Subgraph Detection in Functional Brain Networks During Many Diverse Cognitive Tasks

Jonathon Walters
Department of Psychology
Stanford University
waltersj@stanford.edu

1. Introduction

The overarching goal of this work is to better understand how large-scale brain networks support human cognition and behavior. In particular, we are interested in how functional networks vary as a function of the cognitive and behavioral demands of different task contexts. A widely held view is that successfully performing even simple tasks (e.g., finger tapping) requires the dynamic coordination of information integration and segregation across specialized regions and sets of regions ([3]). A relatively stable underlying structural network of white-matter connections supports and constrains this rich dynamic functional repertoire ([20]).

Using functional magnetic resonance imaging (fMRI), networks are typically constructed by first parcellating the brain into hundreds of regions (from tens of thousands of voxels) and then computing pairwise statistical dependencies between regional time courses. These networks have largely been studied statically — that is, a single adjacency matrix is computed for an entire scan or task. Over the past several years, there has been increasing interest in characterizing context-dependent brain network dynamics (e.g., [2]; [18]; [8]; [19]). However, most studies still focus on specific tasks or small sets of tasks in isolation. There is a need for more comprehensive approaches in assessing how brain network dynamics vary across tasks and are driven by tasks' functional requirements (e.g., perceptual, cognitive, behavioral).

A small number of studies have begun investigating network dynamics by identifying subgraphs (i.e., recurring patterns of temporally co-varying edges). These studies apply non-negative matrix factorization (NMF) to fMRI data to identify subgraphs and their temporal expression ([11]). An advantage to NMF is that it can produce time-dependent coefficients, allowing subgraphs to flexibly co-occur over time. Moreover, because the resulting weights are positive, subgraphs are additive and can be more easily interpreted. One study [10], for example, applied NMF to an fMRI dataset of 28 subjects performing two cognitive control tasks (Stroop and Navon tasks), finding that changes in cognitive demand related to subgraph expression. Another study [4] applied similar methodology to resting-state data, showing how the expression of these subgraphs captures dynamic integration between known cognitive systems. A third study [9] investigated how these subgraphs relate to modularity, demonstrating that subgraphs differ in the extent to which they are contained within a module or span two or more modules. However, because these studies measure brain activity in a small number of tasks, they are limited in their generalizability, including, importantly, the subgraphs that are discovered. Moreover, they do not provide a systematic way of probing the cognitive relevance of subgraphs and their dynamics.

In the current study, we seek to extend prior work on subgraphs by a) applying this methodology to a unique dataset that much more comprehensively samples the space of cognitive tasks, and b) characterizing the cognitive and behavioral relevance of subgraphs using a cognitive encoding model framework, which learns to predict subgraph expression for a given task from the cognitive features associated with the task. We used the MultiTask dataset (Fig. 1) ([13]), which consists of 22 subjects performing 26 tasks (47 total task conditions) designed to broadly sample human cognitive functioning. For the encoding models, we used annotations developed by a group of cognitive scientists who had labeled each of the 47 task conditions with cognitive features from a formal knowledge graph — a crowdsourced cognitive science ontology known as the Cognitive Atlas (www.cognitiveatlas.org; [15]; [16]). We believe this dataset and methodology will enable a unique view into the nature of functional subgraphs in the human brain, as well as reveal finer-grained relationships between subgraph dynamics and ongoing changes in cognitive and behavioral contexts over time.

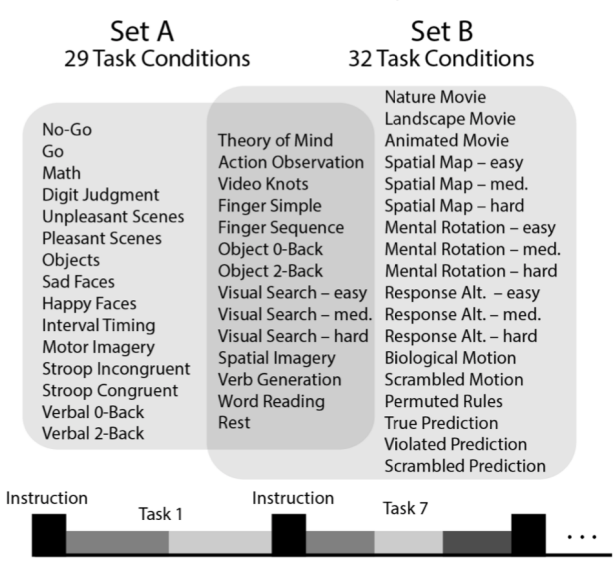

Figure 1. The neuroimaging MultiTask dataset described in [13]. Participants performed 47 task conditions in a battery of 26 tasks while undergoing fMRI. Tasks were chosen to collectively sample broadly from the space of cognitive functions.

