayer connections. Note that community labels are indicated by node color. (b) Maximizing a multi-layer modularity function returns a set of single-subject partitions. Importantly, community labels are preserved across layers; thus, if the label C_1 appears in layers r and s , we assume that the same community has recurred. This property allows us to make several useful measurements. We can calculate, for each node, the mode of its community assignment across subjects to generate a consensus partition. We can also calculate the entropy of each node's community assignments, which measures the variability of communities across subjects. (c) The preservation of community labels also allows for a direct comparison of any one subject to any other subject. Given partitions of subjects (or layers), denoted here with variables r and s , we can generate a bit vector whose values are $\{0, 1\}$ depending on whether a given node has the same/different community assignment. Doing so for all pairs of subjects generates a three-dimensional *entropy tensor*. When averaged over nodes, this tensor generates a $T \times T$ matrix whose elements indicate, in total, the number of non-identical community assignments between pairs of subjects. When averaged over either of its other dimensions, the result is an $N \times T$ matrix, whose elements indicate, in total, the similarity of a node's community assignment within a given subject to that of the remaining $T - 1$ subjects.

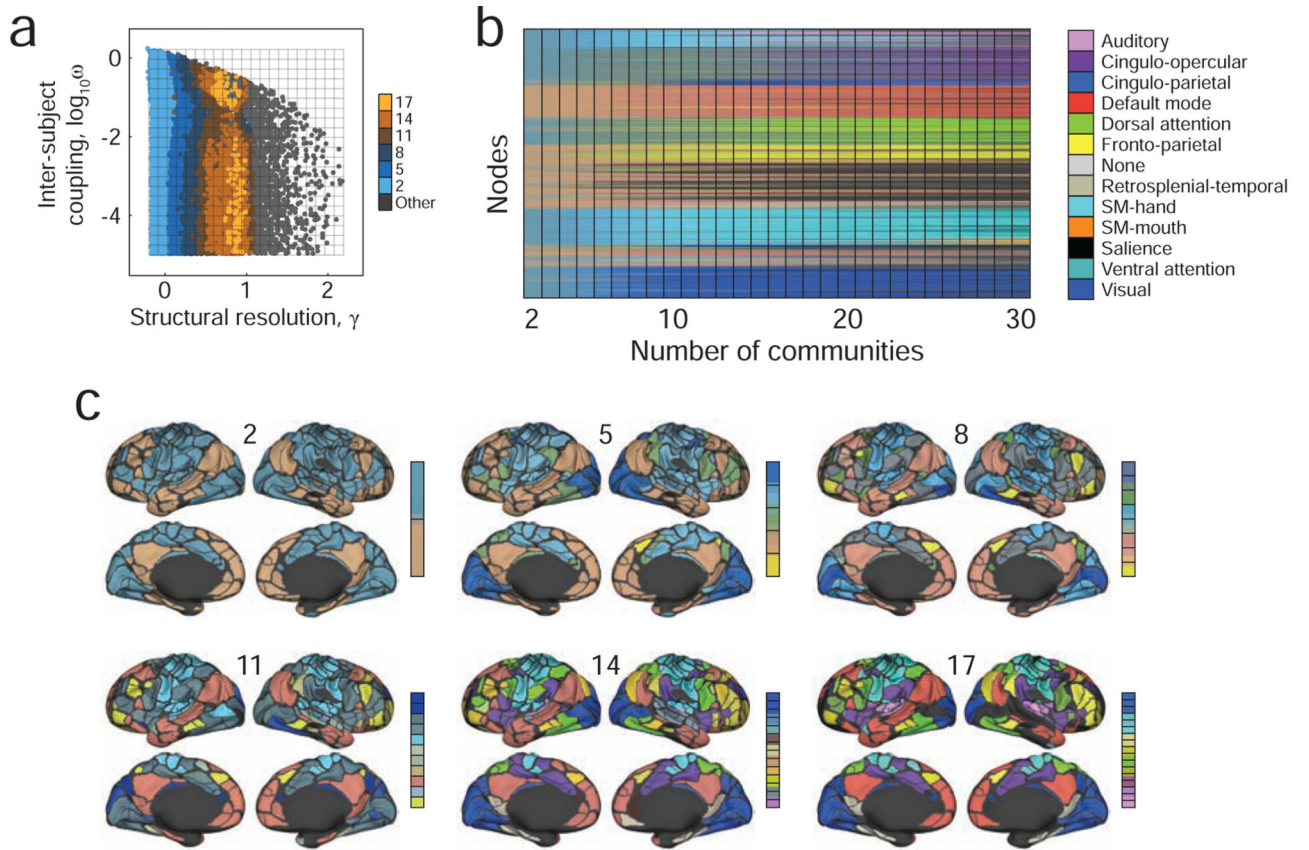


FIG. 2. Examples of detected community structure.

(a) The composition of detected communities depends on the structural resolution parameter, γ , and on the inter-subject coupling parameter, ω . To generate a sample of possible partitions, we chose random combinations of γ and ω and estimate consensus community structure at those points. We show, here, the locations in parameter space where the resulting consensus partitions contained 2, 5, 8, 11, 14, and 17 non-singleton communities per subject. (b) We ordered all consensus partitions in ascending order according to their number of non-singleton communities. Each community was colored by the weighted average of its constituent brain areas' cognitive systems. For example, a community comprised of exclusively DMN brain areas would be assigned the DMN color (red, in this case), whereas a system composed of an equal number of DMN and visual brain areas would have a purple color (the average of the DMN's red and the visual system's blue). (c) We also show example consensus partitions as we vary the number of non-singleton communities to 2, 5, 8, 11, 14, and 17.

