Neuron Article



Jonathan D. Power,^{1,*} Alexander L. Cohen,¹ Steven M. Nelson,² Gagan S. Wig,¹ Kelly Anne Barnes,¹ Jessica A. Church,¹ Alecia C. Vogel,¹ Timothy O. Laumann,¹ Fran M. Miezin,^{1,3} Bradley L. Schlaggar,^{1,3,4,5} and Steven E. Petersen^{1,2,3,5,6,7}

¹Department of Neurology ²Department of Psychology ³Department of Radiology ⁴Department of Pediatrics ⁵Department of Anatomy and Neurobiology ⁶Department of Neurosurgery ⁷Department of Biomedical Engineering Washington University in Saint Louis, St. Louis, MO 63130, USA *Correspondence: powerj@wusm.wustl.edu DOI 10.1016/j.neuron.2011.09.006

SUMMARY

Real-world complex systems may be mathematically modeled as graphs, revealing properties of the system. Here we study graphs of functional brain organization in healthy adults using resting state functional connectivity MRI. We propose two novel brain-wide graphs, one of 264 putative functional areas, the other a modification of voxelwise networks that eliminates potentially artificial short-distance relationships. These graphs contain many subgraphs in good agreement with known functional brain systems. Other subgraphs lack established functional identities; we suggest possible functional characteristics for these subgraphs. Further, graph measures of the areal network indicate that the default mode subgraph shares network properties with sensory and motor subgraphs: it is internally integrated but isolated from other subgraphs, much like a "processing" system. The modified voxelwise graph also reveals spatial motifs in the patterning of systems across the cortex.

INTRODUCTION

Advances in neuroimaging that facilitate the study of brain relationships in humans have stimulated an enormous amount of scientific and medical interest in recent years (Biswal et al., 1995; Bullmore and Sporns, 2009; Deco et al., 2011; Dosenbach et al., 2010). Resting state functional connectivity MRI (rs-fcMRI), which measures spontaneous low-frequency fluctuations in blood oxygen level dependent (BOLD) signal in subjects at rest, has attracted particular attention for its ability to measure correlations in neural activity (via BOLD signal) between distant brain regions. These correlations are of great interest to the medical community because an increasing number of pathologic conditions appear to be reflected in functional connectivity between particular brain regions (Church et al., 2009; Seeley et al., 2009). At the same time, these correlations are of fundamental interest to neuroscientists because they offer the first opportunity to comprehensively and noninvasively explore the functional network structure of the human brain (Bullmore and Sporns, 2009).

Although a variety of methods may be used to study rs-fcMRI data, one of the most powerful and flexible approaches is the graph theoretic approach (Bullmore and Sporns, 2009; Rubinov and Sporns, 2010). Within this framework, a complex system is formalized as a mathematical object consisting of a set of items and a set of pairwise relationships between the items. Items are called nodes, relationships are called ties, and collections of these nodes with their ties are called graphs or networks. A short and incomplete list of established topics in graph theory includes quantifying hierarchy and substructure within a graph, identifying hubs and critical nodes, determining how easily traffic flows in different portions and at different scales of a network, and estimating the controllability of a system (Liu et al., 2011; Newman, 2010). Because graph theoretic analyses can model properties at the level of the entire graph, subgraphs, or individual nodes, and because the brain itself is a complex network, graph theoretic approaches are a natural and attractive choice for rs-fcMRI analysis.

A current obstacle to the graph-based study of functional brain organization is that it very difficult to define the individual nodes that make up a brain network. On first principles, treating a graph as a model of a real system, if the nodes of the graph do not accurately represent real items in the system, the graph itself is a distorted model and graph theoretic properties will diverge from the true properties of the system (Butts, 2009; Smith et al., 2011; Wig et al., 2011). The brain is a complex network with macroscopic organization at the level of functional areas and subcortical nuclei, but the number and locations of these entities in humans is largely unknown. Standard approaches to forming whole-brain rs-fcMRI graphs often ignore this issue and define nodes as voxels (Buckner et al., 2009; Cole et al., 2010; Fransson et al., 2011; Tomasi and Volkow, 2011; van den Heuvel et al., 2008), large parcels from anatomically based brain atlases (Hartman et al., 2011; He et al., 2009; Meunier et al., 2009a; Spoormaker et al., 2010; Tian et al., 2011), or random interpolations between voxels and parcels (Hayasaka and Laurienti, 2010; Meunier et al., 2009b). These approaches are not meant to correspond to macroscopic "units" of brain organization, and thus there is no direct reason to believe that these approaches result in well-formed nodes (Wig et al., 2011).

An overarching goal of this report is to, at least partially, overcome this obstacle. We have developed methods to define, as best we can, a set of more appropriate nodes, and to define a network based upon these nodes (and the ties between them). We also propose a second novel brain network, based on a modification of voxel-wise approaches, and examine some of its properties in relation to the first graph. Before studying these graphs in detail, we are obliged to demonstrate that they (1) display signs of accuracy, and (2) improve upon previous graph definitions.

Our evaluation of rs-fcMRI brain graphs rests upon a simple and fundamental argument. Decades of PET and fMRI experiments have defined functional systems as groups of brain regions that coactivate during certain types of task (e.g., the dorsal attention system, (Corbetta and Shulman, 2002; Corbetta et al., 1995); here and elsewhere we replace common neuroscientific usage of "network" with "system," reserving the word network for the graph theoretic sense, such that "dorsal attention network" becomes "dorsal attention system"). A more recent large literature indicates that rs-fcMRI signal is specifically and highly correlated within these functional systems (e.g., within the visual system, default mode system, dorsal attention system, ventral attention system, auditory system, motor system, etc.) (Biswal et al., 1995; Dosenbach et al., 2007; Fox et al., 2006; Greicius et al., 2003; Lowe et al., 1998; Nelson et al., 2010a). There is a family of methods (subgraph detection) that is used to break large networks into subnetworks of highly related nodes (subgraphs), such that nodes within subgraphs are more densely connected (here, correlated) to one another than to the rest of the graph. We hypothesized that specific patterns of high correlation within functional systems would be reflected as subgraphs within a brain-wide rs-fcMRI network. Thus, the presence of subgraphs that correspond to functional systems is an indication that a graph accurately models some features of brain organization, and the absence of such subgraphs raises suspicions that a graph may not be well-defined.

With this hypothesis in mind, we open this report by studying the subgraph structures of four brain-wide graphs within a single data set. As mentioned above, two novel graphs are studied: a graph of putative functional areas (264 nodes), and a modification of voxelwise networks that excludes short-distance correlations (40,100 nodes). Two other standard graphs are used for comparison: a graph of parcels from a popular brain atlas (90 nodes), and a standard voxelwise graph (40,100 nodes). To presage the results, subgraphs in the areal network are significantly more like functional systems than subgraphs in the atlasbased graph, and subgraphs in the modified voxelwise network are more like functional systems than the standard voxelwise network. Additionally, despite great differences in network size and definition, the areal and modified voxelwise subgraphs are remarkably alike and contain many subgraphs corresponding to known functional systems, bolstering confidence in their accuracy. Given these findings, we perform a variety of further analyses upon the novel graphs to learn more about functional brain organization, with some novel and interesting results.

Comparing Networks: Defining Four Brain-Wide Networks

Two novel and two standard methods of graph definition were examined within a large cohort of healthy young adults (and in a matched replication cohort; see Table S1 available online). To reiterate, graphs are composed of a set of nodes and a set of ties between nodes. Graphs were formed using the nodes described below, and ties were defined using Pearson correlation coefficients between node rs-fcMRI timecourses. The cross correlation matrix of a set of nodes thus defines a graph. Because most graph theoretic techniques are developed (and are most meaningful) in sparse graphs (Newman, 2010), thresholds were applied to the graphs to eliminate weak ties (such that correlations under the threshold were ignored). Because there is no "correct" threshold, all analyses were performed over a range of thresholds, typically beginning around 10% tie density (retaining the strongest 10% of correlations) and rising until the networks became severely fragmented (see Supplemental Experimental Procedures).