2. Methods

2.1. Dataset

fMRI BOLD signal was measured in 22 subjects performing 26 diverse tasks. Each subject completed 32 10-minute runs (approx. 5.3 hrs total), at a temporal resolution of 1 second, with each run consisting of a continuous sequence of 17 tasks (5 second instruction, 30 second task). Tasks varied widely in their functional requirements, from intensive demands like cognitive control and working memory (e.g., stop-signal, Stroop, n-back, CPRO) to more passive perceptual experiences (e.g., viewing a sequence of neutral pictures). See [13] for a complete description of the dataset.

2.2. Preprocessing

All fMRI BOLD timeseries data were preprocessed with standard methods (e.g., confound regression, temporal band pass filtering; see [18] for an example). For computational efficiency and to make inferences at the spatial scale of brain regions and large- scale networks and cognitive systems, we parcellate the brain using the Schaefer 200-region functional parcellation ([17]), which is based on resting-state data and takes into account both local and global signal variability. This parcellation reduces the dimensionality from tens of thousands of voxels to 200 regions. The resulting timeseries of each region is the average of all voxels' timeseries from within a given region.

2.3. Constructing functional networks

Pearson's correlation was computed for each pair of regions in each task instance, creating 9152 signed, weighted, and fully-connected graphs (from 22 subjects, 26 tasks, and 16 samples per task). For each task instance, we discarded data from the first 6 seconds due to the hemodynamic lag (i.e., 6-30 seconds). To construct the configuration matrix for NMF, we concatenated these graphs into a matrix of shape E x T, where $T = 9152$ is the number of time windows (or number of tasks) and $E = N(N-1)/2 = 19000$ total unique edges. Because NMF requires negative values, we vertically stacked this matrix with itself, thresholding each of the two matrices $\mathbf{A}_+$ and $\mathbf{A}_-$ to be above or below 0 (respectively) and taking the absolute value of $\mathbf{A}_-$. This resulted in a final configuration matrix $\mathbf{A}$ of shape $(E * 2)$ x T.

2.4. Decomposing functional brain networks into subgraphs

2.4.1 Non-negative matrix factorization (NMF)

To detect subgraphs, we used an unsupervised learning approach called non-negative matrix factorization (NMF) ([14]), which aims to solve the factorization problem $\mathbf{V} \approx \mathbf{WH}$, where $\mathbf{W}, \mathbf{H} \geq 0$. We apply NMF to our configuration matrix $\mathbf{A}$ to discover a non-orthogonal basis set $\mathbf{W}$ of temporally co-varying edges (i.e, subgraphs) and their coefficients $\mathbf{H}$ across tasks. Thus, the graph at any time window (i.e., for any task) may be expressed as a linear combination of subgraphs, given by the coefficients for that time window. We approximated this factorization with the following cost function:

$$\min_{\mathbf{W},\mathbf{H}} \frac{1}{2}\|\mathbf{A} - \mathbf{WH}\|_F^2 + \alpha\|\mathbf{W}\|_F^2 + \beta \sum_{t=1}^{T} \|\mathbf{H}(:,t)\|_1$$

where $\mathbf{A}$ is the configuration matrix of shape $(E * 2)$ x T, $\mathbf{W}$ is the subgraph matrix of shape E x k, and $\mathbf{H}$ is the matrix of shape k x T with coefficients for each subgraph over all time windows. With our prior concatenation of $\mathbf{A}_+$ and $\mathbf{A}_-$, note that $\mathbf{H}$ is therefore the horizontal stacking of $\mathbf{H}_+$ and $\mathbf{H}_-$. Three parameters need to be set: number of subgraphs k, $L2$ regularization strength α on the subgraphs, and $L1$ regularization strength β on the expression coefficients. To perform NMF, we used alternating non-negative least squares algorithm with the block-pivoting method ([12]), always initializing $\mathbf{W}$ and $\mathbf{H}$ randomly and uniformly from the closed interval $[0, 1]$.

2.4.2 Hyperparameter search

We randomly sampled the parameter space, with k $\sim$U$(5, 30)$, a $\sim$U$(.1, 5)$, b $\sim$U$(0.1, 5)$, using a 5-fold cross-validation scheme. For each of 2000 parameter settings, NMF was run for 100 iterations on each fold, and

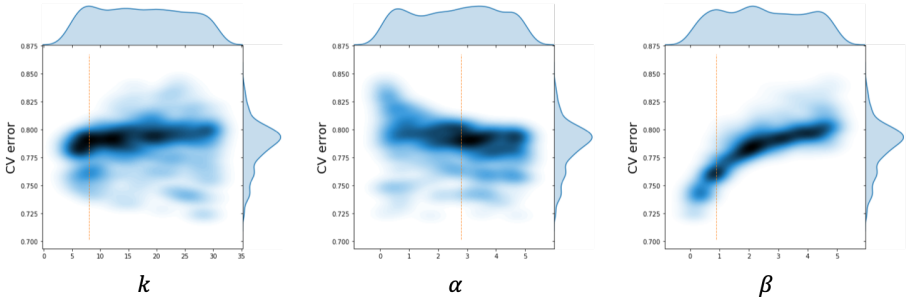

Figure 2. Hyperparameter search for non-negative matrix factorization over number of subgraphs k, subgraph sparsity α, and coefficient sparsity β, for 2000 randomly sampled settings. Plots show kernel density estimation of the marginal distribution of 5-fold (across subjects) CV error for each parameter. Dashed lines indicate optimal values.

the cross-validation error was calculated on the held-out fold using $\|\hat{\mathbf{A}} - \mathbf{WH}\|_F^2$. We sought to identify the setting that minimized this CV error, and to further prevent overfitting ([1]; [10]), we defined the optimal value of each parameter as the average of all values less than or equal to the 25% percentile of the CV error. This procedure led to $k_{opt} = 8$, $a_{opt} = 2.789$, $b_{opt} = 0.896$ (Fig. 2).

2.4.3 Consensus clustering

Because NMF is non-deterministic, we repeated NMF for 50 seeds s using the optimal parameters, concatenating each $\mathbf{W}$ into an ensemble matrix $\mathbf{W}_{ens}$ of shape E x $(s * k_{opt})$. We then ran NMF on $\mathbf{W}_{ens}$ for 100 iterations to identify a final basis set $\mathbf{W}_{opt}$ and then solved for $\mathbf{H}_{opt}$ to recover the expression coefficients.

2.4.4 Null model

We defined the null model as a random configuration matrix $\mathbf{A}_{rand}$ that preserved the degree distribution of each observed graph but randomly shuffled its edges (i.e., for each row of $\mathbf{A}$, randomly shuffle its entries). Then, to identify expression coefficients we used $\mathbf{W}_{opt}$ to solve for $\mathbf{H}_{rand}$.

2.5. System-level topology of subgraphs

To qualitatively assess how a given subgraph might comprise within- and between-system connections amongst previously identified large-scale cognitive systems ([21]), we plot the top 3% of edges in each subgraph, with nodes colored by their system assignment. See Figure 4.

2.6. Subgraph expression

To assess whether a given subgraph was over- or under-expressed on a given task, we calculated average subgraph expression in each task for each subject by aggregating (summing) the coefficients across a subjects 16 observed task graphs (separately for $\mathbf{H}_+$ and $\mathbf{H}_-$) (Fig 4a). We then normalized across subjects for each task (demeaning and scaling to unit variance). We repeated this procedure for $\mathbf{H}_{rand}$. Then, we used a paired sample t-test to assess whether average subgraph expression in a given task significantly differed from the null model (all tests were FDR-corrected for multiple comparisons with an overall significance threshold of 0.05). To visualize deviations from the null model, we calculated the difference between the means of the data and the null model (i.e., data-null) for each subgraph and task (Fig X (center); asterisks indicate significance at FDR-corrected threshold of 0.05).

2.7. Behavioral relevance of subgraphs

We measured overall behavioral accuracy for each subject by calculating mean accuracy across runs for each task that had behavior measurements (17 of the 26 tasks), and then averaging across tasks and normalizing across subjects. To obtain a single expression score per subgraph for each subject, we aggregated (summed) the expression coefficients across tasks. Then, for each subgraph we computed the Pearson's correlation coefficient between behavioral scores and subgraph expression (separately for each edge type).

2.8. Cognitive relevance of subgraphs

Each task had previously been labeled with cognitive features (i.e., entities in the Cognitive Atlas ontology). For each subgraph, we used a cognitive encoding model to learn a mapping from cognitive feature space to expression level. We used multiple linear regression (lasso) to predict subgraph expression from cognitive features, employing a nested cross-validation scheme, optimizing regularization strength α in the inner loop (4-fold CV across the 23 training tasks) and evaluating performance on the 3 held-out tasks.

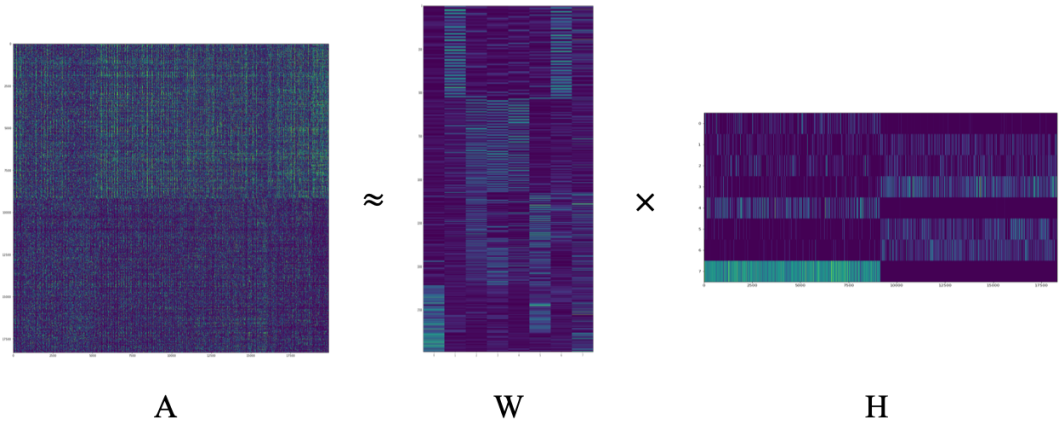

Figure 3. Non-negative matrix factorization results. Each column of the matrix $\mathbf{A}$ is the concatenation of two whole-brain networks extracted from a time window of 24 seconds (starting 6 seconds after a task commenced). For each window, a network of 200 brain regions was constructed by computing the edges as the Pearson's r between the timeseries of every pair of brain regions. This graph was then separated into positive and negative correlations (thresholded above or below zero) and concatenated into a final feature vector. NMF was applied to $\mathbf{A}$ to discover a basis set of subgraphs $\mathbf{W}$ (middle) and their coefficients $\mathbf{H}$ (right).

3. Results

3.1. Subgraph detection

Using the optimal parameters found in the hyperparameter search (Fig. 2), through consensus clustering we identified 8 subgraphs (Figs. 3 and 4).

3.2. Qualitative assessment of subgraph topology at the system-level

Subgraphs appeared to vary in how they connect previously identified large-scale cognitive systems (Fig. 5). By thresholding subgraphs to keep the strongest 3% of edges, we see that subgraphs strongest connections range from almost purely within-system (Subgraph 8) to primarily between-system for a small number of systems (e.g., Subgraph 3).

3.3. Subgraph expression across tasks

For a given subgraph and task, we used a paired t-test (FDR-corrected for multiple comparisons) to determine if average subgraph expression differed from the null model. Plotted in Figure 5 (center left) is the difference in means of the data and null distributions, for the positive (left) and negative (right) components, with an asterisk indicating significance at less than 0.05. Indeed, there were significant relationships for several subgraphs and tasks, for both positive and negative edge types (Fig. 5). That is, subgraphs exhibited unique significance profiles across tasks, suggesting that some subgraphs are more or less relevant for performing particular tasks. A significant relationship in the positive or negative component means that brain regions that are more strongly connected in the given subgraph are more positively or negatively correlated, respectively, in their dynamics. To the extent that positive or negative correlations can be interpreted as reflecting the integration or segregation of information (e.g., [7]), these findings gesture toward high-level constraints in information processing as a function of cognitive and behavioral context.

The strongest overall expression is seen in the positive component of Subgraph 6, with nearly zero expression in the negative component. This subgraph thus represents a pattern of positive correlations present in most (or all) tasks. As the strongest edges in Subgraph 6 are primarily within-system (with some longer-range connections), this subgraph at least qualitatively replicates a prominent pattern of connectivity found not only during the resting-state but also during tasks (e.g., [5]). It remains unclear as to what extent this subgraph can be characterized as task-general, perhaps instead reflecting a large component of connectivity nearly always present in brain function.

3.4. Behavioral relevance

We found subgraph expression to be significantly related to overall behavioral accuracy (Fig. 6). Specifically, the positive expressions of Subgraphs 2, 3, and 7 were positively correlated with behavior, and the negative expression of Subgraph 5 was positively correlated with behavior (FDR-corrected for multiple comparisons, $p < 0.05$).

Interestingly, in Subgraph 2, the strongest edges appear to connect the visual system with three other systems: saliency/ventral attention, dorsal attention, and somatomotor. While follow-up analyses are required to examine these relationships more closely, our result suggests that information integration between the visual system and these three

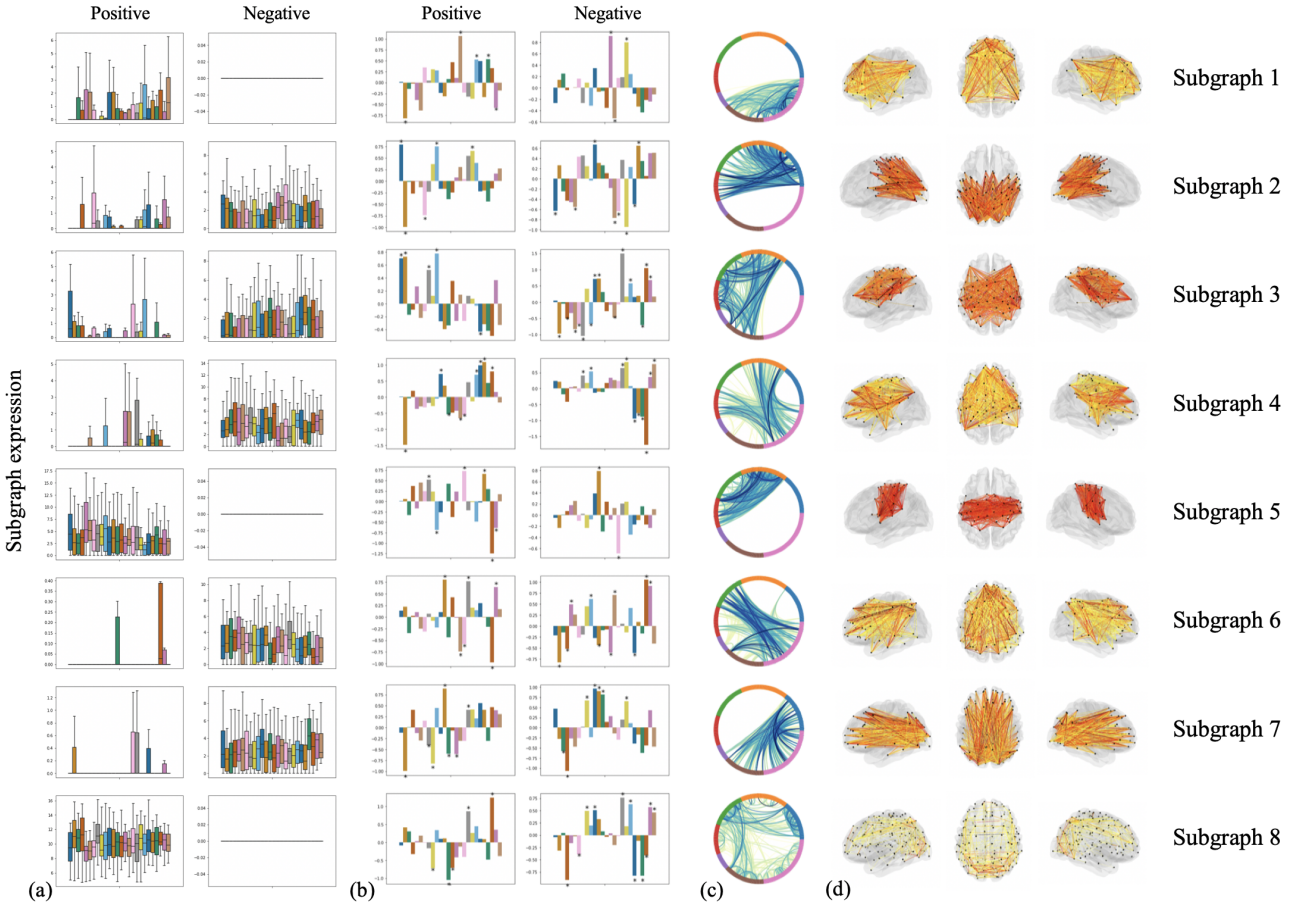

Figure 4. Eight subgraphs detected using non-negative matrix factorization. In (c) and (d), the strongest 3% of connections in each subgraph are shown, colored by strength. In (c), colors around the ring indicate the cognitive system label for a given node (see Fig. 4 for the legend). (a) and (b) indicate subgraph expression across tasks, for the positive and negative components. Each bar shows a distribution of expression across subject. In (a), subgraph expression for a given subject and task is the average across 16 runs. (b) shows the difference in means between the data and null model distributions, and asterisks indicate a significant difference (paired t-test, FDR-corrected for multiple comparisons at a threshold of 0.05).

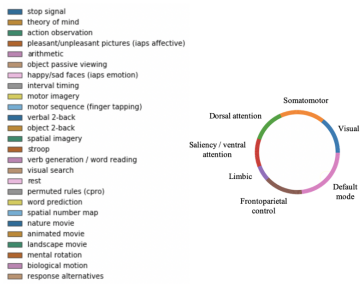

Figure 5. (Left) Task legend. (Right) Color code for the 7 large-scale cognitive systems originally identified in [21]

other systems was beneficial for task performance. Similarly, Subgraph 3 may be behaviorally beneficial by aiding the integration of information among systems that are usually engaged during demanding tasks (frontoparietal control, somatomotor, dorsal attention, and saliency/ventral attention).

3.5. Cognitive relevance

Generalization performance was extremely poor for all values of alpha (see Fig. 7 for a histogram of r-squared values across folds, where $\alpha = 1.0$). Thus, even though subgraphs varied in their expression patterns across tasks, we were unable to find a linear mapping from the cognitive features associated with each task to these expression levels.

4. Discussion

In this study, we used an unsupervised learning approach, non-negative matrix factorization, to discover subgraphs — recurring patterns of temporally co-varying interactions between brain regions — in an fMRI dataset consisting of 117 hours of brain data, with 22 subjects performing

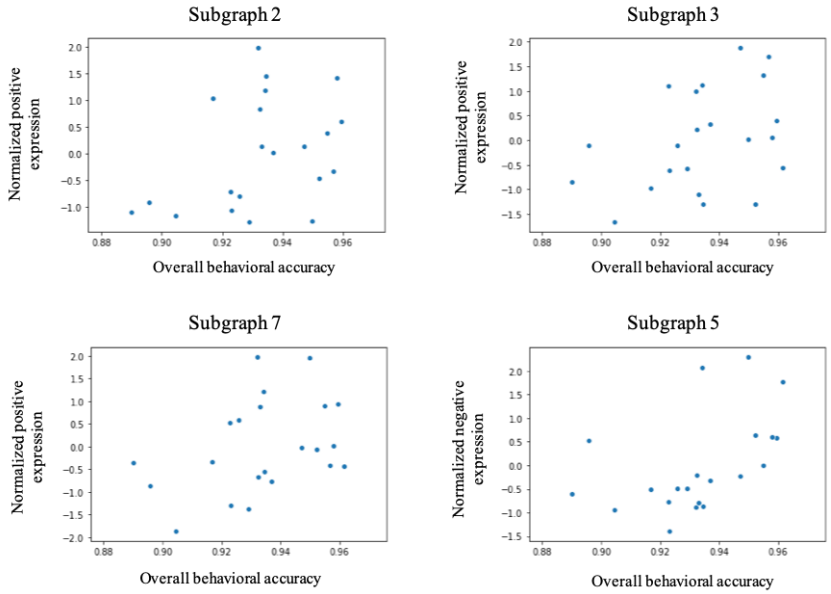

Figure 6. Overall behavioral accuracy in the task battery was related to subgraph expression. The positive expression of Subgraphs 2, 3, and 7 significantly correlated with behavior, and likewise for the negative expression of Subgraph 5 (all FDR-corrected for multiple comparisons at a threshold of 0.05).

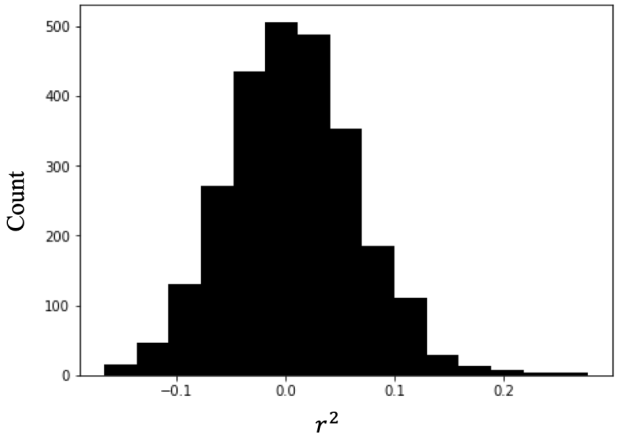

Figure 7. Performance of a cognitive encoding model that attempts to learn a linear mapping from cognitive features to subgraph expression. A group of cognitive scientists had previously labeled the tasks with entities from a knowledge graph (i.e., ontology) known as the Cognitive Atlas. For each subgraph, multiple linear regression was used to predict expression level for a task given a cognitive feature vector. We used nested cross-validation to optimize $L1$ regularization strength α across 4-folds in the inner loop, and computed r-squared values for the three held-out tasks in the outer loop (leave-three-out). The histogram shows r-squared values calculated across splits using the optimal $\alpha = 1.0$.

26 diverse tasks.

First, we found that subgraphs varied in their topological relation to previously identified large-scale cognitive systems. While additional analyses are required to more accurately interpret these results, some tentative qualitative observations can be made. Subgraphs exhibited a preference for patterns of correlated or anticorrelated activity between particular systems. They also appeared to vary in the strength of their within- and between-system connections. Subgraph 8, for instance, showed a more diffuse connectivity pattern, tending to connect regions within the same system (though notably also showed some long-distance connections). Subgraph 5, on the other hand, was highly circumscribed to connections between the somatomotor and saliency / ventral attention systems, and Subgraph 3 showed strong connections amongst systems related to relatively more demanding tasks (frontoparietal control, somatomotor, dorsal attention, saliency/ventral attention). Subgraph 6, which appears to connect the default mode network with the dorsal attention network, was expressed across all tasks, primarily negatively, suggesting a task-general competitive interaction between these two systems. Follow-up work is necessary to more rigorously characterize the topologies of these subgraphs.

Next, we found that subgraphs varied in their prevalence across tasks. Some subgraphs, like the positive component of Subgraph 8, appear to be quite general, while others, like the positive components of subgraphs 1, 2, and 3, were more selectively expressed across tasks. This finding at least partially replicates previous work showing that the brain's community structure (via hard-partitioning of the nodes) has both task-specific and task-general components, and that these network architectures can be viewed as reconfigurations of a set of intrinsic networks present during

both task and non-task states (e.g., [6]; [5]).

We also found that subgraph expression was related to individual differences in overall behavioral performance across the task battery. Specifically, we identified the positive components of three subgraphs and the negative component of one subgraph that positively correlated with overall accuracy. These findings suggest that certain subgraphs may be behaviorally useful in certain contexts, perhaps by facilitating particular patterns of information integration and segregation.

Despite these results, we were not able to explicitly link subgraphs to formal representations of cognitive functions. Our cognitive encoding model, which modeled a subgraph's expression in a given task as a linear function of that task's cognitive features, showed poor generalization accuracy to unseen tasks. This is somewhat surprising given the other results, though future work may be successful as cognitive ontologies improve over time or by using other methods for generating cognitive features and assigning them to tasks.

Collectively, our results provide evidence that subgraphs of functional brain networks may play important, specialized, and context-dependent roles for information integration and segregation. The comprehensiveness of this dataset (i.e., its large number of diverse tasks densely sampled within participants), along with the use of a technique (i.e., NMF) not widely applied to fMRI data, enabled a unique and novel characterization of brain function, presenting numerous avenues for future research to pursue.