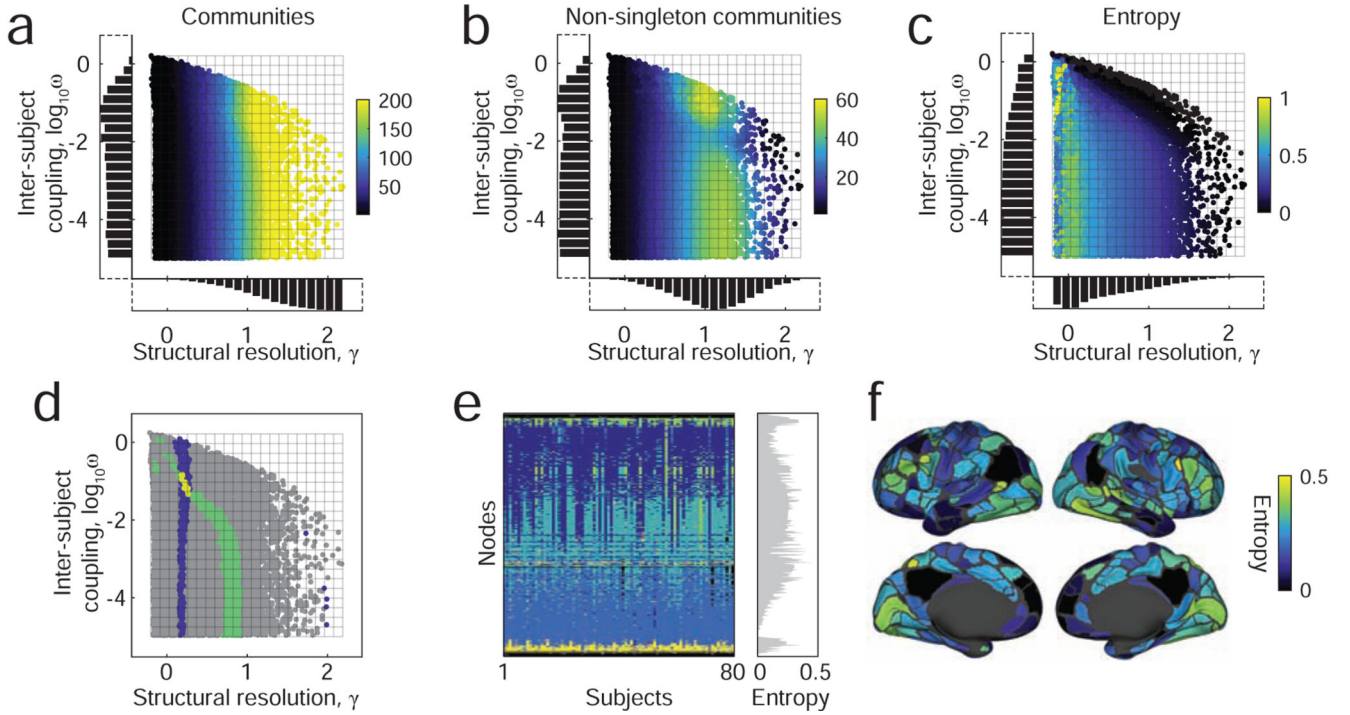


FIG. 3. Multi-scale analysis strategy and schematic.

Points are sampled in a two-dimensional constrained parameter space. The structural resolution parameter, γ , determines the number and size of communities while the inter-layer coupling parameter, ω , tunes the consistency of communities across individuals. Here, we summarize the statistics of communities detected using this sampling approach applied to the HCP333 dataset. (a) The number of communities per layer. (b) The number of communities per layer after excluding singleton communities, which are communities composed of a single node. (c) The mean inter-subject entropy (variability). (d) We can query particular subsets of partitions based on the number of communities, their average entropy, or other statistics, allowing us not only to probe different organizational scales of the network, but also to accommodate varying degrees of heterogeneity across subjects. (e) An example of the detected partitions and their consistency across $T = 80$ subjects. (f) The variability (inter-subject entropy) of community assignments across individuals. Brighter (yellow) coloring indicates greater levels of variability.

