

Decoding Subject-Driven Cognitive States with Whole-Brain Connectivity Patterns

W. R. Shirer¹, S. Ryali², E. Rykhlevskaia², V. Menon^{2,3} and M. D. Greicius^{1,3}

¹Functional Imaging in Neuropsychiatric Disorders (FIND) Lab, Department of Neurology and Neurological Sciences, ²Stanford Cognitive and Systems Neuroscience Lab (SCSNL), Department of Psychiatry and Behavioral Sciences and ³Program in Neuroscience, Stanford School of Medicine, Stanford, CA 94305, USA

Address correspondence to M. D. Greicius. Email: greicius@stanford.edu.

Decoding specific cognitive states from brain activity constitutes a major goal of neuroscience. Previous studies of brain-state classification have focused largely on decoding brief, discrete events and have required the timing of these events to be known. To date, methods for decoding more continuous and purely subject-driven cognitive states have not been available. Here, we demonstrate that free-streaming subject-driven cognitive states can be decoded using a novel whole-brain functional connectivity analysis. Ninety functional regions of interest (ROIs) were defined across 14 large-scale resting-state brain networks to generate a 3960 cell matrix reflecting whole-brain connectivity. We trained a classifier to identify specific patterns of whole-brain connectivity as subjects rested quietly, remembered the events of their day, subtracted numbers, or (silently) sang lyrics. In a leave-one-out cross-validation, the classifier identified these 4 cognitive states with 84% accuracy. More critically, the classifier achieved 85% accuracy when identifying these states in a second, independent cohort of subjects. Classification accuracy remained high with imaging runs as short as 30–60 s. At all temporal intervals assessed, the 90 functionally defined ROIs outperformed a set of 112 commonly used structural ROIs in classifying cognitive states. This approach should enable decoding a myriad of subject-driven cognitive states from brief imaging data samples.

Keywords: classification, functional connectivity, resting-state, subject-driven cognition

Introduction

The ability to decode and distinguish specific cognitive states from brain imaging data constitutes a major goal of neuroscience. Efforts in this realm of brain-state classification have included decoding basic features of motor function (Dehaene et al. 1998); selecting which specific stimulus a subject has perceived (Haxby et al. 2001; Mitchell et al. 2003; Kamitani and Tong 2005), imagined (O'Craven and Kanwisher 2000), or recalled (Polyn et al. 2005; Chadwick et al. 2010) from among several candidate stimuli and separating true from false responses (Davatzikos et al. 2005). Increasingly subtle methods have allowed classification of stimuli subconsciously perceived by subjects (Haynes and Rees 2005a); classification of novel stimuli spatially or thematically related to previously viewed stimuli (Kay et al. 2008; Walther et al. 2009); classification of attentional shifts (Esterman et al. 2009; Chiu et al. 2010); and classification of distinct subjective percepts of the same stimulus (Haynes and Rees 2005b). These standard classification approaches have yielded remarkable results in terms of classification accuracy but have some important limitations (Haynes and Rees 2006). With one exception (Richiardi et al. 2010), all previous classification studies have required some form of cognitive subtraction in

which explicit information regarding the timing of perceptual or cognitive events is required. While these studies point to the power of functional imaging in decoding cognitive events whose timing is strictly controlled by the investigator or explicitly reported by the subject, they have little bearing on the more naturalistic phenomenon of continuous, subject-driven cognitive processing.

A major obstacle to decoding subject-driven cognitive states has been the functional imaging field's reliance on cognitive subtraction (Friston 1998). In standard functional magnetic resonance imaging (fMRI) studies, cognitive subtraction experiments measure blood oxygen level-dependent (BOLD) signal changes across 2 or more states and the precise start and stop times of each state are required. By contrast, functional connectivity (FC) MRI examines BOLD signal correlations across brain regions and can be performed over single free-streaming states. This approach has been applied most commonly to resting-state fMRI data to reveal a host of brain networks, referred to here as intrinsic connectivity networks (ICNs), with distinct spatial and temporal profiles corresponding to canonical functions such as vision, hearing, language, working memory, visuospatial attention, salience processing, and episodic memory (Biswal et al. 1995; Hampson et al. 2002; Greicius et al. 2003; Kiviniemi et al. 2003; Beckmann et al. 2005; Bellec et al. 2006; Damoiseaux et al. 2006; Seeley et al. 2007; Kiviniemi et al. 2009; Smith et al. 2009). These same ICNs can be detected during states of continuous cognitive processing, both investigator-driven (Fransson 2006; Hampson et al. 2006) and subject-driven (Harrison et al. 2008).

The present study extends the reach of brain-state decoding by undertaking a classification of 4 continuous internally generated states based on their distinct FC patterns. We asked whether patterns of FC within and across ICNs could be used to distinguish between 4 subject-driven cognitive states: undirected rest, retrieval of recent episodic memories, serial subtractions, and (silent) singing of music lyrics. The imaging data were acquired in continuous 10-min runs with no stimulus presentation and no investigator-imposed temporal landmarks other than the start and end of the scan. Patterns of within- and between-ICN connectivity were used to train a classification algorithm on data from 14 subjects. The generalizability of the classification algorithm was tested with 2 methods: leave-one-out cross-validation (LOOCV) and cross-validation on an independent cohort. We then examined whether the classifier was sufficiently specific to exclude or reject a novel fifth cognitive state, spatial navigation, from the 4 states on which it was trained. The robustness of the classifier to shorter runs of continuous subject-driven cognitive states was examined. Finally, classification with the 90 functionally defined regions of interest