The first novel graph (referred to as the areal graph) was defined in accord with neurobiological principles. The brain is a complex network with a hierarchical spatial and functional organization (in the cortex) at the level of neurons, local circuits, columns, functional areas, and functional systems. Standard rs-fcMRI analyses use cubic voxels that are a few millimeters on each side, and thus can potentially resolve brain relationships at the level of areas. Centers of putative areas were identified using two independent methods operating on data sets that were not used in graph analyses (see Experimental Procedures). The first method was meta-analytic in nature (as in Dosenbach et al., 2006), and explored a large fMRI data set to identify voxels that were reliably and significantly modulated when certain behaviors were demanded (e.g., buttonpressing) or certain signal types were found (e.g., error-related activity). The second method extended a recently developed technique of mapping cortical areas using rs-fcMRI to entire cortical sheets (fc-Mapping) (Barnes et al., 2011; Cohen et al., 2008; Nelson et al., 2010a). The combination of these methods yielded 264 putative areas spanning the cerebral cortex, subcortical structures, and the cerebellum (see Experimental Procedures, Figure S1, and Table S1 for analysis details, and Table S2 for coordinates). Regions of interest (ROIs) were modeled as 10 mm diameter spheres. Graphs were formed using ROIs as nodes (n = 264) and ties terminating within 20 mm of a source node center were set to zero to avoid possible shared signal between nearby nodes. This procedure yielded graphs of putative functional areas in which each node represented, to the best of our capabilities, an element of brain organization.

The second novel graph that was examined was a modification of voxelwise networks in which all short-distance ties were excluded. This modification arose from several practical observations. First, nearby voxels share nonbiological signal (causing increased rs-fcMRI correlation), a result of unavoidable steps in data processing (e.g., reslicing, blurring). Second, shortdistance relationships are especially susceptible to spurious







Subgraphs replicate across cohorts



Subgraphs change hierarchically over thresholds

Spheres: areal, main cohort Surfaces: modified voxelwise, replication cohort



Figure 1. Areal Subgraph Structure Is Highly Similar across Cohorts and Subgraph Structure Is Similar between Areal and Modified Voxelwise Graphs

Top left: A spring embedded layout of the areal graph at 4% tie density visualizing the graph and the basis for subgraphs.

Top right: For both cohorts, plots are shown of the areal assignments into subgraphs (colors) at tie densities from 10% down to 2% in 1% steps. ROI ordering is identical, and all subgraphs with fewer than four members are colored white. The standard measure of subgraph similarity, normalized mutual information, between node assignments of the cohorts at identical tie densities ranged from 0.86 to 0.92, indicating highly similar patterns across cohorts (1 = identical assignments, 0 = no information shared between assignments).

Bottom: subgraphs from three thresholds are shown for the areal (spheres) and modified voxelwise graphs (surfaces). Note the similarity of subgraph assignments between networks, despite the great difference in network size and cortical coverage, even in different subjects (main versus replication cohorts). All areal subgraphs with fewer than four members are colored white, and all modified voxelwise subgraphs with fewer than 100 voxels are colored white. Areal networks are shown at 10%, 3%, and 2% tie density (r > 0.16, 0.30, and 0.33), and modified voxelwise networks are shown at 5%, 2%, and 0.5% tie density (r > 0.16, 0.23, and 0.31).

shorten addresses of individual nodes). Other algorithms were tested and yielded similar results (Figure S2).

Figure 1 illustrates our methodology and highlights several important results. The first panel

augmentation by subject motion (Power et al., 2011). Third, as will be seen shortly, voxelwise graphs are dominated at higher thresholds by short-distance relationships, which are logically partially artificial based on the above considerations. Modified voxelwise networks are presented in which all ties terminating within 20 mm of a source node are excluded, though other distances (e.g., 15 mm and 25 mm) were also tested, with similar results (data not shown).

The two standard methods of graph formation were parcelbased and voxel-based. The parcel-based graph was formed using the 90-parcel AAL atlas (Tzourio-Mazoyer et al., 2002), a popular method of graph formation. This atlas divides the cortex and subcortical structures into parcels based upon anatomical landmarks. The voxel-based graph was defined using all voxels within the AAL atlas (n = 40,100), and the modified voxelwise graph was also defined using these voxels.

Comparing Networks: Correspondence between Subgraphs and Functional Systems

Subgraphs were determined over a range of thresholds for each graph using one of the best-performing subgraph detection algorithms currently available (Infomap) (Fortunato, 2010; Rosvall and Bergstrom, 2008). This algorithm uses the map equation to minimize information theoretic descriptions of random walks on the graph (essentially assigning zip codes to subgraphs to

depicts the areal graph in a spring embedded layout and maps subgraphs onto nodes using colors, visibly demonstrating the basis for subgraphs. In spring embedded layouts, ties act as springs to position nodes in space such that well-connected groups of nodes are pulled together, providing an intuitive and informative picture of the graph. The second panel shows the subgraph assignments of the areal network in both cohorts over a range of thresholds (each chart consists of 9 columns of 264 color entries). ROIs are ordered identically for both cohorts, and the patterns of subgraph assignment across cohorts are in good agreement. The standard graph theoretic measure of similarity between two sets of node assignments is normalized mutual information (NMI), which measures how much information one set of assignments provides about another set of assignments. Values of 1 indicate identical assignments, and values of 0 indicate that no information is gained about the second set of assignments by knowing the first. Between cohorts, NMI ranges from 0.86 to 0.92 across thresholds, indicating very similar assignments.

The subgraph charts contain subgraphs whose composition remains quite constant over thresholds (e.g., the horizontal bands of blue, red, or yellow) as well as subgraphs that are hierarchically refined as thresholds rise (e.g., cyan becoming cyan, orange, pink, and purple). These patterns can be seen on brain surfaces (Figure 1, bottom) as relatively constant subgraph



Figure 2. Many Modified Voxelwise Subgraphs Replicate across Cohorts and Even within Single Subjects

Select subgraphs from the modified voxelwise analysis are presented from a dorsal view for both cohorts and for an additional single subject. Cohort subgraphs are taken from the 2% tie density analysis and subgraphs in the individual are taken from a 0.5% tie density analysis. The overall NMI between cohort assignments at this threshold was 0.71, and NMI values between subgraphs from different cohorts are shown in the matrix to the right. Additional views of this data and replications of subgraphs from other thresholds are found in Figure S3.

compositions for visual (blue), default (red), or fronto-parietal (yellow) regions over thresholds, and as refinement of the large cyan subgraph into hand somatosensory-motor (cyan), face somatosensory-motor (orange), auditory (pink), and cingulo-opercular (purple) subgraphs. This bottom panel of Figure 1 plots areal assignments (spheres) in the main cohort over the modified voxelwise assignments (surfaces) in the replication cohort, demonstrating the similarity of subgraphs over thresholds across different cohorts and even across graph definitions. As Figure 2 shows, the modified voxelwise graphs also replicate well across cohorts and even in single subjects. Fuller visualizations of these data and replications of subgraphs from other thresholds are found in Figure S3.

We predicted that well-formed graphs would possess wellformed subgraphs corresponding to major functional systems of the brain. Figure 3 gives an overview of how well each network met this prediction. At left, PET and fMRI data defining major functional systems are shown. The next three columns display subgraphs from a single threshold of analysis for each graph (a high threshold, tailored to each graph). In the second column, areal and modified voxelwise assignments are shown simultaneously because they are in such good agreement. The areal and modified voxelwise graphs contain subgraphs that correspond to each of the functional systems, and these subgraphs contain most or all of the brain regions implicated in the functional systems, and sometimes also some extra brain regions. In contrast, the AAL-based graph is incapable of representing most functional systems at this threshold (or any threshold; see Figure S4). The standard voxel-based graph represents some functional systems well (e.g., the default mode system), but others are only incompletely represented. Examination of other thresholds of the standard voxelwise graph (Figure S4) indicates that at low to moderate thresholds, reasonable subgraph representations of some functional systems are found, but that as thresholds rise, portions of functional systems tend to merge, and subgraphs come to resemble a patchwork of local subgraphs across the cortex (see circled regions in Figure S4).