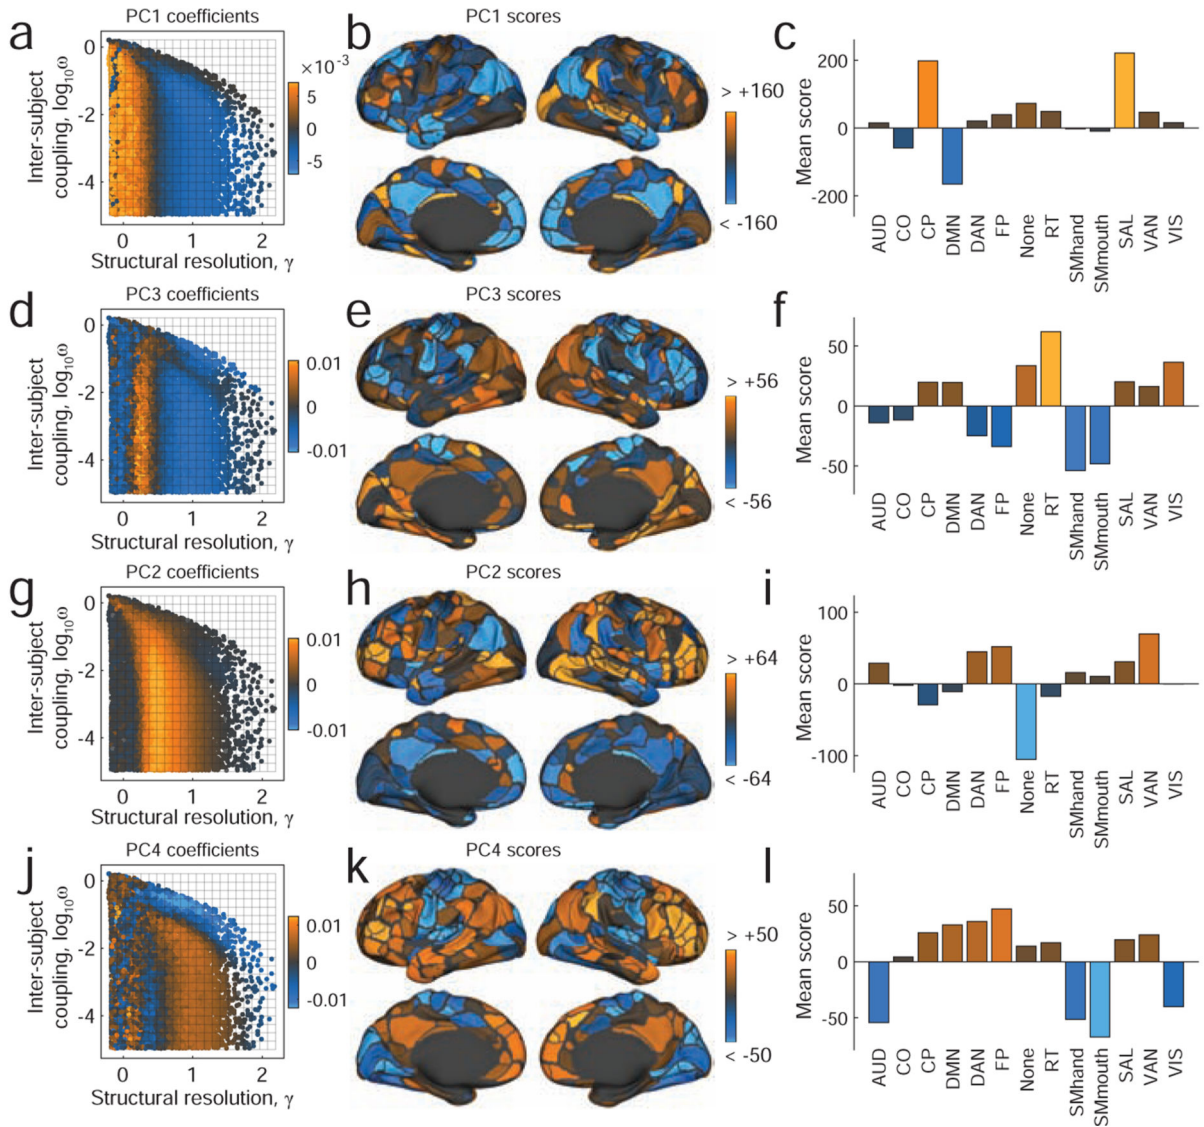


FIG. 4. Modes of inter-subject community variability.

Principal component coefficients and scores for the first four components. (a) PC coefficients for the first component projected into $\{\gamma, \omega\}$ parameter space. (b) PC scores for the same component projected onto the cortical surface. The community assignments of bright orange brain areas are highly variable across subjects at orange points in the parameter space. Conversely, the community assignments of blue brain areas are highly consistent across subjects at those same points. (c) Areal values of principal component scores averaged across thirteen previously-defined cognitive/functional systems [26]. This panel helps shift focus away area-level community variability and onto system level patterns of variation. The remaining panels show corresponding plots for PC_2 , PC_3 , and PC_4 . Note: we present principal components in the order in which they are expressed along the γ axis. This choice results in the following ordering: PC_1 , PC_3 , PC_2 , and PC_4 .