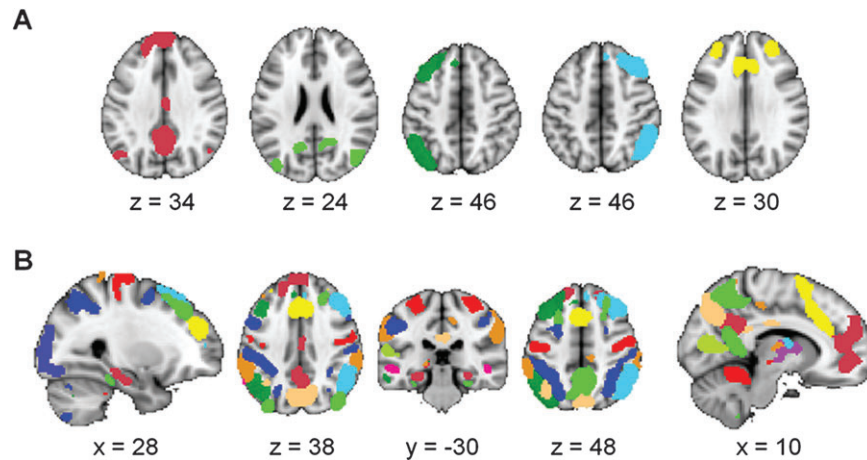


Figure 1. Functional parcellation of the brain into 90 regions of interest that cover the majority of cortical and subcortical gray matter. Group ICA applied to the resting-state data of 15 subjects yielded 14 ICNs of which 5 are shown in *A* (for all 14 ICNs, see Supplementary Fig. S2). Each ICN is thresholded to generate between 2 and 12 ROIs per ICN. When all 90 ROIs across the 14 ICNs are overlaid on a single brain image (*B*) the majority of cortical and subcortical gray matter is covered.

(ROIs) was compared across multiple duration lengths to classification with 112 commonly-used structural ROIs.

Materials and Methods

Subjects

Twenty-seven healthy right-handed subjects (17 females) aged 18–30 participated in this study. These subjects were recruited in 2 cohorts of 15 and 12 subjects separated by 5–6 months and were treated as independent cohorts for training and validating the classifier. Data from the first 15 subjects were used to define the functional (ROIs). Data from 14 of the same 15 subjects were used to train the classifier. One subject was excluded due to excessive motion (>3 mm translation and/or $>3^\circ$ rotation) during the memory scan. Data for the 12 remaining subjects in the validation cohort were collected 5 months later. Ten of these subjects' data were used to validate the classifier (2 subjects from the second cohort were excluded for falling asleep). The experimental protocol was approved by the Institutional Review Board of Stanford University.

The first cohort of subjects completed four 10-min tasks: a resting-state scan, an episodic memory task, a music lyrics task, and a subtraction task. The rest scan was always completed first, and the order of the 3 cognitive tasks was counterbalanced. For the rest task, subjects were instructed to close their eyes, let their minds wander, and try not to focus on any one thing. For the memory task, subjects were asked to recall the events of the day from when they awoke until they lay down in the scanner. For the music task, subjects were asked to sing their favorite songs in their head. For the subtraction task, subjects were asked to count backwards from 5000 by 3s. Subjects were instructed to keep their eyes closed during each of the self-driven cognitive states. The 12 subjects from the validation cohort completed these tasks plus an additional task, in which they were asked to imagine walking through all the rooms of their house or apartment (Owen et al. 2006). Debriefing of subjects confirmed that all but 2 were awake throughout the scans and were able to perform the self-driven tasks for the entire 10 min.

fMRI Acquisition

Functional images were acquired on a 3-T General Electric scanner using an 8-channel head coil. To reduce blurring and signal loss arising from field inhomogeneities, an automated high-order shimming method based on spiral acquisitions was used (Kim et al. 2002). Thirty-one axial slices (4 mm thick, 0.5 mm skip) covering the whole brain were imaged using a T_2^* -weighted gradient-echo spiral pulse sequence (time repetition = 2000 ms, time echo = 30 ms, flip angle = 80° , and 1 interleaved) (Glover and Lai 1998; Glover and Law 2001). The field of

view was 220×220 mm², and the matrix size was 64×64 , giving an in-plane spatial resolution of 3.4375 mm.

fMRI Analysis

Data were preprocessed and analyzed using SPM5 (www.fil.ion.ucl.ac.uk/spm). Images were corrected for movement using least square minimization and normalized (Friston et al. 1995) to the Montreal Neurological Institute template. Images were then resampled every 2 mm using sinc interpolation and smoothed with a 6-mm Gaussian kernel. Resampling and smoothing were done in 3 dimensions yielding a 2-mm³ resolution and effective spatial smoothness (full-width at half-maximum) of $7.2 \times 7.1 \times 8.4$ mm. The difference in the x and y dimensions reflects imprecision in the measurement as calculated by Statistical Parametric Mapping's smoothness algorithm. A high-pass filter was applied to remove low-frequency signal (<0.008 Hz) from the data. A low pass filter is often used in resting-state analyses but was excluded here to retain potentially useful information in the higher frequency bands, particularly during the cognitive tasks. To confirm our hypothesis that high-frequency data might be useful in classifying, we included an analysis using a band-pass filter which resulted in significantly reduced classification accuracy (see Supplementary Text and Supplementary Fig. S1). It is worth noting that cardiac and respiratory signals are known to cause noise in high-frequency bands. To correct for this, we measured the subjects' heart rate and respiration rate while they were being scanned. These data were used to regress the participants' physiological noise from their fMRI data (Chang and Glover 2009).

ROI Creation

We created the ROIs by applying FSL's MELODIC independent component analysis (ICA) software (<http://www.fmrib.ox.ac.uk/fsl/melodic/index.html>) to the group-level resting-state data for the first 15 subjects. Of the 30 components generated, 14 were selected visually as being ICNs based on previous reports by our group and others (Greicius et al. 2003; Fox et al. 2005; Damoiseaux et al. 2006; Seeley et al. 2007; Kiviniemi et al. 2009; Smith et al. 2009). Each of the 14 ICNs was thresholded independently and arbitrarily to generate distinct moderately sized ROIs in the cortex and subcortical gray matter ($z = 7.0 \pm 0.47$; $z_{\max} = 9$; $z_{\min} = 4.6$; voxels ≥ 25). The subcortical clusters in most ICNs are less robust and a lower threshold was used to isolate these clusters ($z = 3.8 \pm 0.40$; $z_{\max} = 5$; $z_{\min} = 2.5$; voxels ≥ 15). This thresholding step resulted in 90 ROIs across the 14 ICNs covering most of the brain's gray matter (Fig. 1 and Supplementary Fig. S2). ROI generation was done prior to classification training and so was not driven by classification results. These 90 ROIs are available for download from the corresponding author's website (<http://findlab.stanford.edu/research>).