References

[1] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb):281–305, 2012. 3

[2] Urs Braun, Axel Schäfer, Henrik Walter, Susanne Erk, Nina Romanczuk-Seiferth, Leila Haddad, Janina I Schweiger, Oliver Grimm, Andreas Heinz, Heike Tost, et al. Dynamic reconfiguration of frontal brain networks during executive cognition in humans. *Proceedings of the National Academy of Sciences*, 112(37):11678–11683, 2015. 1

[3] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience*, 10(3):186, 2009. 1

[4] Lucy R Chai, Ankit N Khambhati, Rastko Ciric, Tyler M Moore, Ruben C Gur, Raquel E Gur, Theodore D Satterthwaite, and Danielle S Bassett. Evolution of brain network dynamics in neurodevelopment. *Network Neuroscience*, 1(1):14–30, 2017. 1

[5] Michael W Cole, Danielle S Bassett, Jonathan D Power, Todd S Braver, and Steven E Petersen. Intrinsic and task-evoked network architectures of the human brain. *Neuron*, 83(1):238–251, 2014. 4, 7

[6] Michael W Cole, Jeremy R Reynolds, Jonathan D Power, Grega Repovs, Alan Anticevic, and Todd S Braver. Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature neuroscience*, 16(9):1348, 2013. 7

[7] Alex Fornito, Ben J Harrison, Andrew Zalesky, and Jon S Simons. Competitive and cooperative dynamics of large-scale brain functional networks supporting recollection. *Proceedings of the National Academy of Sciences*, 109(31):12788–12793, 2012. 4

[8] Peter Fransson, Björn C Schiffler, and William Hedley Thompson. Brain network segregation and integration during an epoch-related working memory fmri experiment. *Neuroimage*, 178:147–161, 2018. 1

[9] Ankit N Khambhati, Marcelo G Mattar, Nicholas F Wymbs, Scott T Grafton, and Danielle S Bassett. Beyond modularity: Fine-scale mechanisms and rules for brain network reconfiguration. *NeuroImage*, 166:385–399, 2018. 1

[10] Ankit N Khambhati, John D Medaglia, Elisabeth A Karuza, Sharon L Thompson-Schill, and Danielle S Bassett. Subgraphs of functional brain networks identify dynamical constraints of cognitive control. *PLoS computational biology*, 14(7):e1006234, 2018. 1, 3

[11] Ankit N Khambhati, Ann E Sizemore, Richard F Betzel, and Danielle S Bassett. Modeling and interpreting mesoscale network dynamics. *NeuroImage*, 180:337–349, 2018. 1

[12] Jingu Kim, Yunlong He, and Haesun Park. Algorithms for nonnegative matrix and tensor factorizations: A unified view based on block coordinate descent framework. *Journal of Global Optimization*, 58(2):285–319, 2014. 2

[13] Maedbh King, Carlos R Hernandez-Castillo, Russell A Poldrack, Richard B Ivry, and Jörn Diedrichsen. Functional boundaries in the human cerebellum revealed by a multi-domain task battery. *Nature neuroscience*, 22(8):1371–1378, 2019. 1, 2

[14] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788, 1999. 2

[15] Russell A Poldrack, Aniket Kittur, Donald Kalar, Eric Miller, Christian Seppa, Yolanda Gil, D Stott Parker, Fred W Sabb, and Robert M Bilder. The cognitive atlas: toward a knowledge foundation for cognitive neuroscience. *Frontiers in neuroinformatics*, 5:17, 2011. 1

[16] Russell A Poldrack and Tal Yarkoni. From brain maps to cognitive ontologies: informatics and the search for mental structure. *Annual review of psychology*, 67:587–612, 2016. 1

[17] Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and BT Thomas Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral Cortex*, 28(9):3095–3114, 2017. 2

[18] James M Shine, Patrick G Bissett, Peter T Bell, Oluwasanmi Koyejo, Joshua H Balsters, Krzysztof J Gorgolewski, Craig A Moodie, and Russell A Poldrack. The dynamics of functional brain networks: integrated network states during cognitive task performance. *Neuron*, 92(2):544–554, 2016. 1, 2

[19] James M Shine, Michael Breakspear, Peter T Bell, Kaylena A Ehgoetz Martens, Richard Shine, Oluwasanmi Koyejo, Olaf Sporns, and Russell A Poldrack. Human cognition involves the dynamic integration of neural activity and neuro-

modulatory systems. *Nature neuroscience*, 22(2):289, 2019. 1

[20] Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: a structural description of the human brain. *PLoS computational biology*, 1(4):e42, 2005. 1

[21] BT Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*, 106(3):1125–1165, 2011. 3, 5