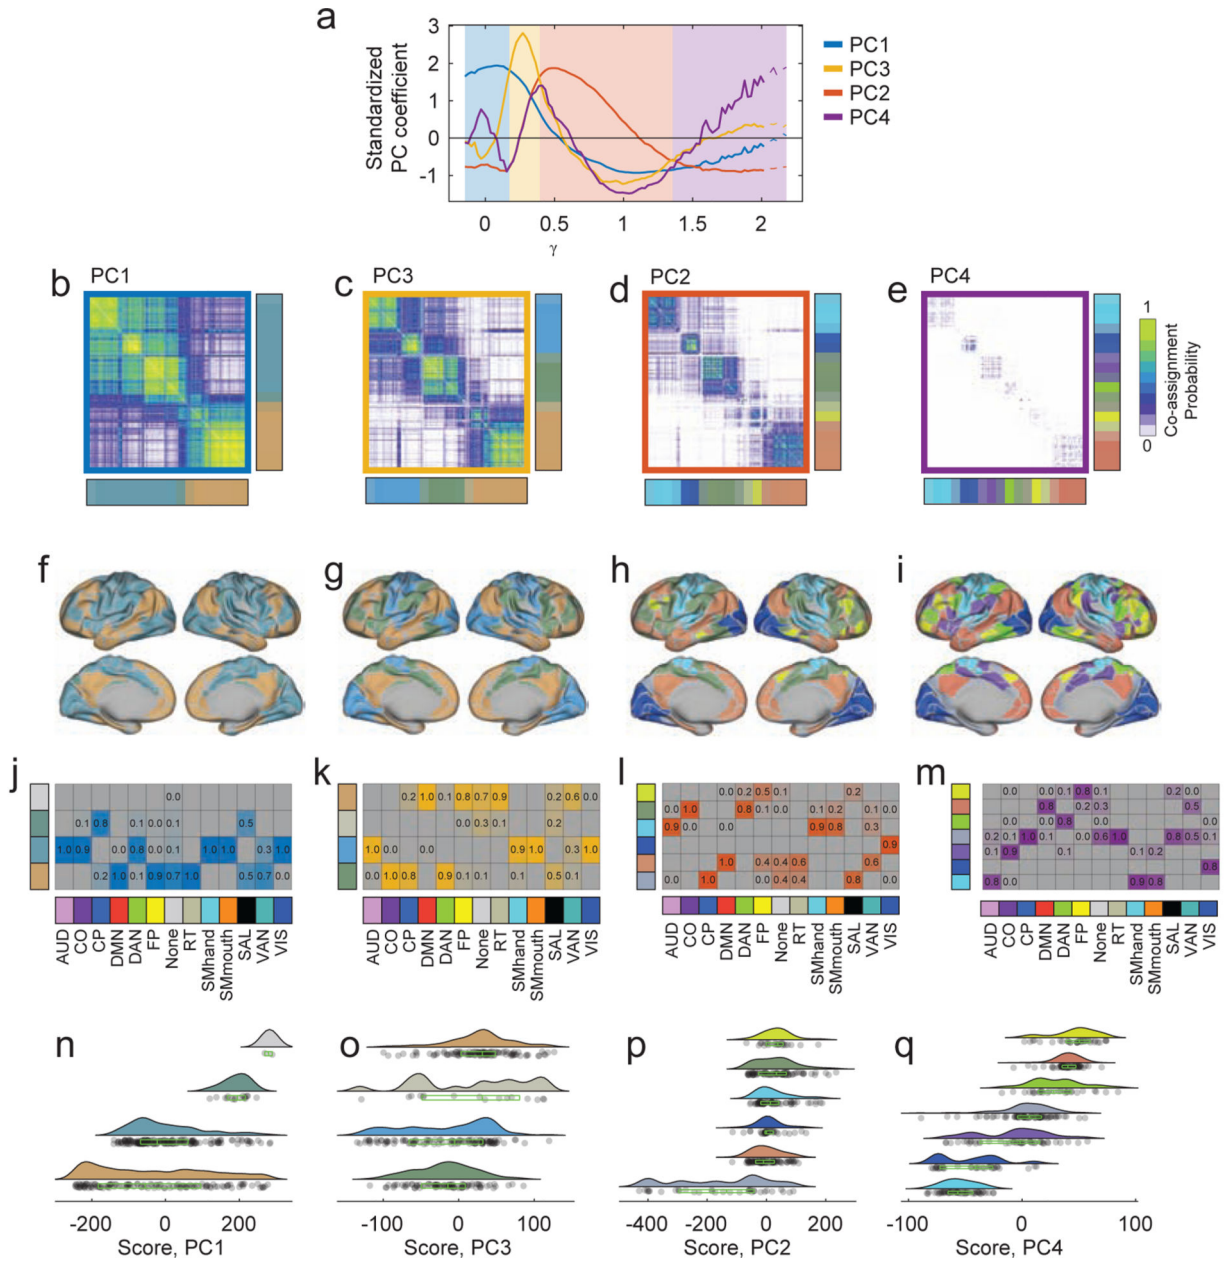


FIG. 5. Community structure changes as a function of γ .