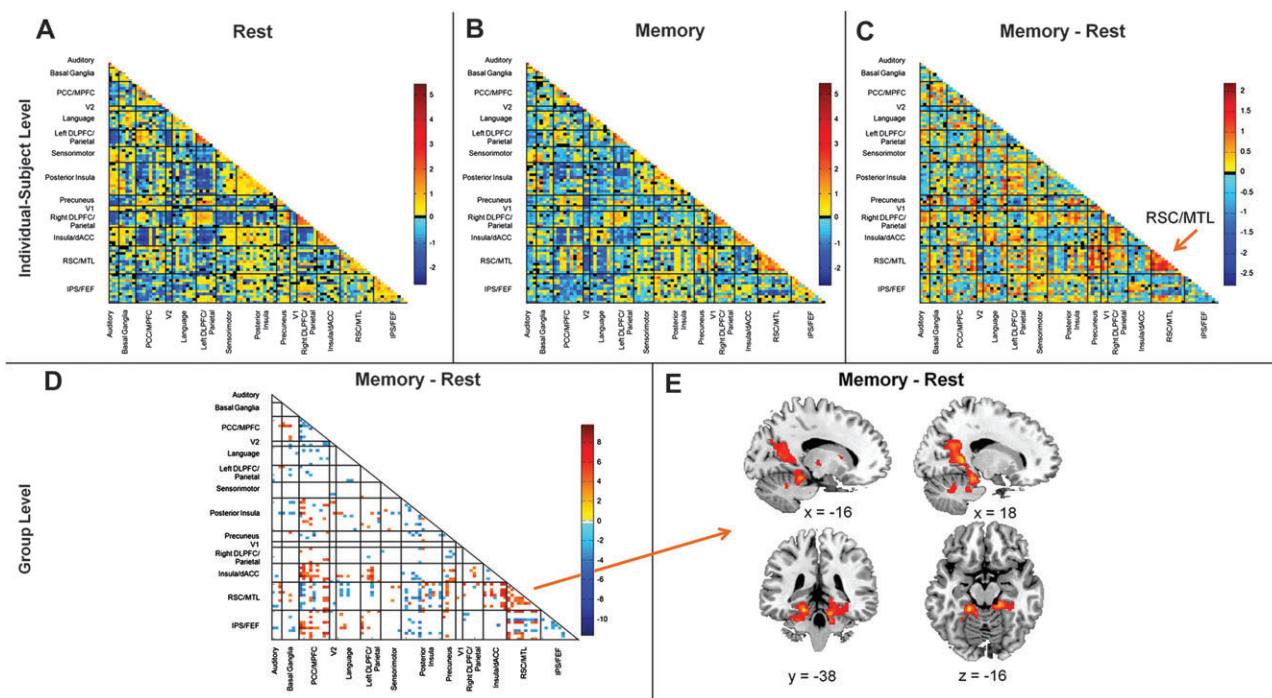


Figure 2. Subject-driven episodic memory recall drives changes in whole-brain functional connectivity. A single subject's connectivity matrix is shown for the rest scan in *A*. Cells colored in yellow-red indicate a positive pairwise correlation between 2 ROIs; green-blue cells indicate negative pairwise correlations. Coarse anatomic labels for each ICN are indicated along the *x*- and *y*-axes (more detailed anatomic information is available in Supplementary Table S1). Each ICN is bracketed by black bars and divided into 2–12 ROIs. The strong within-ICN correlations are evident along the diagonal. The same subject's memory state connectivity matrix is shown in *B*. Subtracting the rest state matrix from the memory state matrix provides the difference matrix shown in *C* where connectivity within the RSC/MTL network is shown to increase during the memory task. A paired-sample *t*-test of the state matrices across all 14 subjects (*D*) reveals changes in connectivity both between and within ICNs. These within-ICN changes (orange arrow) can also be detected by performing a paired-sample *t*-test on the individual subject ICA data (*E*). This analysis reveals clusters in the RSC/MTL network whose connectivity increases significantly during the memory scan compared with the rest scan.

Individual Subject Functional Connectivity Matrices

Fourteen subjects had usable data in the resting-state scan and the 3 additional subject-driven cognitive tasks: memory, subtraction, and music. We measured the FC between the 90 ROIs during rest and the 3 different cognitive tasks (Fig. 2). For each ROI time series, we regressed out the global mean and the confounding effects of CSF and white matter. We then calculated the Pearson correlation coefficient between the time series of all ROIs and converted these correlation coefficients to *z*-scores by applying the Fisher transformation. This resulted in a 3960 cell matrix of FC for each of the 4 cognitive states in every subject. Individual subject FC matrices were created in the same manner for the spatial navigation task in the validation cohort.

Group-Level State Matrices

We created our classification algorithm by selecting cells of interest for each of the 4 cognitive states studied in the first cohort of subjects. The classifier was not trained on the spatial navigation task. For each cognitive state, we performed a one-sample *t*-test across all subjects for each of the 3960 cells and retained cells that were significant at a false discovery rate (FDR) corrected *P* value of 0.01. Any cells that were significant for more than one cognitive state were excluded. This resulted in state-specific cells with strong positive or negative correlations that were consistent across subjects and unique to a particular cognitive state. These criteria identified 187 cells of interest for rest, 147 cells of interest for memory, 114 cells of interest for music, and 265 cells of interest for subtraction (Fig. 3). The classifier parameters were developed on the full 14-subject training data set and then validated in both a LOOCV analysis and on the independent cohort.

Classification of 4 Subject-Driven Cognitive States

We attempted to classify an individual's 4 cognitive states by deriving an overall measure of their FC within each of the 4 group-level state matrices. We tested this with 2 different cohorts of participants: the

original cohort of 14 subjects using LOOCV and the independent validation cohort of 10 subjects. On a subject-by-subject basis, each of an individual's 4 scan matrices was assigned to the group-level state matrix that it best matched based on a spatial correlation fit score. To calculate the fit of a given individual scan matrix to a specific group-level state matrix, we first examined FC within the cells of interest for the group-level state matrix and determined whether the sign of the individual cell FC agreed with the sign of the group-level cell FC. Cells whose FC sign agreed with the group-level matrix's cell sign were identified as "correct" and cells whose sign did not agree as "incorrect." We then multiplied each cell within the individual state matrix by the *z*-score in the corresponding cell of the group-level state matrix. This allowed us to weight each cell in the individual state matrix by the FC strength predicted by the group-level state matrix. We then took the sum of the absolute values for all correct cells multiplied by the proportion of correct cells and subtracted the sum of the absolute values of all incorrect cells multiplied by the proportion of incorrect cells. Because the algorithm calculates the fit score from the average connectivity in the cells of interest, the algorithm is unbiased by the number of cells in each group matrix. For each subject, each scan was assigned to that group-level state matrix for which it had the highest fit score. A binomial test was used to determine the significance of the classification accuracy. A flow chart of the classification algorithm is provided in Supplementary Figure S3.