To more quantitatively assess subgraph correspondence to functional systems, we used NMI to compare groups of coordinates from functional systems with the subgraph identities of the nodes nearest to the coordinates under each network definition. A one-factor ANOVA of NMI demonstrates an effect of graph ($p < 10^{-7}$; see Figure S5). The AAL-based graph displays the lowest correspondence (NMI = 0.37 ± 0.04 , significantly lower than all other graphs) across thresholds, and the variable structure of the voxelwise graph is reflected in NMI that ranges widely over thresholds (0.58-0.86), in contrast to the stable and high NMI found in the areal (0.72 ± 0.06) and modified voxelwise graphs (0.87 ± 0.15). Importantly, as thresholds rise, NMI between functional systems and subgraphs increases for the standard voxelwise analysis.

Choosing Network Definitions for Further Analysis

The areal and modified voxelwise graphs best meet our predictions about the correspondence between functional systems and subgraphs within brain-wide networks. The poorer correspondence in the AAL-based and standard voxelwise graphs likely results from coarse, nonfunctionally based nodes in the AAL-based graph, and the effects of millions of artificially high short-range correlations between nearby voxels in the standard voxelwise graph. We turn now from our focus upon confirmatory findings to novel observations about functional brain organization that can be drawn from the areal and modified voxelwise graphs. We shall continue to focus on the network at the level of subgraphs. We begin by discussing the identities of subgraphs, then examine the relationships and properties of particular subgraphs, and end with observations about relationships between all subgraphs.



Figure 3. Graph Definition Dictates Fidelity to Functional Brain Organization

At left, the task-defined locations of four established functional systems. The next three columns display, for the main cohort, the single subgraph that best corresponds to each functional system under the four graph definitions. Circles are placed around small portions of subgraphs that might otherwise be overlooked (there are small green regions within green circles). Data from a single threshold tailored to each graph are shown. The threshold was the next-to-highest threshold that each graph can achieve before the graph becomes severely fragmented (defined by the giant component containing <50% of the nodes in the graph). Tailored thresholds were 3% for the areal graph, 5% for the AAL-based graph, and 2% for both voxel-based graphs. Correspondence between these functional systems and subgraphs is good for the areal and modified voxelwise graphs, intermediate for the voxelwise graph, and poor for the AAL-based graph. Note especially the correspondence between areal (spheres) and modified voxelwise (surface) subgraphs, despite great differences in network size (n = 264 versus n = 40,100). See Figure S4 and Figure S5 for more comprehensive and quantitative presentations of subgraph assignments. Images in the left column are modified from (Corbetta et al., 2008; Corbetta and Shulman, 2002; Dosenbach et al., 2007; Shulman et al., 1997).

Subgraph Identities

The identities of the red (default), yellow (fronto-parietal task control), green (dorsal attention), and teal (ventral attention) subgraphs are already clear. The remaining major subgraphs are now considered.

Several subgraphs correspond to sensory and motor regions (Figure 4, left). A visual system (blue) was identified, spanning most of occipital cortex, often including a small portion of superior parietal cortex and a portion of the postero-lateral thalamus (potentially lateral geniculate nucleus [LGN], see horizontal sections). At moderate thresholds, somatosensory-motor (SSM) cortex (S1, M1, and some pre- and postcentral-gyrus cortex) was divided into dorsal (cyan) and ventral (orange) subgraphs. These subgraphs also included voxels in the parietal operculum that likely correspond to the second somatosensory area (S2) (Burton et al., 2008), as well as a portion of the thalamus possibly corresponding to ventral posterior thalamus (VP). At high thresholds, an auditory subgraph (pink) emerged from the purple cingulo-opercular subgraph.

Rather than a division between somatosensory and motor regions, a division between dorsal and ventral SSM regions is found. Although motor and sensory function are typically localized to the pre- and postcentral gyri, respectively, classic descriptions of stimulus-evoked responses and sensations in humans indicate that these processes are not exclusively localized to either side of the central sulcus (Penfield and Boldrey, 1937), a finding consistent with recent investigations of primary motor and somatosensory cortex in rodents (Matyas et al., 2010). The division into ventral and dorsal subgraphs roughly separates the face from the rest of the body, a distinction confirmed by button-pushing and verb generation meta-analysis data (Figure S1). Similar dorsal/ventral distinctions have recently been found (Yeo et al., 2011). Intriguingly, correlations between meta-analytic face SSM (orange) and auditory (pink) ROIs are higher than correlations between body SSM (cyan) and auditory ROIs (auditory-face r = 0.16, auditory-hand r = 0.05, p < 0.001, significant in both cohorts). These differential correlations are unlikely to reflect only anatomical connectivity, but instead might be related to the history of coactivation that these regions surely share as a function of oral/aural language. Thus, it appears that somatosensory and motor cortex are functionally divided into a ventral facial representation and a dorsal representation of the rest of the body (called "hand" for brevity).

Two cingulo-opercular subgraphs (black and purple, Figure 4, middle) are identified, both encompassing regions in anterior cingulate/medial superior prefrontal cortex (aCC), anterior prefrontal cortex (aPFC), and the anterior insula (al) (with additional regions in inferior and middle frontal gyrus and supramarginal gyrus at multiple thresholds). Two distributed functional systems have been ascribed to cingulo-opercular



Figure 4. Subgraph Identities

Left: visual (blue), auditory (pink), and hand (cyan) and face (orange) sensory-somatomotor (SSM) subgraphs are shown for the areal network at 2% (spheres) and the modified voxelwise network at 0.5% tie density (surface). The mean correlations in the main cohort between auditory processing (pink, MNI: -38 - 33 17) and hand (cyan, -40 - 19 54) and face (orange, -49 - 11 35) regions are shown below. Auditory-face correlations are significantly higher than auditory-hand correlations in both cohorts (p < 0.001, two-sample two-tail t test).

Bottom: slices from the 4% tie density modified voxelwise analysis, with labels on relevant thalamic nuclei (numbers are z coordinates).

Middle: two cingulo-opercular subgraphs shown from the 3% areal (spheres) and 2% tie density modified voxelwise analysis (surface). Middle, published ROIs (cingulo-opercular task control [Dosenbach et al., 2007]; salience [Seeley et al., 2007]) or modified voxelwise subgraphs, with an overlaid heat map of on-cue meta-analysis activation. On-cue activity localizes to the purple subgraph.

Bottom: very strong fc-Mapping gradients are displayed separating the black and purple subgraphs, indicating that they possess distinct rs-fcMRI signals. Right: at top, three unknown subgraphs from the 0.5% tie density modified voxelwise analyses are shown. The salmon subgraph (gray in all other figures, here salmon for contrast) is reproduced with a 2% areal subgraph overlaid as spheres, and the strongest activations from the memory retrieval meta-analysis are shown below. The light blue subgraph is also reproduced and the coordinates of a putative functional system from Nelson et al. (2010a) are overlaid as tan spheres.

cortex: a cingulo-opercular control system first described by Dosenbach et al. (2006) as the "core" of a task performance system, which is thought to instantiate and maintain set during task performance, and the salience system of Seeley et al. (2007). Relative to the black subgraph, the purple subgraph lies anterior and ventral in aCC, lateral in aPFC, and dorsal in the al. Three pieces of data hint at the identities of these subgraphs. First, the coordinates reported for the task control network are dorsal to salience coordinates in the insula (Dosenbach et al., 2007; Seeley et al., 2007), although most other coordinates do not distinguish the competing functional systems. Second, on-cue activity localizes to the purple subgraph in the al, aCC, and aPFC (the task control system was defined over a range of tasks by on-cue activity entering a task block, sustained activity during a task block, and error-related activity). Finally, the fc-Mapping technique detects a strong border between the black and purple subgraphs at many locations, indicating that rs-fcMRI signal differs strongly between these subgraphs, consistent with prior reports (Nelson et al., 2010b). We suggest that the purple subgraph more closely represents the cingulo-opercular task control system, whereas the black

670 Neuron 72, 665–678, November 17, 2011 ©2011 Elsevier Inc.

subgraph more likely relates to a salience system, though the evidence for such assignments is provisional.

At least three distributed subgraphs with previously unknown functional identities are also found (Figure 4, right). The first subgraph (salmon in Figure 4, gray in Figure 1) includes parts of posterior cingulate, posterior medial parietal, and lateral parietal cortex. We are unaware of any earlier characterizations of this collection of brain regions as a coherent functional system, but we found that these regions display the strongest activation in our memory retrieval meta-analysis. Another distributed subgraph (light blue) is found in frontal, parietal, and temporal cortex at higher thresholds of the modified voxelwise analysis. This set of regions is not a commonly described functional system, but recent work (fMRI and rs-fcMRI) (Nelson et al., 2010a) has indicated that a very similar set of regions (tan spheres in Figure 4) interposed between fronto-parietal and default regions may be a functional system, also implicated in memory retrieval. Another novel subgraph is shown in plum, with representation in fusiform cortex, the precuneus, lateral and medial posterior parietal cortex, and superior frontal cortex.



Figure 5. The "Task Positive System" Consists of Multiple Subgraphs, Including Dorsal Attention, Fronto-Parietal Task Control, and Cingulo-Opercular Task Control Systems

At left, the "task+ system" of Fox et al. (2005). At right, three subgraphs from the 0.5% tie density modified voxelwise analysis. The task+ system is composed of at least three subgraphs, corresponding to the fronto-parietal task control, cingulo-opercular task control, and dorsal attention systems.

We now shift from examining individual subgraphs to collections of subgraphs and their relationships to one another.

The "Task-Positive System" Is Composed of Multiple Subgraphs Whereas the "Task-Negative System" Is Composed of a Single Subgraph

In an influential article, Fox et al. (2005) described a task-positive network that is broadly activated across tasks, and a task-negative network that is broadly inactivated across tasks (Figure 5). Seed timecourses demonstrated that rs-fcMRI signal in one network tended to rise as the signal in the other network fell, and the authors used seed correlation maps to suggest that large portions of the brain are organized into two anticorrelated networks. This framework is a useful heuristic, but the present results suggest a more complicated picture.

The "task-negative system" corresponds predominantly to a single subgraph (the default mode system), with possible additional correspondence to the memory retrieval (salmon) subgraph described above. The "task-positive system" is, from a graph theoretic perspective, composed of at least three major subgraphs: the dorsal attention system (green), the fronto-parietal task control system (yellow), and the cinguloopercular task control system (purple). Because subgraphs are formed of nodes that are more related to one another than to the rest of the network, the rs-fcMRI timecourses of these subgraphs must be distinct from one another.

This highlights a fundamental difference between "resting state networks" defined by seed map analyses and the subgraphs defined by graph-based approaches. Seed maps measure only the relationships between a seed ROI and other brain regions (usually voxels), whereas a graph of N nodes integrates the information of N seed maps to capture not only the relationships of a seed region to other brain regions, but also the second-order relationships among those other brain regions. In other words, seed maps measure relationships in isolation, whereas graphs capture these relationships and their context. There is no necessary conflict in saying that a seed from dorsal attention regions highlights broad swaths of cortex (the seed's voxelwise neighbors) and that graph-based analyses indicate that some of these neighbors belong to other discrete subgraphs. Thus, the "task-positive system" seems to be composed of at least three subgraphs, corresponding to distinct attentional and task control systems.

The Default Mode System Is, from a Graph Theoretic Perspective, Like Sensory and Motor Systems

Classic models of cognitive control posit that sensory information is received, processed according to the demands of a task, and an output is generated (Norman and Shallice, 1986). Processing at the input and output stages is thought to be relatively modular (not strictly in the graph theoretic sense), whereas cognitive control mechanisms must flexibly adapt processing to a wide range of task sets (Posner and Petersen, 1990). On such an account, within a graph theoretic context, subgraphs thought to be responsible for task set or "control" ought to maintain a relatively diverse set of relationships, whereas sensory or motor "processing" systems ought to have relatively compartmentalized sets of relationships.

The compartmentalization and diversity of relationships in graphs can be measured by two related, standard graph measures: the local efficiency and participation coefficients of nodes. Local efficiency is a measure of integration among the neighbors of a node (the nodes a node has ties with): high local efficiency means that a node is embedded within a richly connected environment, and low local efficiency means that the neighbors of the target node are sparsely connected to one another. The participation coefficient measures the extent to which a node connects to subgraphs other than its own. Low participation coefficients indicate that nodes are confined to interactions within their own subgraphs, whereas higher coefficients indicate that nodes connect to a variety of subgraphs. Figure 6 plots subgraphs, local efficiency, and participation coefficients for the areal graph over a range of thresholds. "Processing" systems ought to have high local efficiency and low participation coefficients, reflected as hot colors in the middle panel and cool colors in the right panel of Figure 6. The visual (blue) and hand SSM (cyan) subgraphs meet this prediction, as expected, and, intriguingly, so does the default mode system (red). The more diverse relationships of "control" systems, on the other hand, ought to be reflected in lower local efficiencies and higher participation coefficients, seen as cooler colors in the middle panel and warmer colors in the right panel. In comparison to "processing" systems, the fronto-parietal task control (yellow) subgraph has significantly lower local efficiency and higher participation indices, as one would expect. ANOVA and t tests confirm that these findings hold over a range of thresholds (see Figure 6).

These findings have several implications. Viewed from a graph theoretic perspective, sensory and motor systems and the default mode system have similar levels of self-integration and self-containment. From the cognitive control perspective



Figure 6. Default, Visual, and Somatosensory-Motor Systems Are Well-Integrated on Local Scales but Are Relatively Isolated in Relation to Other Functional Systems

At top, the subgraphs, local efficiencies, and participation coefficients for all nodes in the areal network over a range of thresholds are shown. The local efficiency of each node indicates the extent to which a node is embedded in a richly connected local environment. High (hot color) values indicate a richly connected local environment. The participation coefficient of each node indicates the extent to which a node has ties to other subgraphs. Here, low (cool color) values indicate that nodes are connected almost exclusively to members of their own subgraph. One-factor ANOVAs indicate a significant effect of subgraph at all thresholds for both indices (all with $p < 10^{-6}$), and post hoc t tests indicate that the cyan, blue, and red subgraphs have significantly higher local efficiencies and lower participation coefficients at most or all thresholds than the yellow subgraph. Node assignments for a single threshold (4% tie density) are shown on a brain and in a spring embedded layout, and the local efficiencies and participation coefficients of relevant subgraphs at this threshold are shown. Note that local efficiency is independent of subgraph assignment, whereas participation coefficients depend upon subgraph assignment.

outlined above, these similarities would suggest that the default mode system acts more as a "processing system" than a "control system" (in contrast with the fronto-parietal system). Viewed from a perspective of temporal dynamics, the high similarity of node relationships within SSM and visual systems and the default mode system might indicate that these systems in particular are relatively stationary, whereas other subgraphs such as task control systems might have more dynamic sets of relationships. It should also be noted that several studies (Buckner et al., 2009; Cole et al., 2010) have implicated the default mode system as the seat of the most prominent "hubs" in rs-fcMRI brain graphs. Although default mode nodes may indeed have many ties, the isolated nature of the default mode subgraph recasts the meaning of these nodes as hubs in the context of brain-wide rs-fcMRI networks.

Functional Systems Are Arranged in Topological Motifs across the Cortex

One of the more striking features of the modified voxelwise analysis is that subgraphs appear to be arranged in spatial motifs throughout the cortex. Figure 7 demonstrates the presence of motifs at a single threshold of the modified voxelwise analysis. For each subgraph, the distribution of its spatial interfaces (defined as en face voxels) with other subgraphs is plotted,



Figure 7. Functional Systems Are Arranged into Topological Motifs across the Cortex

In charts, particular subgraphs at a single threshold are selected, the spatial boundaries of that subgraph are found, and the distribution of spatial interfaces (en face voxels) to other subgraphs are calculated. The most frequent interfaces are plotted as percents of the total subgraph interface volume. Motifs are inferred by finding instances where subgraphs interfacing with a subgraph are themselves very unlikely to interface. For instance, in the top chart, the light blue subgraph interfaces most frequently with the yellow and red subgraphs, but red is only 3.6% of yellow's interface, and yellow is only 2.6% of red's interface. Below each chart, plots of relevant subgraphs on brain surfaces visually demonstrate the repeated spatial patterns of subgraphs. Data from the modified voxelwise analysis at 1% tie density in the replication cohort are presented.

and then these neighboring subgraphs are examined to see whether they are themselves unlikely to interface (implying a 3-step motif). For example, the light blue subgraph interfaces predominantly with red and yellow subgraphs, which are themselves miniscule portions of each others' borders (red is 3.5% of yellow's border, and yellow is 2.6% of red's border), implying a yellow-light blue-red motif. Plots of relevant subgraphs on brain surfaces visually confirm the presence of motifs. Three instances of this motif are demonstrated, for the light blue, black (salience), and green (dorsal attention) subgraphs. Other 3-step motifs are present but not shown (e.g., red-teal-purple), and these motifs can be found up and down subgraph hierarchies (i.e., thresholds).

A principal concern about such spatial motifs is that they are artifactual—that they arise as intermediate mixtures of adjacent signals, particularly when averaging over subjects. Although these concerns cannot be entirely excluded, several interposed subgraphs (e.g., the green dorsal attention system or the teal ventral attention system) have firm and extensive experimental bases. If these are not considered artifactual, then other subgraphs deserve similar consideration.

DISCUSSION

Task-Free Approaches Delineate Functional Systems across the Cortex

At the onset of functional neuroimaging some 25 years ago, investigators made educated guesses about the types of operations that the human brain must perform, and designed experimental paradigms to elicit such operations (Lueck et al., 1989; Pardo et al., 1991; Petersen et al., 1988; Posner et al., 1988). Over time, evidence accumulated implicating collections of brain regions that were assumed to share the burden of some set of cognitive operations, defining functional systems (Corbetta and Shulman, 2002; Dosenbach et al., 2006; Raichle et al., 2001). Until the study of spontaneous BOLD activity, however, the association of regions within a functional system was to some extent dependent upon sets of task paradigms. Task-based approaches left functional systems open to an interpretation that rather than being a fundamentally related group of brain regions within a brain-wide context, a functional system thus defined might be just a transient and task-specific association of brain regions.

The subgraphs presented herein were derived in task-free data using methods with no prior information about node identity. There is substantial agreement between aspects of paradigm-driven functional system definition in neuroimaging, and paradigm-free subgraphs derived in task-free activity. Even if one were to object that the areal network included functional assumptions via meta-analytic localizers, the modified voxelwise analysis, which returned very similar results, made no such assumptions. In a brain-wide context, several functional systems are distinguished from each other by spontaneous activity.

This task-free definition of brain functional organization can inform perspectives on cognitive function. For example, dorsal and lateral frontal cortex appears to be apportioned among a variety of distributed subgraphs, many of which correspond to functional systems with known characteristics (Figure 2). This organization does not appear consistent with accounts of cognition that posit rostro-caudal gradients or hierarchies across frontal cortex (Badre and D'Esposito, 2009; O'Reilly, 2010). In a related manner, the finding of similar graph properties (relatively dense internal relationships and relatively few external relationships) in visual, SSM, and default mode systems may inform the degree to which the default mode system is seen as a processing type of system versus a control type of system. Such a finding need not contradict the description of posterior members of the default mode system as cortical hubs (Buckner et al., 2009), but it may alter the understanding of what it means to be a hub.

Integrating the Present Findings with Other Approaches to Whole-Brain rs-fcMRI Analysis

Recent investigations into the structure of functional brain organization using a variety of methods (Erhardt et al., 2010; Yeo et al., 2011) have found some similar (but not identical) sets of resting state networks as the subgraphs reported here. We consider convergence across methods to be a key indicator of the validity of findings. We find the graph theoretic framework to be especially useful, because it is capable of describing the overall graph (no such measures are presented in this article, but small-world measures are an example), portions of the system (e.g., subgraphs), or individual nodes of the system (e.g., local efficiency) within a common framework.

Our findings have substantial implications for past and future graph-based analyses. By examining multiple network definitions within a single data set, we were able to show how network definition profoundly affects the properties of a network, and therefore the conclusions one would draw about the brain. Our results demonstrate drawbacks in some previous approaches, while offering new approaches that appear to more plausibly represent brain organization.

It is important to recognize that these new approaches to graph definition are not equivalent or interchangeable. Note that in this article we examine several graph theoretic properties of the areal graph, but restrict our discussions of modified voxelwise data to spatial observations. The areal graph is formed using our best estimates of the functional "units" in the brain, and many properties of this network should be fairly direct reflections of functional brain organization. On the other hand, the modified voxelwise graph is defined using volumetric elements (voxels), and this graph reflects volumetric properties of functional organization. In this graph, most functional areas are probably represented by many voxels, and large functional areas (and functional systems) will dominate the graph structure regardless of their roles in information processing relative to smaller areas or systems. This volume-based definition thus warps representations of information processing, limiting the conclusions that can be drawn from this graph.

Directions for Future Work

The analyses presented here suggest several avenues for future inquiry. Within graphs that possess many subgraphs with strong correspondence to functional systems, we have detected additional subgraphs with no such identity but with hints of shared activity in certain contexts (e.g., memory retrieval activity in the salmon and light blue subgraphs). Unifying functional attributes among these subgraphs should be sought and tested. Our results demonstrate strong within-subgraph connectivity in sensory, motor and default mode systems, especially in contrast to task control systems, suggesting that these systems may differ in the dynamics of their relationships with other subgraphs over time. Our analyses only examined static pictures of graphs obtained by summarizing activity over entire epochs into a single correlation coefficient, and future work should explore if and how these relationships change over time. Perhaps the most obvious avenue for future work will lie in the comparison of graphs across the lifespan and in disease. A recognized limitation within graph theoretic investigations of structural and functional brain networks is the current lack of validated parcellation strategies (see Fornito et al., 2010; Wig et al., 2011; Zalesky et al., 2010) for comprehensive discussions). We have derived and presented a graph of 264 putative functional areas that displays a plausible functional structure that should be sensitive to the organization of many functional systems. If the locations of functional areas do not greatly differ across populations (Barnes et al., 2011), this graph should be applicable to a wide variety of populations, such as clinical or developmental cohorts.

Limitations

The present study should be considered a preliminary draft of functional brain networks and has many limitations. The methods of locating putative functional areas may certainly have overlooked, misplaced, or fabricated some areas. Additionally, the spherical ROIs used to model functional areas do not reflect the true shapes of functional areas. However, because subgraph structures in areal and modified voxelwise networks were remarkably alike, this does not seem to have crippled the endeavor. This study used a single signal (BOLD) with known susceptibility artifacts in temporal and orbitofrontal cortex. Accordingly, much remains to be discovered about the organization of the ventral surface of the brain, as well as subcortical and cerebellar organization (see Buckner et al., 2011). One additional limitation inherent to fMRI is resolution: voxels are 3 mm on each side, and partial voluming as well as the smoothing inherent in data processing limit the resolution that these studies can achieve. To offset these undesired effects, short-distance relationships were eliminated from areal and modified voxelwise analyses, and single subjects were examined. Future efforts that refine rs-fcMRI techniques and integrate findings from other modalities, such as structural imaging, EEG, or MEG, will provide valuable additions and refinements to our observations, both in terms of identifying the functional "units" of the human brain and in more completely modeling functional brain networks in space and time.

CONCLUSIONS

We close with two broad points. First, there is a growing trend to examine healthy and pathological brain activity in terms of networks (Bullmore and Sporns, 2009; Church et al., 2009; Seeley et al., 2009). The sensitivity and specificity of such analyses is directly linked to the comprehensiveness and accuracy of the framework used to examine brain networks. The framework used in this report appears to be reasonably accurate, and is capable of describing networks as a whole, as subgraphs, or as individual nodes, making it a powerful tool for examining functional relationships in the human brain. Second, the accuracy of connectivity analyses depends upon the isolation of relevant or unique signals. As the areal and modified voxelwise analyses demonstrate, the human cortex possesses a complex and dense topography of functional systems, underscoring the need for "tedious anatomy" in neuroimaging studies (Devlin and Poldrack, 2007).

EXPERIMENTAL PROCEDURES

Subjects

Healthy young adults were recruited from the Washington University campus and the surrounding community. All subjects were native English speakers and right-handed. All subjects gave informed consent and were compensated for their participation.

Data Sets and Data Collection

This study utilized multiple data sets. The first and second data sets were used for meta-analytic and fc-mapping analyses, respectively. The third data set was used for rs-fcMRI network analysis. The first (n > 300, detailed in Table S1) and second data sets (n = 40) were acquired on a Siemens 1.5 Tesla MAGNETOM Vision MRI scanner (Erlangen, Germany) as described in Dosenbach et al. (2010). The third data set (n = 106: a 53 subject cohort, 52 subject cohort, and an additional single subject) was acquired on a Siemens MAGNETOM Tim Trio 3.0T Scanner with a Siemens 12 channel Head Matrix Coil (Erlangen, Germany) as described in Dosenbach et al. (2010). See Supplemental Experimental Procedures for acquisition details.

Data Processing

Functional images underwent standard fMRI preprocessing to reduce artifacts, register subjects to a target atlas, and resample the data on a 3 mm isotropic grid (Shulman et al., 2010). See Supplemental Experimental Procedures for further details.

rs-fcMRI Processing

For rs-fcMRI analyses, several additional preprocessing steps were utilized to reduce spurious variance unlikely to reflect neuronal activity (Fox et al., 2009). These steps included: (1), a temporal band-pass filter (0.009 Hz < f < 0.08 Hz) and spatial smoothing (6 mm full width at half maximum); (2), regression of six parameters obtained by rigid body head motion correction; (3), regression of the whole brain signal averaged across the whole brain; (4), regression of white matter signal averaged from white matter ROIs; and (5), regression of white matter signal averaged from white matter ROIs. The first derivatives of these regressors were also regressed.

Meta-Analytic ROI Definition

The first method of identifying putative functional areas searched a large fMRI data set acquired in a single scanner (data set 1) for brain regions that reliably displayed significant activity when certain tasks were performed (e.g., button-pressing) or certain signal types (e.g., error-related activity) were expected (see Table S1). Meta-analyses identified 322 ROIs (10 mm diameter spheres, see Figure S1), which were reduced to a final collection of 151 nonoverlapping meta-analytic ROIs. Full details of meta-analyses are available in Supplemental Experimental Procedures.

fc-Mapping ROI Definition

fc-Mapping techniques were applied to eyes-open fixation rs-fcMRI data from 40 healthy young adults (data set 2: 27 M/13 F, average age = 26.4 years old, average RMS movement = 0.42 mm, average number of volumes = 432). See Cohen et al. (2008) and Nelson et al. (2010a) for full conceptual and technical descriptions of fc-Mapping on cortical patches. Here, patches extending over the entire cortical surface (one per hemisphere) were used to define putative functional areas. This technique generated 254 ROIs across the cortex, which were reduced to a final set of 193 nonoverlapping ROIs. See Supplemental Experimental Procedures for further details.

Areal ROI Set Formation

Meta-analytic ROIs and fc-Mapping ROIs were merged to form a maximallyspanning collection of ROIs. Meta-analytic ROIs were given preference, and nonoverlapping fc-Mapping ROIs were then added, resulting in 264 independent ROIs.

Parcel-Based, Voxel-Based, and Modified Voxelwise Network Formation

A 90-node parcel-based network was formed by using the 90-parcel automated anatomical labeling (AAL) atlas (Tzourio-Mazoyer et al., 2002) to assign all voxels (n = 44,100) within the atlas into 90 parcels. An average timecourse was formed for each parcel by averaging the timecourses of all nodes within the parcel. A 44,100-node voxelwise network was defined from all voxels within the AAL atlas (Tzourio-Mazoyer et al., 2002). The modified voxelwise networks arose by masking out ties that terminated within 20 mm of the source voxel. Distances of 15–25 mm were tested, with similar results across networks. Analyses were performed on all voxels in both hemispheres (n = 44,100), and also on all voxels within a single hemisphere (n = 22,050). Single hemisphere analyses were much less computationally demanding, permitting a wider range of analysis), and results between single- and dual-hemisphere analyses were similar. All figures except Figure 3 (both hemispheres were used for consistency with the voxelwise analysis and the rest of the literature in this figure) in the article portray single-hemisphere analyses.

Formation of Two Subject Cohorts for rs-fcMRI Network Analysis

rs-fcMRI networks were studied in continuous eyes-open fixation data from two cohorts (data set 3) of healthy young adults, matched for age, sex, movement and number of volumes in scans, as shown in Table S1. These subjects underwent a rigorous quality control process to correct for subject motion (Power et al., 2011). See Supplemental Experimental Procedures for details. Reported numbers of volumes (time frames of rs-fcMRI data) and RMS are for the final, usable, data (Table S1). Data cleaning for subject movement during the scan removed 6% of the data from subjects (range 4%–8%), and each cohort contained a mean of 350 frames of data per subject (range 215–501 frames). The single subject in Figure 2 had 1181 frames of data.

rs-fcMRI Graph Formation

Given a collection of N ROIs (parcels, voxels, or putative areas), within each subject, timecourses are extracted for all ROIs and an N \times N correlation matrix is calculated. An average matrix is formed across all subjects in a cohort, and the diagonal is set to zero. This defines a weighted graph.

Typical graph analyses of weighted networks ignore negative ties and are obliged to explore a range of thresholds to characterize the properties of a network (Power et al., 2010; Rubinov and Sporns, 2010). Recent proposals to incorporate negative weights into analyses of subgraph detection have been made (Rubinov and Sporns, 2011; Traag and Bruggeman, 2009), but here we follow the traditional approach. Many real-world networks have tie densities of a few percent or less (Newman, 2010), and the graph analytic techniques utilized here were developed upon such networks (Fortunato, 2010; Newman, 2010; Rosvall and Bergstrom, 2008). Accordingly, the analyses presented here typically span a threshold range on the order of 10% down to 1% tie density though the precise range depends upon the network (for example, the AAL-based parcel network becomes severely fragmented below 4% tie density and we do not present results from such thresholds). In general, results are presented over a range of thresholds to give the reader a sense of the dependence of a property upon thresholds, and no formal definition of threshold ranges is proposed, because it is essentially arbitrary.

As noted in the text, short-range correlations can arise from shared patterns of local neuronal activity, but they can also arise from aspects of data processing (e.g., reslicing, blurring), as well as motion-induced artifacts (Power et al., 2011). Local correlations are thus combinations of neurobiological and artifactual signal. To minimize the effects of questionable correlations on network structure, ties terminating within 20 mm of the source ROI are set to zero in all areal network analyses and in the modified voxelwise analysis. Although this process does not completely remove the effect of reslicing and blurring on correlations in the data (consider a voxel's correlations of questionable origin. This procedure eliminated 635 (4.1%) of the 15,375 positive ties in the areal network, and 15.3 million (4.2%) of 470 million ties in the single hemisphere voxelwise network.

Subgraph Detection and Graph Analysis

For a given network at a given threshold, the correlations below the threshold were set to zero, and the resulting matrix was subjected to subgraph detection algorithms. We utilized the Infomap algorithm, one of the best-performing algorithms on multiple benchmark networks (Fortunato, 2010; Lancichinetti and Fortunato, 2009). Other algorithms were tried, with similar results. Subgraph assignments were returned as numbers, which were then mapped onto nodes and ROIs as colors.

Local efficiency was calculated after (Latora and Marchiori, 2001). Participation coefficients were calculated after (Guimerà et al., 2005). Binary networks were used for calculations.

Computations and Visualizations

MRI images were processed using in-house software. Network calculations were performed using MATLAB (The MathWorks, Natick, MA). The Infomap algorithm was provided by Rosvall and Bergstrom (2008). Network visualizations were created using the Social Network Image Animator (SoNIA) software package (Bender-deMoll and McFarland, 2006). Brain surface visualizations were created using Caret software and the PALS surface (Van Essen, 2005; Van Essen et al., 2001).

SUPPLEMENTAL INFORMATION

Supplemental Information includes five figures, two tables, and Supplemental Experimental Procedures and can be found with this article online at doi:10.1016/j.neuron.2011.09.006.

ACKNOWLEDGMENTS

We thank Nico Dosenbach, Thomas Pearce, Bradley Miller, and our reviewers for their attentive reading of this manuscript. We thank Olaf Sporns and Mika Rubinov for technical help with graph analysis, and Joe Dubis for help with meta-analyses. This work was supported by NIH R21NS061144 (S.P.), NIH R01NS32979 (S.P.), a McDonnell Foundation Collaborative Action Award (S.P.), NIH R01HD057076 (B.L.S.), NIH F30NS062489 (A.L.C.), NIH U54MH091657 (David Van Essen), and NSF IGERT DGE-0548890 (Kurt Thoroughman).

Accepted: September 2, 2011 Published: November 16, 2011

REFERENCES

Badre, D., and D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? Nat. Rev. Neurosci. *10*, 659–669.

Barnes, K.A., Nelson, S.M., Cohen, A.L., Power, J.D., Coalson, R.S., Miezin, F.M., Vogel, A.C., Dubis, J.W., Church, J.A., Petersen, S.E., and Schlaggar, B.L. (2011). Parcellation in left lateral parietal cortex is similar in adults and children. Cereb. Cortex., in press. Published online August 1, 2011. 10.1093/cer-cor/bhr189.

Bender-deMoll, S., and McFarland, D.A. (2006). The art and science of dynamic network visualization. J. Soc. Struct. 7.

Biswal, B., Yetkin, F.Z., Haughton, V.M., and Hyde, J.S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. *34*, 537–541.

Buckner, R.L., Sepulcre, J., Talukdar, T., Krienen, F.M., Liu, H., Hedden, T., Andrews-Hanna, J.R., Sperling, R.A., and Johnson, K.A. (2009). Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to Alzheimer's disease. J. Neurosci. *29*, 1860–1873.

Buckner, R.L., Krienen, F.M., Castellanos, A., Diaz, J.C., and Yeo, B.T.T. (2011). The organization of the human cerebellum estimated by intrinsic functional connectivity. J. Neurophysiol., in press. Published online July 27, 2011. 10.1152/jn.00339.2011.

Bullmore, E., and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. Nat. Rev. Neurosci. *10*, 186–198. Burton, H., Sinclair, R.J., Wingert, J.R., and Dierker, D.L. (2008). Multiple parietal operculum subdivisions in humans: tactile activation maps. Somatosens. Mot. Res. *25*, 149–162.

Butts, C.T. (2009). Revisiting the foundations of network analysis. Science 325, 414–416.

Church, J.A., Fair, D.A., Dosenbach, N.U., Cohen, A.L., Miezin, F.M., Petersen, S.E., and Schlaggar, B.L. (2009). Control networks in paediatric Tourette syndrome show immature and anomalous patterns of functional connectivity. Brain *132*, 225–238.

Cohen, A.L., Fair, D.A., Dosenbach, N.U.F., Miezin, F.M., Dierker, D., Van Essen, D.C., Schlaggar, B.L., and Petersen, S.E. (2008). Defining functional areas in individual human brains using resting functional connectivity MRI. Neuroimage *41*, 45–57.

Cole, M.W., Pathak, S., and Schneider, W. (2010). Identifying the brain's most globally connected regions. Neuroimage *49*, 3132–3148.

Corbetta, M., and Shulman, G.L. (2002). Control of goal-directed and stimulusdriven attention in the brain. Nat. Rev. Neurosci. 3, 201–215.

Corbetta, M., Shulman, G.L., Miezin, F.M., and Petersen, S.E. (1995). Superior parietal cortex activation during spatial attention shifts and visual feature conjunction. Science *270*, 802–805.

Corbetta, M., Patel, G., and Shulman, G.L. (2008). The reorienting system of the human brain: from environment to theory of mind. Neuron 58, 306–324.

Deco, G., Jirsa, V.K., and McIntosh, A.R. (2011). Emerging concepts for the dynamical organization of resting-state activity in the brain. Nat. Rev. Neurosci. *12*, 43–56.

Devlin, J.T., and Poldrack, R.A. (2007). In praise of tedious anatomy. Neuroimage 37, 1033–1041, discussion 1050–1058.

Dosenbach, N.U.F., Visscher, K.M., Palmer, E.D., Miezin, F.M., Wenger, K.K., Kang, H.C., Burgund, E.D., Grimes, A.L., Schlaggar, B.L., and Petersen, S.E. (2006). A core system for the implementation of task sets. Neuron *50*, 799–812.

Dosenbach, N.U.F., Fair, D.A., Miezin, F.M., Cohen, A.L., Wenger, K.K., Dosenbach, R.A.T., Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., et al. (2007). Distinct brain networks for adaptive and stable task control in humans. Proc. Natl. Acad. Sci. USA *104*, 11073–11078.

Dosenbach, N.U., Nardos, B., Cohen, A.L., Fair, D.A., Power, J.D., Church, J.A., Nelson, S.M., Wig, G.S., Vogel, A.C., Lessov-Schlaggar, C.N., et al. (2010). Prediction of individual brain maturity using fMRI. Science *329*, 1358–1361.

Erhardt, E.B., Rachakonda, S., Bedrick, E.J., Allen, E.A., Adali, T.I., and Calhoun, V.D. (2010). Comparison of multi-subject ICA methods for analysis of fMRI data. Hum. Brain Mapp., in press. Published online December 15, 2010. 10.1002/hbm.21170.

Fornito, A., Zalesky, A., and Bullmore, E.T. (2010). Network scaling effects in graph analytic studies of human resting-state FMRI data. Front. Syst. Neurosci. *4*, 22.

Fortunato, S. (2010). Community detection in graphs. Phys. Rep. 486, 75-174.

Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., and Raichle, M.E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proc. Natl. Acad. Sci. USA *102*, 9673–9678.

Fox, M.D., Corbetta, M., Snyder, A.Z., Vincent, J.L., and Raichle, M.E. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. Proc. Natl. Acad. Sci. USA *103*, 10046–10051.

Fox, M.D., Zhang, D., Snyder, A.Z., and Raichle, M.E. (2009). The global signal and observed anticorrelated resting state brain networks. J. Neurophysiol. *101*, 3270–3283.

Fransson, P., Aden, U., Blennow, M., and Lagercrantz, H. (2011). The functional architecture of the infant brain as revealed by resting-state fMRI. Cereb. Cortex *21*, 145–154.

Greicius, M.D., Krasnow, B., Reiss, A.L., and Menon, V. (2003). Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proc. Natl. Acad. Sci. USA *100*, 253–258.

Guimerà, R., Mossa, S., Turtschi, A., and Amaral, L.A.N. (2005). The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. Proc. Natl. Acad. Sci. USA *102*, 7794–7799.

Hartman, D., Hlinka, J., Palus, M., Mantini, D., and Corbetta, M. (2011). The role of nonlinearity in computing graph-theoretical properties of resting-state functional magnetic resonance imaging brain networks. Chaos *21*, 013119.

Hayasaka, S., and Laurienti, P.J. (2010). Comparison of characteristics between region-and voxel-based network analyses in resting-state fMRI data. Neuroimage *50*, 499–508.

He, Y., Wang, J., Wang, L., Chen, Z.J., Yan, C., Yang, H., Tang, H., Zhu, C., Gong, Q., Zang, Y., and Evans, A.C. (2009). Uncovering intrinsic modular organization of spontaneous brain activity in humans. PLoS ONE *4*, e5226.

Lancichinetti, A., and Fortunato, S. (2009). Community detection algorithms: a comparative analysis. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *80*, 056117.

Latora, V., and Marchiori, M. (2001). Efficient behavior of small-world networks. Phys. Rev. Lett. 87, 198701.

Liu, Y.-Y., Slotine, J.-J., and Barabási, A.L. (2011). Controllability of complex networks. Nature 473, 167–173.

Lowe, M.J., Mock, B.J., and Sorenson, J.A. (1998). Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. Neuroimage 7, 119–132.

Lueck, C.J., Zeki, S., Friston, K.J., Deiber, M.P., Cope, P., Cunningham, V.J., Lammertsma, A.A., Kennard, C., and Frackowiak, R.S.J. (1989). The colour centre in the cerebral cortex of man. Nature *340*, 386–389.

Matyas, F., Sreenivasan, V., Marbach, F., Wacongne, C., Barsy, B., Mateo, C., Aronoff, R., and Petersen, C.C.H. (2010). Motor control by sensory cortex. Science *330*, 1240–1243.

Meunier, D., Achard, S., Morcom, A., and Bullmore, E. (2009a). Age-related changes in modular organization of human brain functional networks. Neuroimage *44*, 715–723.

Meunier, D., Lambiotte, R., Fornito, A., Ersche, K.D., and Bullmore, E.T. (2009b). Hierarchical modularity in human brain functional networks. Front. Neuroinform. *3*, 37.

Nelson, S.M., Cohen, A.L., Power, J.D., Wig, G.S., Miezin, F.M., Wheeler, M.E., Velanova, K., Donaldson, D.I., Phillips, J.S., Schlaggar, B.L., and Petersen, S.E. (2010a). A parcellation scheme for human left lateral parietal cortex. Neuron *67*, 156–170.

Nelson, S.M., Dosenbach, N.U., Cohen, A.L., Wheeler, M.E., Schlaggar, B.L., and Petersen, S.E. (2010b). Role of the anterior insula in task-level control and focal attention. Brain Struct. Funct. *214*, 669–680.

Newman, M.E.J. (2010). Networks: An Introduction (Oxford: Oxford University Press).

Norman, D.A., and Shallice, T. (1986). Attention to action: willed and automatic control of behavior. In Consciousness and Self-Regulation, R.J. Davidson, G.E. Schwartz, and D. Shapiro, eds. (New York: Plenum Press), pp. 1–18.

O'Reilly, R.C. (2010). The what and how of prefrontal cortical organization. Trends Neurosci. *33*, 355–361.

Pardo, J.V., Fox, P.T., and Raichle, M.E. (1991). Localization of a human system for sustained attention by positron emission tomography. Nature *349*, 61–64.

Penfield, W., and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. Brain *60*, 389–443.

Petersen, S.E., Fox, P.T., Posner, M.I., Mintun, M., and Raichle, M.E. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. Nature *331*, 585–589.

Posner, M.I., and Petersen, S.E. (1990). The attention system of the human brain. Annu. Rev. Neurosci. *13*, 25–42.

Posner, M.I., Petersen, S.E., Fox, P.T., and Raichle, M.E. (1988). Localization of cognitive operations in the human brain. Science *240*, 1627–1631.

Power, J.D., Fair, D.A., Schlaggar, B.L., and Petersen, S.E. (2010). The development of human functional brain networks. Neuron 67, 735–748.

Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., and Petersen, S.E. (2011). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage, in press. Published online October 14, 2011. 10.1016/j.neuroimage.2011.10.018.

Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., and Shulman, G.L. (2001). A default mode of brain function. Proc. Natl. Acad. Sci. USA 98, 676–682.

Rosvall, M., and Bergstrom, C.T. (2008). Maps of random walks on complex networks reveal community structure. Proc. Natl. Acad. Sci. USA *105*, 1118–1123.

Rubinov, M., and Sporns, O. (2010). Complex network measures of brain connectivity: uses and interpretations. Neuroimage *52*, 1059–1069.

Rubinov, M., and Sporns, O. (2011). Weight-conserving characterization of complex functional brain networks. Neuroimage *56*, 2068–2079.

Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L., and Greicius, M.D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. J. Neurosci. *27*, 2349–2356.

Seeley, W.W., Crawford, R.K., Zhou, J., Miller, B.L., and Greicius, M.D. (2009). Neurodegenerative diseases target large-scale human brain networks. Neuron 62, 42–52.

Shulman, G.L., Fiez, J.A., Corbetta, M., Buckner, R.L., Miezin, F.M., Raichle, M.E., and Petersen, S.E. (1997). Common blood flow changes across visual tasks: II. Decreases in cerebral cortex. J. Cogn. Neurosci. 9, 648–663.

Shulman, G.L., Pope, D.L.W., Astafiev, S.V., McAvoy, M.P., Snyder, A.Z., and Corbetta, M. (2010). Right hemisphere dominance during spatial selective attention and target detection occurs outside the dorsal frontoparietal network. J. Neurosci. *30*, 3640–3651.

Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., and Woolrich, M.W. (2011). Network modelling methods for FMRI. Neuroimage *54*, 875–891.

Spoormaker, V.I., Schröter, M.S., Gleiser, P.M., Andrade, K.C., Dresler, M., Wehrle, R., Sämann, P.G., and Czisch, M. (2010). Development of a large-scale functional brain network during human non-rapid eye movement sleep. J. Neurosci. *30*, 11379–11387.

Tian, L., Wang, J., Yan, C., and He, Y. (2011). Hemisphere- and gender-related differences in small-world brain networks: a resting-state functional MRI study. Neuroimage *54*, 191–202.

Tomasi, D., and Volkow, N.D. (2011). Functional connectivity hubs in the human brain. Neuroimage 57, 908–917.

Traag, V.A., and Bruggeman, J. (2009). Community detection in networks with positive and negative links. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *80*, 036115.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage *15*, 273–289.

van den Heuvel, M.P., Stam, C.J., Boersma, M., and Hulshoff Pol, H.E. (2008). Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. Neuroimage *43*, 528–539.

Van Essen, D.C. (2005). A Population-Average, Landmark- and Surface-based (PALS) atlas of human cerebral cortex. Neuroimage 28, 635–662.

Van Essen, D.C., Drury, H.A., Dickson, J., Harwell, J., Hanlon, D., and Anderson, C.H. (2001). An integrated software suite for surface-based analyses of cerebral cortex. J. Am. Med. Inform. Assoc. *8*, 443–459.

Wig, G.S., Schlaggar, B.L., and Petersen, S.E. (2011). Concepts and principles in the analysis of brain networks. Ann. N Y Acad. Sci. *1224*, 126–146.

Yeo, B.T.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zollei, L., Polimeni, J.R., et al. (2011). The organization of the human cerebral cortex estimated by functional connectivity. J. Neurophysiol. *106*, 1125–1165.

Zalesky, A., Fornito, A., Harding, I.H., Cocchi, L., Yücel, M., Pantelis, C., and Bullmore, E.T. (2010). Whole-brain anatomical networks: does the choice of nodes matter? Neuroimage *50*, 970–983.