We restricted our analysis to the region of parameter space with $\omega < 10^{-2}$. Within this space, virtually all variation in community structure occurs as a function of γ . We then calculated the average PC coefficient for each of the four PCs as a function of γ and z-scored these values for each PC independently. (a) This procedure enabled us to partition γ values into four segments according to which PC is dominant at that point in parameter space. (b-e) For each segment, which corresponded to a different PC's dominance, we calculated the co-assignment probability for all pairs of nodes. (f-i) We also derived each segments' consensus communities. (j-m) We also assessed how brain systems described in [26] are distributed across detected communities. In these panels, each column corresponds to one of the thirteen systems and each row corresponds to a detected community. Columns were

normalized so that they sum to unity. Within a column the values of cells indicate the fraction of that system's regions that were assigned to each of the detected communities. ($n-q$) We break down the PC scores by detected communities.

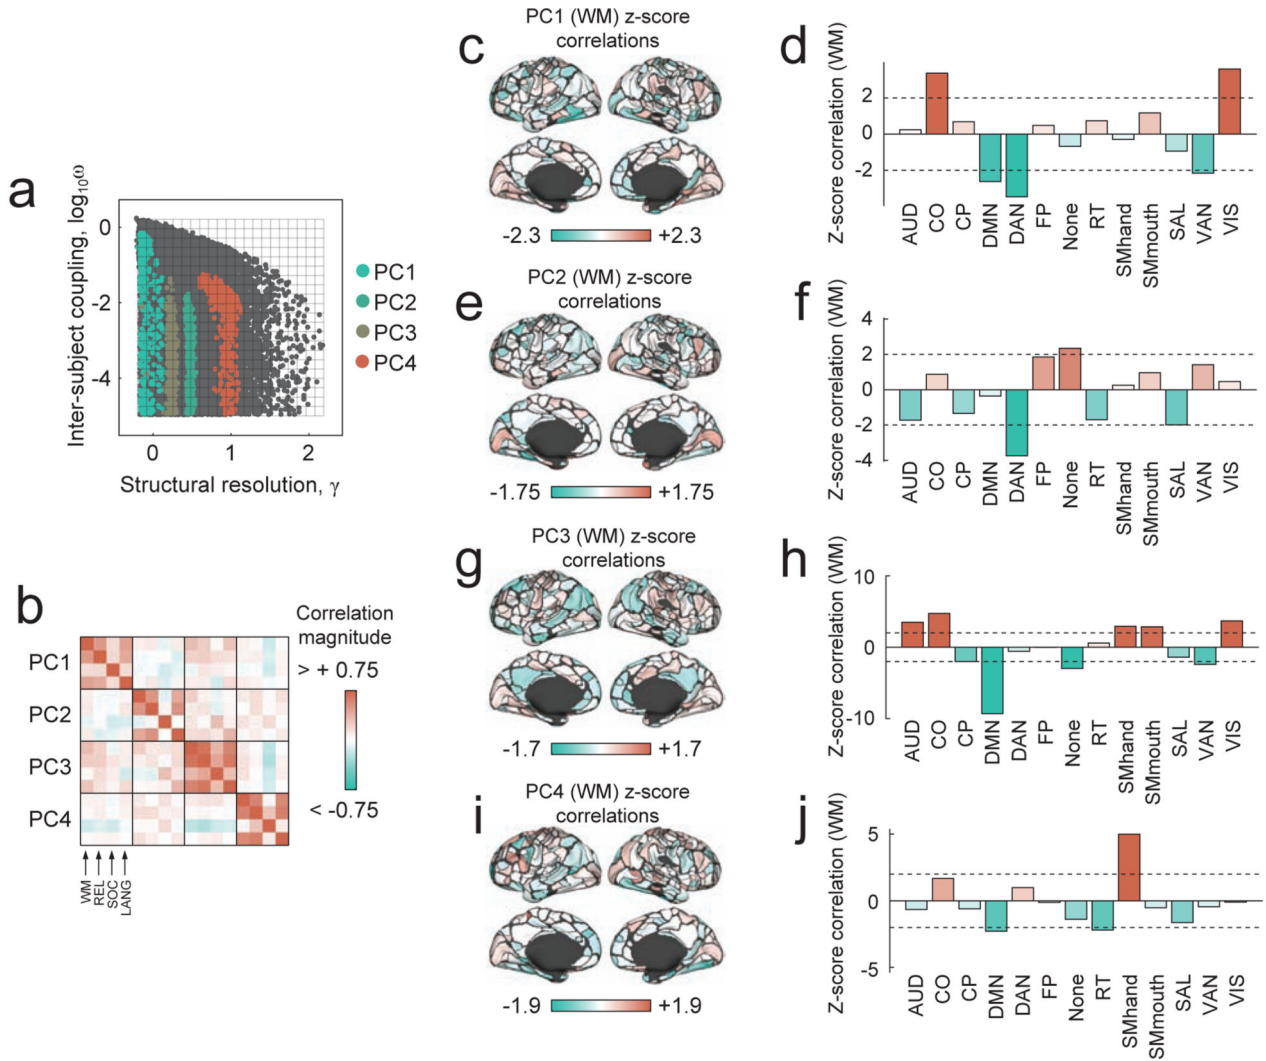


FIG. 6. Correlation of community structure with measures of task performance. (a) For each PC, we studied the sub-sample of partitions corresponding to the 1% largest PC coefficients. (b) For each subsample, we calculated the Pearson correlation coefficient between subjects' community entropy scores and four measures of in-scanner performance on cognitively demanding tasks: working memory (WM), relational (REL), social (SOC), and language (LANG), in the HCP terminology. In panels (c, e, g, i), we show the brain-behavior z-scored correlation coefficients for the first four PCs associated with performance on the WM task, plotted on the cortical surface (permutation test in which behavioral measures were randomly and uniformly shuffled). In panels (d, f, h, j), we show the mean brain-behavior correlation coefficients for the first four PCs, z-scored within each cognitive system. Larger z-scores indicate that the average correlation over all brain areas in a given system is greater than expected in random systems of the same size. Here, as an example, we show correlations for the WM task. Results for other tasks are included in the Supplementary Materials. Note that the dotted lines in panels (d), (f), (h), and (j) correspond to $z = \pm 2$.