For the LOOCV, the 4 group-level state matrices were calculated 14 different times such that a given subject's scans were compared with group-level state matrices generated from the other 13 subjects. Although LOOCV is a standard method for demonstrating classification generalizability (Mitchell et al. 2003; Mourao-Miranda et al. 2005), it is prone to cohort effects as the classifier may be overfit to a data set that is not fully representative of the broader population (Davatzikos et al. 2005; Hastie et al. 2009). Accordingly, we also applied this classification algorithm to a completely independent cohort of 10 new subjects acquired several months after the original cohort described above. For the validation

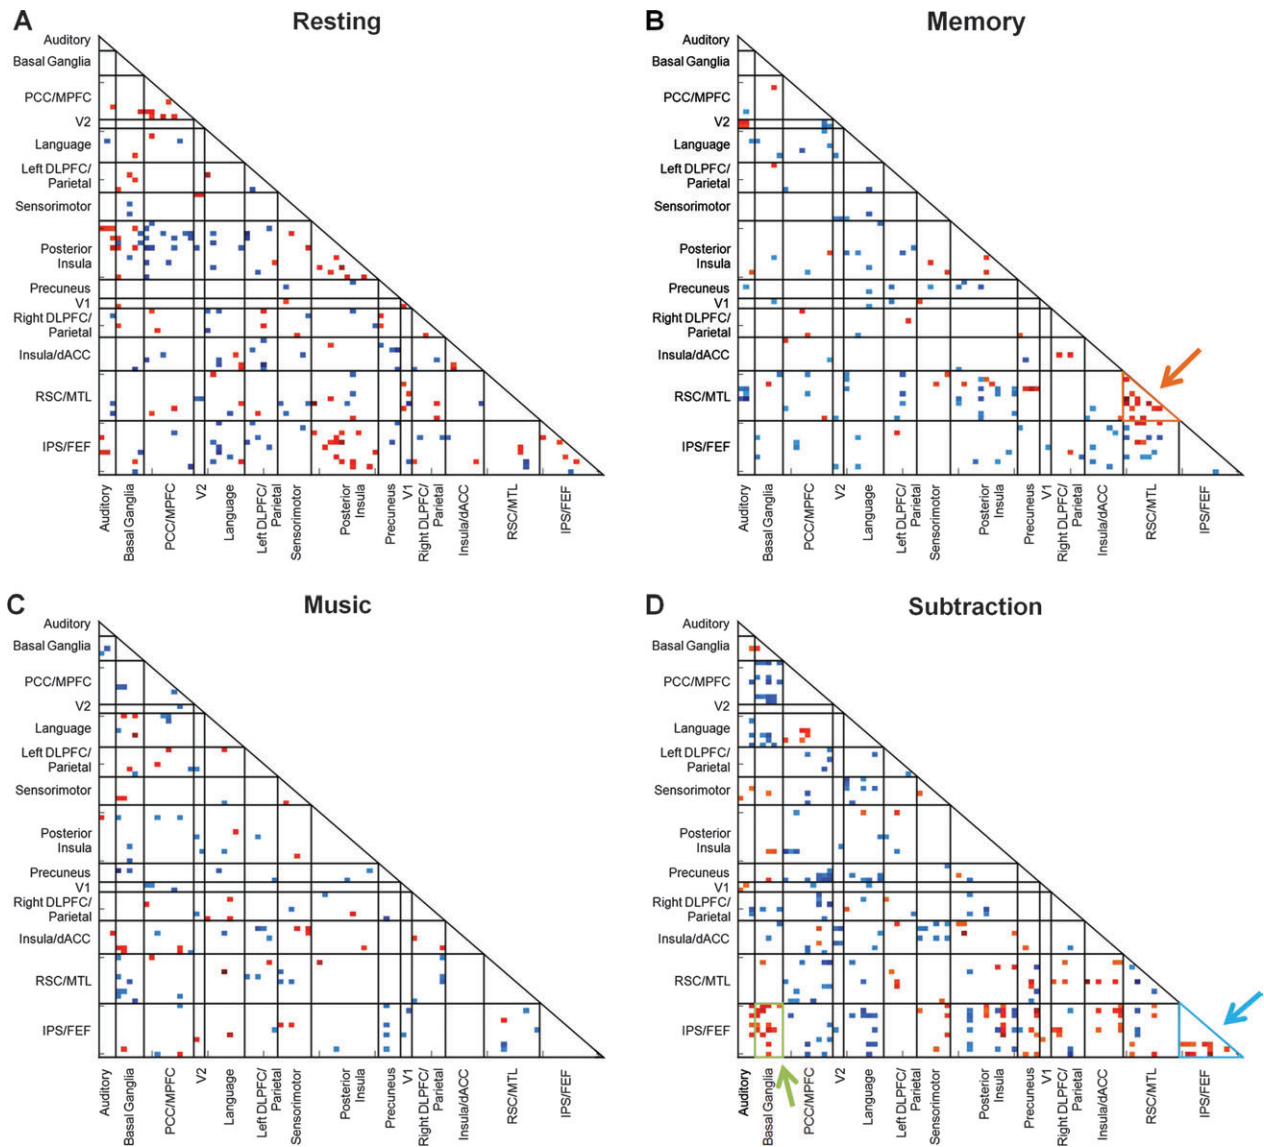


Figure 3. Distinct cross-subject patterns of whole-brain connectivity for 4 subject-driven cognitive states. For each of the 4 states, cells of interest which showed significant state-specific positive or negative correlations were included in the group-level state matrix. These state matrices are shown in *A–D*. The orange arrow in *B* indicates strong connectivity within the RSC/MTL network in the group-level memory state matrix. In the subtraction task (*D*), connectivity within the IPS/Prefrontal ICN is increased (blue arrow) but the classification algorithm also highlights increased connectivity between this ICN and the basal ganglia ICN (green arrow).

cohort classification, we used the group-level state matrices shown in Figure 3 derived from all 14 subjects in the initial cohort.

Classification Accuracy as a Function of Scan Length

To determine the influence of scan duration on classification accuracy, we repeated classification of the validation cohort at 11 increasingly shorter scan durations ranging from 10 min down to the first 30 s (Fig. 4).

Rejecting a Novel Fifth Cognitive State

The 10 subjects from the validation cohort also completed a self-driven spatial navigation task in which they were asked to imagine walking through the rooms of their home. This task was used to assess whether the classifier was sufficiently specific to exclude or reject a novel cognitive state from the 4 states on which it was trained. We calculated an individual subject spatial navigation matrix for each of the 10 subjects and included this matrix with the 4 other scan matrices in a best-fit analysis. On a subject-by-subject basis, each of the 5 individual scans was assigned a fit score to each of the 4 group-level state

matrices. In this classification analysis, given that there were only 4 group-level state matrices and 5 scans, we forced unique assignments of the individual scans to the group-level state matrices using a “winner-take-all” approach. If 2 of an individual’s scans matched to the same group-level matrix, the better match was selected and the second scan was assigned to its second-best match. The individual scan that did not fit any of the group-level state matrices better than the other individual scans was classified as the spatial navigation scan. Note that this winner-take-all algorithm is less stringent than the “best-fit” algorithm used for our main classification analyses and described above under “Classification of 4 Subject-Driven Cognitive States.” A one-sample *t*-test for the spatial navigation scan matrix is provided in Supplementary Figure S4.

Classification with Structural ROIs

The best-fit algorithm described above was implemented to create group-level state matrices for the original cohort and classify the 4 cognitive states in the validation cohort using 112 structurally defined ROIs from the Automated Anatomical Labeling Atlas (Tzourio-Mazoyer




Figure 4. Classification accuracy remains high with scans as short as 1 min. The classification algorithm was tested initially on the full 10-min scans but then on increasingly shorter scan lengths. In each case, the shorter scan lengths are taken from the beginning of the scan (i.e. 0.5 min refers to the first 30 s of the scan). Eleven different scan lengths from 30 s to 10 min were evaluated. The orange line refers to the overall accuracy in distinguishing all 4 states. Accuracy for individual states is shown in the other 4 colors. An accuracy of 25% reflects chance level classification. The overall accuracy remains at 80% with just 1 min of data. With scan lengths below 1 min, overall accuracy tends to decrease, though all 4 scans were identified with significant accuracy with only 30 s of data ($P < 0.001$).

et al. 2002). We used a binomial test to determine the significance of the classification accuracy and performed a paired-samples t -test to compare accuracy when using structural ROIs with accuracy when using functional ROIs.

Group-Level Contrasts of Rest and Memory States

We compared connectivity between and within ICNs using a paired-samples t -test for the memory state and the rest state of the 14 subjects used to train the classifier (Fig. 2D,E). To compare connectivity between all 90 ROIs in the rest and memory states (Fig. 2D), we performed a paired-samples t -test between the states for each of the 3960 pairwise correlation cells. Significant cells were determined by using an FDR-corrected P value of 0.05. To compare connectivity within the retrosplenial cortex (RSC)/medial temporal lobe (MTL) ICN (Fig. 2E), we performed ICA on each subject's rest and memory states. We fixed the number of independent components at 30 for each subject. We then used an automated template-matching procedure to select the RSC/MTL ICN for each scan (Greicius et al. 2004) using the group-level RSC/MTL as a template and compared the connectivity within this ICN for the subjects' rest and memory scans by performing a paired-samples t -test in SPM5. This analysis was masked to a one-sample t -test of the network derived from both states so that results would only reflect changes within the RSC/MTL network. Significant clusters of connectivity within the group-statistical map were determined by using the joint expected probability distribution (Poline et al. 1997) with height ($P < 0.01$) and extent ($P < 0.01$) thresholds, corrected at the whole-brain level.

Results

Subject-Driven Tasks Drive Connectivity Changes

For each of the 4 scans in each of the 14 subjects, a 90×90 matrix of pairwise ROI correlations was calculated (Fig. 2A,B). These matrices can be compared directly within a subject to reveal changes in connectivity strength between 2 subject-driven cognitive states as highlighted for the rest and memory states in Figure 2C. Group-level analyses confirmed these findings, revealing significant connectivity differences across the 90×90 matrix both within and between ICNs (Fig. 2D, $P < 0.05$, FDR corrected). In Figure 2E, we highlight a specific ICN that includes ROIs in the RSC

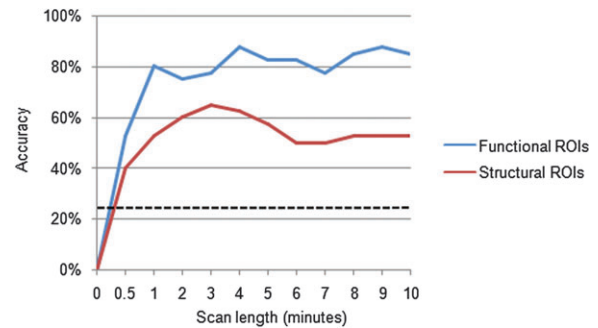


Figure 5. Functional ROIs outperform structural ROIs. We performed classification with 112 structural ROIs from the Automated Anatomical Labeling (AAL) Atlas and 90 functional ROIs identified by ICA on resting-state data from an independent sample. Classification was performed with both sets of ROIs at 11 different scan lengths. In all comparisons, classification with functional ROIs was substantially more accurate than classification with the AAL Atlas ROIs.

and MTL and that showed significantly increased connectivity in the memory state compared with rest ($P < 0.01$, corrected). There were no clusters in this ICN that showed significantly increased connectivity in the opposite contrast (rest > memory).

Group-Level State Matrices

In the group-level memory state matrix (Fig. 3B, orange arrow), several cells corresponding to correlations within the RSC/MTL ICN survive as would be expected from the group-level increases in this ICN during the memory state compared with rest (Fig. 2D,E). Equally important to the cells showing state-specific within-ICN correlations are the numerous cells showing state-specific between-ICN correlations. In Figure 3D (blue arrow), we emphasize increased connectivity during the subtraction task within an ICN that includes intraparietal sulcus (IPS) and prefrontal regions. In addition to increased intranetwork connectivity, the subtraction task elicited increased connectivity between the IPS/prefrontal ICN and the basal ganglia ICN (Fig. 3D, green arrow).

Classification of 4 Subject-Driven Cognitive States

The pattern-recognition classifier correctly identified 84% of the states in the LOOCV analysis (47 of 56 states; $P < 0.001$). Additionally, we used the group-level state matrices generated from the original cohort, shown in Figure 3, to classify the 4 cognitive states in an independent cohort of 10 new subjects acquired several months after the original cohort. In the independent cohort, 85% of the states were correctly classified (34 of 40 scans; $P < 0.001$). The mean state matrix fit scores for each of the 4 scan types across the 10 subjects are shown in Supplementary Figure S5.

Classification Accuracy as a Function of Scan Length

With a goal of applying this approach to more naturalistic (briefer) subject-driven cognitive states, we next examined the classifier accuracy over shorter scan durations in the independent cohort. Classification accuracy remained as high as 80% using only the first minute of data. Classification accuracy by scan length is shown in Figure 4, indicating that a high level of accuracy can be obtained with scan lengths as short as 30 s.

Rejecting a Novel Fifth Cognitive State

In this analysis, where the spatial navigation scans were added, 46 of 50 scans in the validation cohort were correctly classified yielding a classification accuracy of 92% ($P < 0.001$). Note that classification accuracy is higher here than in our main 4-way classification because we used a winner-take-all approach. When applied to the 4-way classification, the winner-take-all approach results in 100% accuracy for both the LOOCV and the independent cohort classification analyses (Supplementary Fig. S6). The mean state matrix fit scores for each of the 4 scan types were significantly greater than the novel cognitive state (Supplementary Fig. S7). For one participant, the spatial navigation task was confounded with the memory task; for another, the spatial navigation task was confounded with the subtraction task. The group-level state matrix for the spatial navigation task was not used to train the classifier but is shown in Supplementary Figure S4.

Comparison of Functional and Structural ROIs in Classification

Classification accuracy with the structural ROIs reached significance for all scan lengths ($P < 0.001$); however, the highest classification accuracy achieved with the structural ROIs was 65% (26 of 40 states correctly classified, Fig. 5). Additionally, a paired-samples *t*-test revealed that classification with the structural ROIs was significantly less accurate than classification with the functional ROIs ($P < 0.001$).

Discussion

Free-streaming subject-driven cognitive states account for a large, arguably the largest, portion of conscious processing (James 1918). Despite this principle, for methodological reasons human brain imaging research has focused mainly on discontinuous cognitive processing of discrete events whose explicit timing parameters are required. Most standard functional imaging studies therefore provide little insight into the fundamental, but difficult to interrogate, internal cognitive states which an individual moves between throughout the day. The results presented here demonstrate the potential for whole-brain FC analyses to provide an experimental window into naturalistic, continuous subject-driven cognitive processing. Face validity for this approach is provided by the finding that connectivity is significantly increased in an ICN connecting known memory regions like the RSC and MTL during the task in which subjects freely recalled the events of their day (Fig. 2D,E). Similarly, careful examination of Figure 3D reveals that the subtraction task increased connectivity in an ICN featuring bilateral IPS clusters typically activated in standard task activation studies of calculation (Dehaene et al. 1999; Menon et al. 2000; Gruber et al. 2001; Zago et al. 2001). In addition to changes in within-ICN connectivity, our approach also reveals state-specific patterns of between-ICN connectivity. Figure 3D shows that the same IPS network that increases its within-ICN connectivity during the subtraction task also increases its between-network connectivity with an ICN featuring the basal ganglia, deep gray matter regions frequently activated in arithmetic tasks (Dehaene et al. 1999; Menon et al. 2000). The FC measures used here do not allow us to distinguish the direct interactions between 2 regions from the temporally correlated but independent coactivation of 2 regions. The subtraction task, for example, may be driving

direct interactions between the IPS and basal ganglia networks in an integrated processing model. Alternatively, subtraction may be driving increased activity within each network separately in a parallel processing model. With that as an important limitation, these results support a theoretical model of cognitive processing in which cognitive state changes can be understood as a reshuffling of FC within and between a fixed set of distributed brain networks (McIntosh 2000).

These shifting patterns of connectivity associated with distinct cognitive states prove robust and consistent across subjects. In both a LOOCV analysis and in an analysis of an independent cohort, our algorithm classified 4 subject-driven cognitive states with high accuracy (84% in LOOCV and 85% in the independent cohort). Importantly, classification accuracy remained at 80% when only the first minute of data was used and remained well-above chance level with as little as 30 s of data. These observations demonstrate that our approach can classify cognitive states over short intervals of self-driven cognition that may better reflect naturalistic processing than the 10-min period used in our main analysis. The temporal resolution of this approach is not likely to improve much further than the 30–60 s range below which our accuracy decreases (Fig. 4). The temporal resolution is limited ultimately by the very low frequency of the BOLD signal fluctuations studied here (Cordes et al. 2001). Recent work in the temporal domain demonstrates prominent nonstationarity of resting-state FC over time, particularly as it relates to negative correlations (Chang and Glover 2010). The classification algorithm developed here could potentially be improved upon by accounting for such variability in connectivity strength.

We anticipate that the spatial resolution of this approach will continue to improve with advances in fMRI acquisition and analysis. Many of the ROIs used here are still relatively large (Fig. 1 and Supplementary Fig. S2) and can likely be subdivided further with increasingly sophisticated parcellation approaches. Combining resting-state fMRI with diffusion tensor tractography (Rushworth et al. 2006; Greicius et al. 2009) and self-clustering FC algorithms (Cohen et al. 2008) represent 2 promising approaches to more finely parcellating gray matter into increasingly indivisible mesoscopic functional units. A third approach would entail mandating a higher model order in the group ICA, so that instead of identifying 14 networks from 30 components as was done here, one might, for example, identify 20 networks in 50 components (Kiviniemi et al. 2009; Smith et al. 2009). With the whole-brain connectivity matrix approach defined here, doubling the number of ROIs from 90 to 180 would increase the matrix size exponentially from 3960 cells to 16 020 cells which may further enhance discriminability between cognitive states.

A number of recent studies have used structurally defined ROIs to gain insights into the basic properties of both structural and functional brain networks (Hagmann et al. 2008; Supekar et al. 2008; Honey et al. 2009) and to classify brain states (Richiardi et al. 2010). While comparisons between different structural parcellation schemes have been undertaken (Zalesky et al. 2010), we are not aware of any prior studies comparing functionally defined ROIs with structural ROIs. The results presented here strongly suggest that the use of functionally defined ROIs is superior to the use of structurally defined ROIs in whole-brain connectivity analyses. Across all scan durations tested in this study, the functional ROIs consistently and significantly outperformed the structural ROIs. The main risk in

using structural ROIs is that data from 2 or more distinct functional regions may be combined into a single time series. When this occurs, classification potential is likely weakened by diluting meaningful information from 2 distinct ROIs. In addition, combining 2 functional ROIs into a single structural ROI has the potential to introduce errors by creating novel hybrid structural ROI time series that do not reflect the true functional information of either functional ROI.

By focusing on FC rather than functional activation, the current study, like that of Richiardi et al. (2010) has overcome a substantial limitation of previous classification approaches: the reliance on investigator-imposed or subject-reported temporal landmarks (Haynes and Rees 2006). More critically, this is the first study to demonstrate that whole-brain FC can be used to classify subject-driven cognitive processing. The 4 states classified here were paced not externally by the investigators but internally by the subjects themselves, more closely reflecting the free-streaming cognitive processing characteristic of human thought. Refinements of this approach should ultimately allow for the classification of a wide array of subject-driven cognitive states from relatively short samples of brain imaging data.

Supplementary Material

Supplementary material can be found at: <http://www.cercor.oxfordjournals.org/>

Funding

Dana Foundation; John Douglas French Alzheimer's Foundation; National Institutes of Health (AT005733, HD059205, HD057610, NS073498, NS058899).

Notes

We thank J. Damoiseaux, A. Etkin, B. Q. Ford, A. Jafari, M. Molino, and W. Seeley for their comments on the manuscript. *Conflict of Interest*: None declared.

References

Beckmann CF, DeLuca M, Devlin JT, Smith SM. 2005. Investigations into resting-state connectivity using independent component analysis. *Philos Trans R Soc Lond B Biol Sci.* 360:1001-1013.

Bellec P, Perlbarg V, Jbabdi S, Pelegri-Issac M, Anton JL, Doyon J, Benali H. 2006. Identification of large-scale networks in the brain using fMRI. *Neuroimage.* 29:1231-1243.

Biswal B, Yetkin FZ, Haughton VM, Hyde JS. 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med.* 34:537-541.

Chadwick MJ, Hassabis D, Weiskopf N, Maguire EA. 2010. Decoding individual episodic memory traces in the human hippocampus. *Curr Biol.* 20:544-547.

Chang C, Glover GH. 2009. Effects of model-based physiological noise correction on default mode network anti-correlations and correlations. *Neuroimage.* 47:1448-1459.

Chang C, Glover GH. 2010. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage.* 50:81-98.

Chiu YC, Esterman M, Han Y, Rosen H, Yantis S. 2010. Decoding task-based attentional modulation during face categorization. *J Cogn Neurosci.* 23:1198-1204.

Cohen AL, Fair DA, Dosenbach NU, Miezin FM, Dierker D, Van Essen DC, Schlaggar BL, Petersen SE. 2008. Defining functional areas in individual human brains using resting functional connectivity MRI. *Neuroimage.* 41:45-57.

Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, Quigley MA, Meyerand ME. 2001. Frequencies contributing to

functional connectivity in the cerebral cortex in "resting-state" data. *AJNR Am J Neuroradiol.* 22:1326-1333.

Damoiseaux JS, Rombouts SA, Barkhof F, Scheltens P, Stam CJ, Smith SM, Beckmann CF. 2006. Consistent resting-state networks across healthy subjects. *Proc Natl Acad Sci U S A.* 103:13848-13853.

Davatzikos C, Ruparel K, Fan Y, Shen DG, Acharyya M, Loughhead JW, Gur RC, Langleben DD. 2005. Classifying spatial patterns of brain activity with machine learning methods: application to lie detection. *Neuroimage.* 28:663-668.

Dehaene S, Le Clecq HG, Cohen L, Poline JB, van de Moortele PF, Le Bihan D. 1998. Inferring behavior from functional brain images. *Nat Neurosci.* 1:549-550.

Dehaene S, Spelke E, Pineda P, Stanescu R, Tsivkin S. 1999. Sources of mathematical thinking: behavioral and brain-imaging evidence. *Science.* 284:970-974.

Esterman M, Chiu YC, Tamber-Rosenau BJ, Yantis S. 2009. Decoding cognitive control in human parietal cortex. *Proc Natl Acad Sci U S A.* 106:17974-17979.

Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME. 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci U S A.* 102:9673-9678.

Fransson P. 2006. How default is the default mode of brain function? Further evidence from intrinsic BOLD signal fluctuations. *Neuropsychologia.* 44:2836-2845.

Friston KJ. 1998. Imaging neuroscience: principles or maps? *Proc Natl Acad Sci U S A.* 95:796-802.

Friston KJ, Ashburner J, Frith CD, Poline JB, Heather JD, Frackowiak RSJ. 1995. Spatial registration and normalization of images. *Hum Brain Mapp.* 3:165-189.

Glover GH, Lai S. 1998. Self-navigated spiral fMRI: interleaved versus single-shot. *Magn Reson Med.* 39:361-368.

Glover GH, Law CS. 2001. Spiral-in/out BOLD fMRI for increased SNR and reduced susceptibility artifacts. *Magn Reson Med.* 46:515-522.

Greicius MD, Krasnow B, Reiss AL, Menon V. 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc Natl Acad Sci U S A.* 100:253-258.

Greicius MD, Srivastava G, Reiss AL, Menon V. 2004. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. *Proc Natl Acad Sci U S A.* 101:4637-4642.

Greicius MD, Supekar K, Menon V, Dougherty RF. 2009. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cereb Cortex.* 19:72-78.

Gruber O, Indefrey P, Steinmetz H, Kleinschmidt A. 2001. Dissociating neural correlates of cognitive components in mental calculation. *Cereb Cortex.* 11:350-359.

Hagmann P, Cammoun L, Gigandet X, Meuli R, Honey CJ, Wedeen VJ, Sporns O. 2008. Mapping the structural core of human cerebral cortex. *PLoS Biol.* 6:e159.

Hampson M, Driesen NR, Skudlarski P, Gore JC, Constable RT. 2006. Brain connectivity related to working memory performance. *J Neurosci.* 26:13338-13343.

Hampson M, Peterson BS, Skudlarski P, Gatenby JC, Gore JC. 2002. Detection of functional connectivity using temporal correlations in MR images. *Hum Brain Mapp.* 15:247-262.

Harrison BJ, Pujol J, Ortiz H, Fornito A, Pantelis C, Yucel M. 2008. Modulation of brain resting-state networks by sad mood induction. *PLoS One.* 3:e1794.

Hastie T, Tibshirani R, Friedman J. 2009. The elements of statistical learning. 2nd ed. New York: Springer.

Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science.* 293:2425-2430.

Haynes JD, Rees G. 2005a. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci.* 8:686-691.

Haynes JD, Rees G. 2005b. Predicting the stream of consciousness from activity in human visual cortex. *Curr Biol.* 15:1301-1307.

- Haynes JD, Rees G. 2006. Decoding mental states from brain activity in humans. *Nat Rev Neurosci.* 7:523-534.
- Honey CJ, Sporns O, Cammoun L, Gigandet X, Thiran JP, Meuli R, Hagmann P. 2009. Predicting human resting-state functional connectivity from structural connectivity. *Proc Natl Acad Sci U S A.* 106:2035-2040.
- Jay W. 1918. *The principles of psychology.* Vol. 1. New York: Holt and Company.
- Kamitani Y, Tong F. 2005. Decoding the visual and subjective contents of the human brain. *Nat Neurosci.* 8:679-685.
- Kay KN, Naselaris T, Prenger RJ, Gallant JL. 2008. Identifying natural images from human brain activity. *Nature.* 452:352-355.
- Kim DH, Adalsteinsson E, Glover GH, Spielman DM. 2002. Regularized higher-order in vivo shimming. *Magn Reson Med.* 48:715-722.
- Kiviniemi V, Kantola JH, Jauhiainen J, Hyvarinen A, Tervonen O. 2003. Independent component analysis of nondeterministic fMRI signal sources. *Neuroimage.* 19:253-260.
- Kiviniemi V, Starck T, Remes J, Long X, Nikkinen J, Haapea M, Veijola J, Moilanen I, Isohanni M, Zang YF, et al. 2009. Functional segmentation of the brain cortex using high model order group PICA. *Hum Brain Mapp.* 30:3865-3886.
- McIntosh AR. 2000. Towards a network theory of cognition. *Neural Netw.* 13:861-870.
- Menon V, Rivera SM, White CD, Glover GH, Reiss AL. 2000. Dissociating prefrontal and parietal cortex activation during arithmetic processing. *Neuroimage.* 12:357-365.
- Mitchell TM, Hutchinson R, Just MA, Niculescu RS, Pereira F, Wang X. 2003. Classifying instantaneous cognitive states from fMRI data. *AMIA Annu Symp Proc.* 465-469.
- Mourao-Miranda J, Bokde AL, Born C, Hampel H, Stetter M. 2005. Classifying brain states and determining the discriminating activation patterns: support Vector Machine on functional MRI data. *Neuroimage.* 28:980-995.
- O'Craven KM, Kanwisher N. 2000. Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *J Cogn Neurosci.* 12:1013-1023.
- Owen AM, Coleman MR, Boly M, Davis MH, Laureys S, Pickard JD. 2006. Detecting awareness in the vegetative state. *Science.* 313:1402.
- Poline JB, Worsley KJ, Evans AC, Friston KJ. 1997. Combining spatial extent and peak intensity to test for activations in functional imaging. *Neuroimage.* 5:83-96.
- Polyn SM, Natu VS, Cohen JD, Norman KA. 2005. Category-specific cortical activity precedes retrieval during memory search. *Science.* 310:1963-1966.
- Richiardi J, Eryilmaz H, Schwartz S, Vuilleumier P, Van De Ville D. 2010. Decoding brain states from fMRI connectivity graphs. *Neuroimage.*
- Rushworth MF, Behrens TE, Johansen-Berg H. 2006. Connection patterns distinguish 3 regions of human parietal cortex. *Cereb Cortex.* 16:1418-1430.
- Seeley WW, Menon V, Schatzberg AF, Keller J, Glover GH, Kenna H, Reiss AL, Greicius MD. 2007. Dissociable intrinsic connectivity networks for salience processing and executive control. *J Neurosci.* 27:2349-2356.
- Smith SM, Fox PT, Miller KL, Glahn DC, Fox PM, Mackay CE, Filippini N, Watkins KE, Toro R, Laird AR, et al. 2009. Correspondence of the brain's functional architecture during activation and rest. *Proc Natl Acad Sci U S A.* 106:13040-13045.
- Supekar K, Menon V, Rubin D, Musen M, Greicius MD. 2008. Network analysis of intrinsic functional brain connectivity in Alzheimer's disease. *PLoS Comput Biol.* 4:e1000100.
- Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, 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.
- Walther DB, Caddigan E, Fei-Fei L, Beck DM. 2009. Natural scene categories revealed in distributed patterns of activity in the human brain. *J Neurosci.* 29:10573-10581.
- Zago L, Pesenti M, Mellet E, Crivello F, Mazoyer B, Tzourio-Mazoyer N. 2001. Neural correlates of simple and complex mental calculation. *Neuroimage.* 13:314-327.
- Zalesky A, Fornito A, Harding IH, Cocchi L, Yucel M, Pantelis C, Bullmore ET. 2010. Whole-brain anatomical networks: does the choice of nodes matter? *Neuroimage.* 50:970-983.