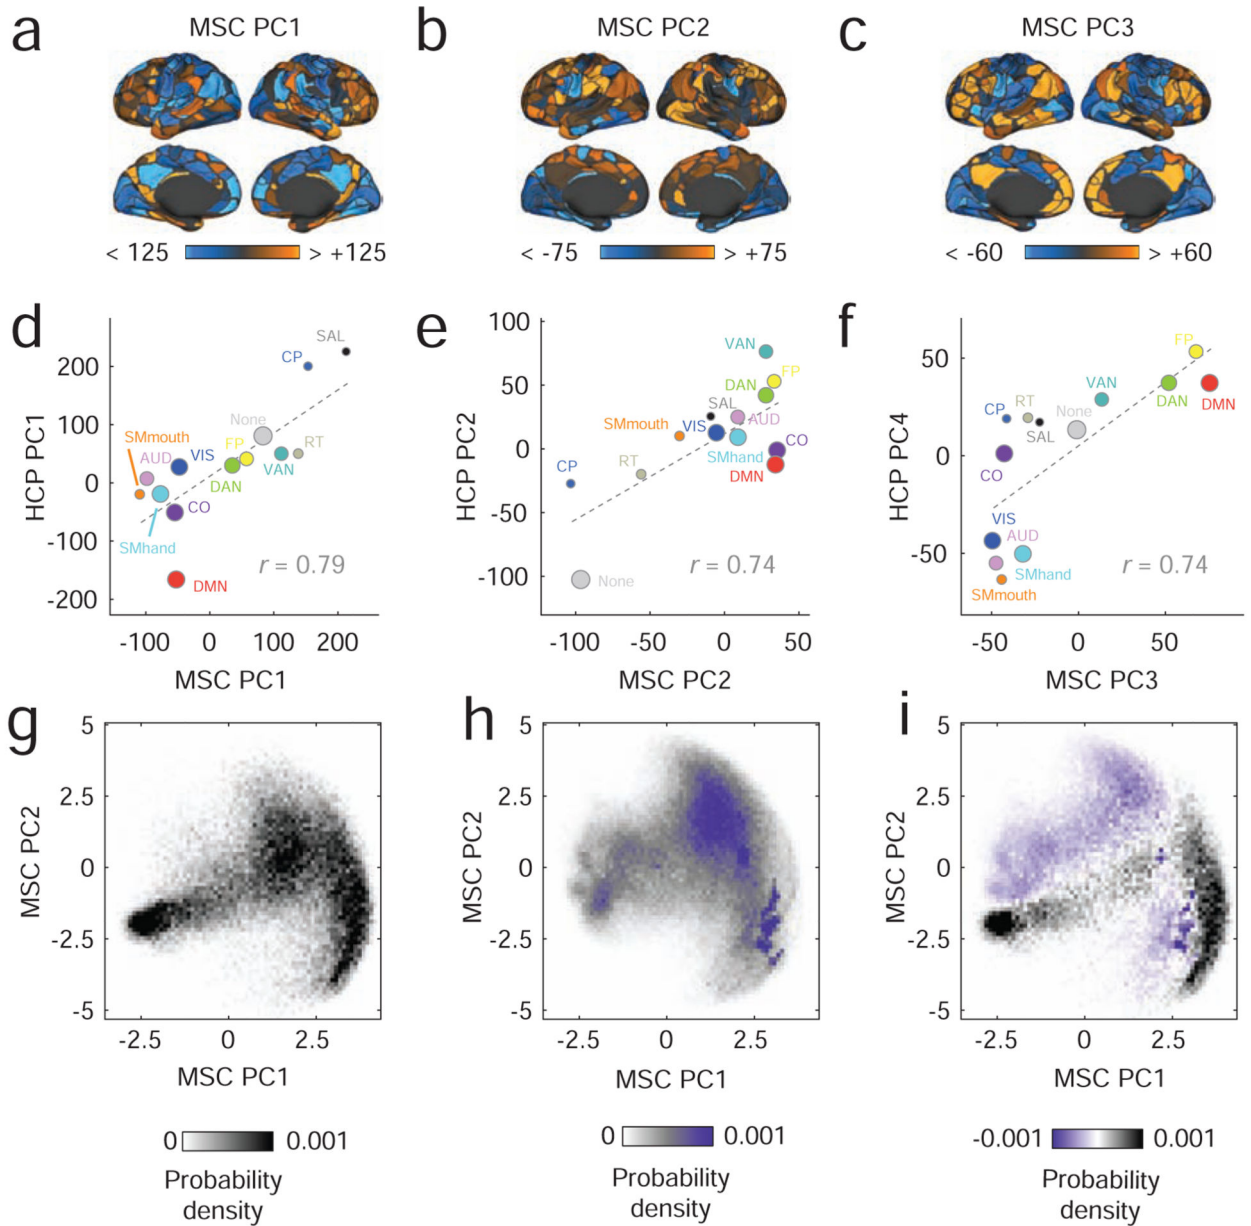


FIG. 7. Summary and analysis of the midnight scan club dataset.

In panels (a)-(c), we show surface maps depicting the first three components generated from PCA analysis of the midnight scan club (MSC) dataset. These principal components correspond, broadly, to the components detected in the HCP dataset. We compare the two datasets by averaging principal components within brain systems and computing system-average correlations. The results are shown in panels (d)-(f). The size of dots is proportional to the number of nodes assigned to each system. We perform a similar analysis of *intra*-subject variability, in which we characterize variation in community structure within subjects across scan sessions. To visualize these results, we project within- and between-subject community entropy scores into the two-dimensional space defined using the first two dimensions of a joint PCA. In panel g, we show the distribution of inter-subject patterns.

Panel *h* shows the same for *intra*-subject patterns, and *i* shows the difference between the two.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript