

Title:

Large-scale multivariate activation states of the human brain

Authors:

Richard H. Chen¹, Takuya Ito¹, Kaustubh R. Kulkarni¹, and Michael W. Cole¹

Author affiliations:

¹ Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ, USA

Corresponding author:

Michael W. Cole
197 University Ave.
Rutgers University
Newark, NJ 07102
michael.cole@rutgers.edu
973-353-3249

Keywords: fMRI, resting-state, dynamics, machine learning, graph theory, community detection

Abstract

Much of our lives are spent in unconstrained rest states, yet cognitive brain processes are primarily investigated experimentally using task-constrained states. It may be possible to utilize the insights gained from experimental control of task processes as reference points for investigating unconstrained rest. To facilitate comparison of rest and task functional MRI (fMRI) data we focused on activation amplitude patterns, commonly used for task but not rest analyses. During rest, we identified spontaneous changes in temporally extended whole-brain activation pattern states. This revealed a hierarchical organization of rest states. The top consisted of two competing states consistent with previously identified “task-positive” and “task-negative” activation patterns. These states were composed of a variety of more specific states that repeated over time and across individuals. In contrast to the classic view that rest consists of only task-negative states, task-positive states were present over 40% of the time while individuals “rested.” This suggests individuals conduct extensive task-oriented mental activity during rest. Consistent with this interpretation, classifier decoding of rest brain states based on thousands of task fMRI studies identified a variety of matched active tasks. Further, an analysis of task data revealed a similar hierarchical structure of brain states. Together these results suggest brain activation dynamics form a general hierarchy across both task and rest, with a small number of dominant general states reflecting basic modes of brain function along with a variety of specific states likely reflecting the rich variety of cognitive states experienced in everyday life.

Significance Statement

Neuroimaging has revolutionized understanding of human brain function, yet most brain activation studies have used instructed tasks, limiting understanding of the full range of human brain activation states. When unconstrained (“rest”) states have been investigated, activation of a “default mode” network involved in self-reflective cognition has been detected. This suggests rest periods consist of self-reflection, yet everyday experience suggests task-focused thoughts – such as attending to the environment and efforts to solve the problems of everyday life – also occur during “rest”. We confirmed this, identifying the broad organizational structure of spontaneous brain states in humans. We found that rest consists of approximately 60% self-reflection and 40% task-focused events, with multiple cognitively meaningful sub-states identified using activity pattern decoding.

Introduction

The brain is a distributed information-processing system with rich spatiotemporal dynamics underlying complex cognitive dynamics. A major goal of cognitive neuroscience is to create a mapping between these two forms of dynamics to better understand the neural basis of cognition. Recent insights in human neuroimaging research have improved this mapping by considering activity in more than one spatial location at a time. These approaches include multivariate pattern analysis (MVPA) of brain activity patterns corresponding to cognitive task events (1, 2), as well as functional connectivity (FC) analysis of brain network dynamics during rest and task (3–5). Building on these advances, here we use what we term “dynamic MVPA” (dMVPA) – MVPA applied to the temporal evolution of brain processes (6–9) – with functional MRI (fMRI) to more comprehensively characterize the repertoire of brain states across a variety of resting and task cognitive states.

In particular, we used a state space characterization of dynamics previously used to gain insight into other real-world complex systems (10, 11). This involves conceptualizing each whole-brain image in time as a single point in a high-dimensional feature space. Brain state changes are thus equivalent to movement through that state space. We primarily utilize distance between these points in state space, using spatial correlations (a standard distance metric) between time points to temporally cluster them into brain states extending through time. We perform further analysis using graph theoretical formulations. Relative to classical clustering approaches, graph theoretical community detection algorithms can better assign patterns near the edges of a cluster (12). This can provide a more comprehensive characterization of a state space’s large-scale organization.

Recent investigations into moment-to-moment changes in FC have been important for characterizing brain dynamics (5). Unlike these studies of FC dynamics, we identify activation pattern-based brain states, which are not limited in their temporal precision by the need for temporally extended windows. More importantly, identifying whole-brain activation pattern states (rather than FC/temporal covariance patterns (13, 14)) facilitates the functional interpretability of those states. This is due to whole-brain fMRI activation patterns being more directly related to the vast fMRI task activation literature, which associates cognitive task manipulations with whole-brain spatial activation maps (15, 16). In contrast, previous studies relating resting-state FC to the fMRI activation literature focused on across-study covariance between brain regions (e.g., the similarity in activation level changes across all tasks) (13, 14), rather than focusing on any specific activation patterns (e.g., the whole-brain activation pattern during the N-back task). These studies revealed resting-state networks that were similar to patterns of across-study covariance, yet single brain activation states are known to activate multiple resting-state networks at once (14, 17). Several recent studies have focused on resting-state activation patterns (18, 19). However, these studies focused on relating these patterns to resting-state networks, rather than task activations. Here we compare resting-state activation patterns to specific task activation maps – potentially involving coactivation of multiple resting-state networks – to identify the likely functional relevance of spontaneous brain states during rest. We expected this approach to better characterize the organization of human brain states and potentially identify task-related states present during “rest” periods.

We propose the use of the term dMVPA given the similarity of this approach to existing MVPA approaches with fMRI and M/EEG. However, whereas standard MVPA uses supervised machine learning (1) to classify activation patterns during experimentally induced cognitive events, dMVPA can be applied to the moment-by-moment temporal evolution of brain states to characterize even spontaneous cognitive events. A well-established example of dMVPA is the identification of M/EEG microstates, as defined by spontaneous changes in the spatial configuration of the scalp electromagnetic field (20). Critically, however, the identification of M/EEG microstates suffers from strong limitations on spatial certainty caused by the smearing

of neuronal sources as the electromagnetic field is conducted to the scalp surface (21, 22). This spatial smearing suggests identified M/EEG microstates are biased toward states that are spatially extreme (i.e., spatially distinct despite smearing) rather than both temporally and spatially distinct. Consistent with this, attempts to map M/EEG microstates to more spatially accurate fMRI data has yielded states with only limited similarity to well-known states from the vast fMRI task activation literature (7). Thus, we use a complementary approach here with fMRI to identify more spatially precise brain states during rest. While others have recently begun investigating spontaneous activation patterns with fMRI (19, 23), these patterns have not been directly compared to task activation patterns. Importantly, as mentioned above, comparison to task activation patterns is a major advantage of this approach, since using fMRI allows comparison of spontaneous states to the vast fMRI task activation literature, which associates cognitive manipulations with whole-brain spatial activation maps (15, 16).

We primarily utilize rest data given the possibility that many brain states are visited in this unconstrained context (3). This allowed us to obtain a broad sampling of possible brain states across many (N=97) individuals. We also supplemented this broad repertoire of spontaneous states with experimentally controlled states identified from a variety of tasks involving distinct cognitive functions. We hypothesized that activation pattern states form a hierarchical structure, with a small number of dominant high-level states reflecting basic modes of brain function and a wide variety of lower-level states reflecting the rich variety of cognitive states humans are capable of visiting.

In addition to this broad hypothesis, we sought to test the hypothesis that “task-positive” brain states occur regularly during rest periods. This contrasts with the highly replicated finding that a “task-negative” brain state with high default-mode network activity is present during rest (3, 24, 25). We expected that the task-negative state dominates rest (consistent with previous results), but that task-positive states also occur frequently. This expectation is based on the common self-reported experience of trying to solve the problems of everyday life during periods of “rest” (26). Such problem solving effort is associated with cognitive control demands, which involve task-positive activations in networks such as the frontoparietal control network and the cingulo-opercular network and deactivation of the default-mode network (27–29). Additionally, everyday experience suggests some portion of “rest” consists of passive attending to external sensory and motor events, suggesting occasional activation of sensory/motor networks and attention networks (30). We expected that identifying the overall organization of spontaneous brain state dynamics would allow us to test the possibility of both task-positive and task-negative brain states occurring during rest periods, as well as allowing us to estimate the proportion of time devoted to each class of brain states.

MATERIALS & METHODS

Participants

Data were collected as part of the Washington University-Minnesota Consortium Human Connectome Project (31). The participants were recruited from the Washington University campus and surrounding area. All participants supplied informed consent. The data were from the “500 Subjects” public data release. We used data from the “100 Unrelated Subjects” set as we wanted a sample representative of the general population (excluding family relations).

We used resting-state fMRI and task-state fMRI data from 100 subjects, with 3 outlier subjects removed for a subset of analyses (see Results). The resting-state dataset consisted of four separate runs, each spanning 14.4 minutes in length. Analyses were performed separately for each rest run. The task data involved seven diverse tasks (32). These seven tasks were selected to tap into different cognitive processes as well as the different neural circuitry that supports those functions. The tasks were related to emotion perception, reward learning,

language processing, motor responses, relational reasoning, social cognition, and working memory. See (32) for more details on the task fMRI data.

MRI parameters

Whole-brain echo-planar scans were acquired with a 32 channel head coil on a modified 3T Siemens Skyra with TR = 720 ms, TE = 33.1 ms, flip angle = 52°, BW = 2290 Hz/Px, in-plane FOV = 208 × 180 mm, 72 slices, 2.0 mm isotropic voxels, with a multi-band acceleration factor of 8 (33). Data were collected across two days. On each day 28 minutes of rest (eyes open with fixation) fMRI data was collected across two runs (56 minutes total), followed by 30 minutes of task fMRI data collection (60 minutes total). Each of the 7 tasks was completed over two consecutive fMRI runs. Details regarding the resting-state data collection for this dataset can be found elsewhere (34), as well as details about the tasks (32).

fMRI preprocessing

We used a minimally preprocessed version of the data, which was the result of standard procedures including: spatial normalization to a standard template, motion correction, and intensity normalization. These steps have been described previously (35). We performed analyses on the volume version of these minimally preprocessed data using AFNI (36). We removed variables of no interest from the time series using linear regression, including: motion estimates, ventricle and white matter signals, as well as derivatives. Ventricle, white matter, gray matter, and anatomical structures were identified for each subject using FreeSurfer (37, 38). Note that whole-brain global signal was not removed due to controversy regarding this preprocessing step (39). The linear trend was removed from the signal and the data were spatially smoothed (FWHM = 4 mm). Resting fMRI data is also typically temporally filtered to isolate the low frequency component of the time series. We did not apply a temporal filter to the data due to the possibility of relatively rapid brain state transitions.

Further data analysis was completed by sampling from a set of 264 regions in order to capture and explore large-scale regional and system-level questions. These 264 regions were independently identified in a previous study using functional criteria (40). Using this approach reduces the chance of blurring signal from neighboring regions with different functional profiles (41). Specifically, the 264 regions were identified and classified using resting-state FC parcellation (42) and a task-based neuroimaging meta-analysis (40). The mean time series from all of the voxels within each of these 264 regions was then calculated and used in subsequent analyses. Subsequent data analyses were conducted with MATLAB 2014b (Mathworks) or Python.

Brain state identification

We conceptualized time in terms of a weighted graph, with each time point as a graph node and edges as the similarities of whole-brain spatial activation patterns at those time points. Activation pattern similarity was calculated using Pearson correlation. Pattern similarity can be equivalently considered as (the inverse of) distance in state space, which has a long methodological history in mathematics and other fields (43). Pearson correlation was used given that it is a well-established distance measure (43), is invariant to scale changes (unlike, e.g., Euclidean distance), and can be conveniently conceptualized as (the square root of) linear variance explained. Note that Pearson correlation does not assume the underlying data are normally distributed when used as a similarity measure, given that its associated p-value is not calculated (44). These correlation-based associations were summarized in a temporal adjacency matrix, consisting of all pairwise similarities among time points.

Time points with strong edges between them were considered as instances of the same brain state. Clusters of similar time points were identified using the Infomap community detection algorithm (45–47). This community detection algorithm was first applied to each

subject's initial resting-state fMRI run (14.4 minutes in duration, consisting of 1200 fMRI time points). Given the likely presence of noise, thresholded the temporal adjacency matrix prior to this stage, removing edges with low similarity. We used a standard consensus approach (48) in which versions of the matrix at thresholds from 1% to 50% density (in 1% increments) were combined based on their modularity score (49, 50). This produced 4 unique brain states (graph communities) per subject on average, with a total of 412 brain states across 98 subjects. Note that two of the 100 subjects were excluded because the clustering algorithm returned more than three standard deviations above the average number of states per subject. We reduced the many time points contributing to each brain state to a single "prototype" vector for each brain state via averaging.

We next conducted a group-level analysis. This involved first computing an adjacency matrix based on the similarities among the brain state prototypes across all subjects. We then applied the Infomap algorithm to that matrix, this time with no thresholding, to produce a group-level brain state solution. Later, we tested for the presence of a hierarchy of brain states, with the no-threshold solution being the top of the hierarchy. We produced lower levels of the hierarchy by re-running the algorithm (independently of the previous level) using a series of thresholds. We began with 100% density (no threshold), going down by 10% increments until the 10% density level. We also conducted a separate analysis going from the 10% density level down to the 2% density level in 1% increments.

Task relevance – correspondence between brain states and behavioral/mental states

We sought to identify correspondences between the identified brain states and known behavioral/mental states based on task manipulations. This allowed assessment of the likely functional relevance of the identified brain states. Task-rest correspondence was assessed using pairwise Pearson's correlations between each group-level brain state and every activation pattern in each time point across all subjects for both resting fMRI and task fMRI. This returned a time series of correlation scores between each subject and the group-level brain states. The time series of all subjects were then visualized and assessed for temporal dynamics of brain state transitions.

Assessing the amount of time spent in each brain state

For resting-state data, each individual's Infomap clustering solution was re-mapped into a two-state solution based on the group-level clustering. The durations in each state were then calculated based on this two-state solution for each individual subject. For task-state data, each time point was classified as either State A or State B using a support vector machine (SVM) classifier trained on resting-state data. Specifically, the SVM classifier was trained on rest runs 2-4 and validated by testing on rest run 1 prior to classification of task state data (for each task separately). Note that validation with rest run 1 data was based on above-chance classification accuracy, using the InfoMap clustering labels as the ground truth. The two-state classification results were used to determine the amount spent in each state for each of the seven tasks.

Neurosynth state decoding

Brain states were decoded using the Neurosynth decoder tool (15). Neurosynth is a meta-analytical tool that contains a database with brain activation patterns and peak signal coordinates paired with the associated cognitive terms from the fMRI scientific literature. The decoder function takes in our voxelwise representation of brain states, cross-references with the database, and returns a list of cognitive terms each paired with a correlation score indicating how well each brain state is associated with each cognitive term. The decoder returned a list of

approximately 3400 cognitive terms and their correlation values with the tested brain state. Of the top 50 highest correlated terms, all anatomical terms, redundant terms, nonsensical terms (i.e., “cortexmpfc”, “networkdmn”), and methodological terms were excluded from the list. The remaining terms were visualized as word clouds and the relative font sizes of each term were determined by the correlation scores.

We sought to validate the results returned by Neurosynth by running a series of permutation tests for each brain state. Specifically, within each subject, we permuted the order of the brain state temporal windows (consecutive time points of the same state) as returned by the Infomap clustering process. This approach allowed us to preserve many of the temporal properties of brain states within each subject, while still permuting the data being averaged for each unique brain state label. The newly “permuted” brain state prototypes were averaged across subjects according to the group-level clustering solution to produce a “permuted” State A and State B. The “permuted” State A and State B were run through the Neurosynth decoder to obtain correlation scores for each of the 3400 terms. A total of 1000 permutations cycles were performed for each brain state, returning a total of 340,000 correlation scores. This allowed us to create a null distribution of r-values to test the statistical significance of the observed r-value for each cognitive term. Note that these permutation tests controlled for multiple comparisons, since all 3400 comparisons were computed during each permutation (51). The same analysis was also performed for each of the 12 states at the 20% density tier of the hierarchy (see Supporting Information for methods, detailed results, and figure).

RESULTS

Identifying multivariate brain activation states

We sought to characterize the repertoire of spontaneous human brain activation patterns using a data-driven method of temporally clustering resting-state fMRI activity. We hypothesized that whole-brain fMRI activation patterns would consist of discrete “states” – configuration patterns that extend (and repeat) through time. To test this possibility we clustered whole-brain resting-state fMRI activations in time. A standard set of functionally defined regions (52) were used for computational tractability, as well as for the previously identified functional system assignments (to potentially increase interpretability of observed states) (Figure 1D). Spatial correlations were used as a similarity/distance measure across individual time points, with brain states defined as temporal clusters of similar activation patterns (Figure 1A).

Brain states were identified first at the individual subject level. Of the 100 subjects, two subjects were excluded from subsequent analyses (unless noted otherwise) because the clustering algorithm returned more than three standard deviations above the average number of states per subject. On average, four unique brain states were identified for each subject for a total of 412 unique brain state prototypes across the remaining 98 subjects. These 412 brain states were re-clustered in a group level analysis. The same processing steps performed at the individual subject level were applied at the group level for determining group-level brain states (see Methods for details).

At the group level two brain states were identified, which we labeled “State A” and “State B” (Figure 1B & 1C). The states were summarized by averaging all prototypes with the same clustering label and visualized in 264 ROI space (Figure 1B) and in voxel space (Figure 1C). An example subject’s temporal adjacency matrix along with the group-level clustering solution is illustrated in Figure 1A. Periods of highly similar activation patterns can be observed by the blocked structure on the diagonal of the temporal adjacency matrix, which are grouped as the same state in the clustering solution. We replicated the current findings using the same approach for the remaining three rest runs and found similar results (second run shown in Figure 1B). Specifically, State A identified in any given rest run was highly correlated on

average with State A identified in other rest runs: rank correlation $\rho=0.92$ ($p<0.0001$ for all pairwise comparisons). This was also the case for State B being correlated with State B identified in other rest runs: average rank correlation $\rho=0.91$ ($p<0.0001$ for all pairwise comparisons). Note that we used Spearman's rank correlation (ρ) to calculate p-values for data that were not approximately normally distributed. Additionally, State A and State B were highly anti-correlated at $\rho=-0.97$ on average ($p<0.0001$ for all pairwise comparisons). Note that the clustering approach used to identify distinct activation states prioritizes identification of maximally distinct (e.g., anti-correlated) states. To confirm the statistical significance of the observed anti-correlation despite this prioritization of maximally distinct states, we performed a permutation test with 10000 iterations, finding that the observed result was $p<0.0001$ (see Supporting Information section for methods and detailed results).

Functionally, State A is highly similar to observed “task-positive” activation patterns reported in the literature (52–54). This activation pattern includes active regions in many task-related functional networks, such as the dorsal and ventral attention networks, salience network, and sensory-motor networks (visual, auditory, somatomotor). Conversely, State B may be a “task-negative” state, given the strong activation of the default mode network (DMN) (53, 55) (along with a portion of the fronto-parietal cognitive control network (FPN)). Note, however, that we used only “task-negative” data – fMRI data collected during rest – to define both of these states.

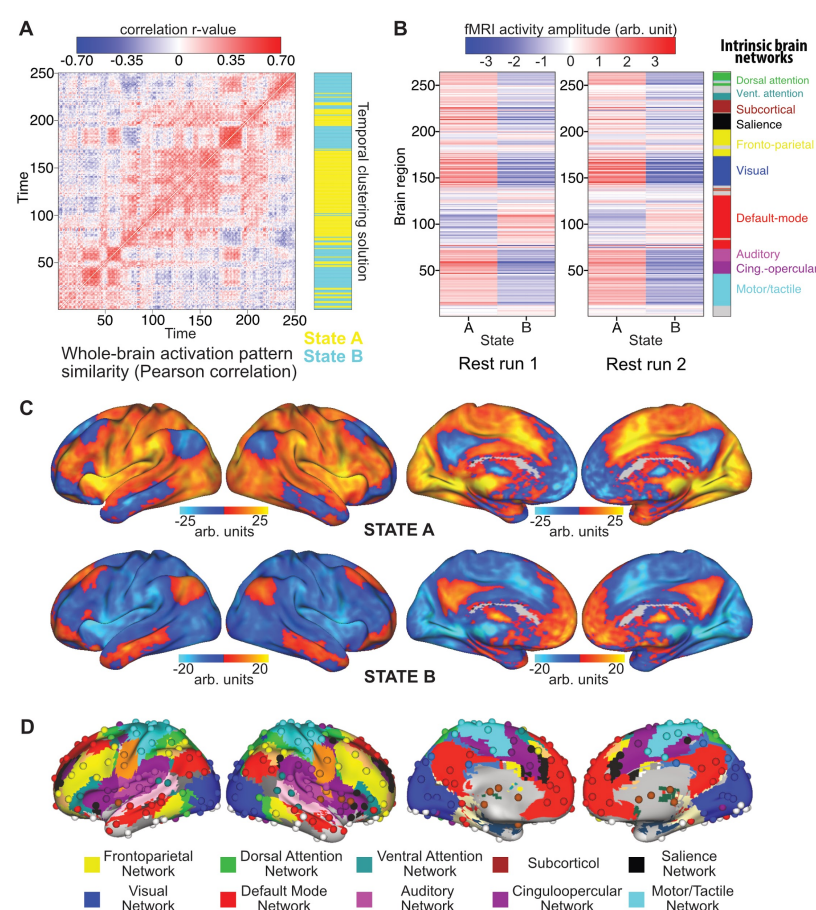


Figure 1 – The repertoire of brain states based on resting-state fMRI. Each individual's prototypical brain states are correlated and clustered using the same Infomap algorithm. **A.** An example of individual time point similarity based on spatial correlations. 250 time points from a single subject are shown for

illustration. The across-subject clustering result (mapped back to this subject's data) is shown in blue and yellow. **B.** Brain state prototypes for state A and state B, averaged across all activation vectors across all subjects. The two-state results replicated across all 4 of the rest runs (results from 2 runs are shown in **B**). **C.** Voxelwise representations of state A and state B. Note that state B involves activation of the default-mode network as well as the tan/salmon-colored portion of the frontoparietal network in panel D. **D.** Functionally defined set of 264 regions and the associated functional network assignments.

Spatial activation patterns vs. FC

Using resting-state data, we identified a “task-positive” state A and a “task-negative” state B. These labels were based on the functions assigned to these activation patterns by the task activation fMRI literature (13, 29, 56, 57). As mentioned above, several aspects of these state A and state B patterns (Figure 1B) are consistent with networks identified in the resting-state FC literature (52, 53, 58). However, the approach used here was quite distinct from FC analyses, since it is based on spatial activation covariance, which is ignored by FC approaches in favor of temporal covariance. Therefore, in order to test whether the current spatial activation pattern approach reveals additional insights independent of traditional FC analyses, we analyzed resting-state FC patterns across time periods identified as either State A or State B (Fig 3). Comparing between the FC patterns calculated by averaging across all subjects (Figure 2), we found that the overall resting-state FC pattern (calculated across the entire rest run) correlated with State A FC at $r=0.948$ and State B FC at $r=0.951$. Additionally, to test whether results were consistent across subjects, FC comparisons were also performed for each subject separately. We then performed random-effects statistical analyses to test for across-subject consistency. The overall resting-state FC pattern correlated with State A FC on average at $r=0.55$ ($t(96)=42.4$, $p<0.0001$) and State B FC on average at $r=0.54$ ($t(96) = 49.1$, $p<0.0001$). Note that we calculated the overall resting-state FC pattern based on a separate run (e.g., run 2 when comparing to run 1's State A FC) to remove circularity from the analysis (59). All non-circular combinations (between the first and second resting-state run) were computed and the reported results are based on averages across these comparisons.

The strong correlation between the individual states' FC and overall rest FC suggests that resting-state FC architecture remains consistent across distinct activation brain states. Further supporting this conclusion, across-subject averaged State A FC correlated with across-subject averaged State B FC at $r=0.951$. This was also the case when computed based on averaging within-subject comparisons: $r=0.49$ ($t(96) = 48.6$, $p<0.0001$), suggesting across-subject consistency. Note that this was calculated by comparing across rest runs, as with the overall resting-state FC pattern comparisons above. The resulting average similarity is very high relative to chance, but also in contrast to the strong anti-correlation ($\rho=-0.97$) between State A and State B activation patterns. This suggests the possibility that whole-brain activation pattern dynamics are largely independent of FC dynamics. This supports our choice to analyze activation patterns independently of FC, as well as identifying brain states based on activation pattern similarities rather than FC dynamics.

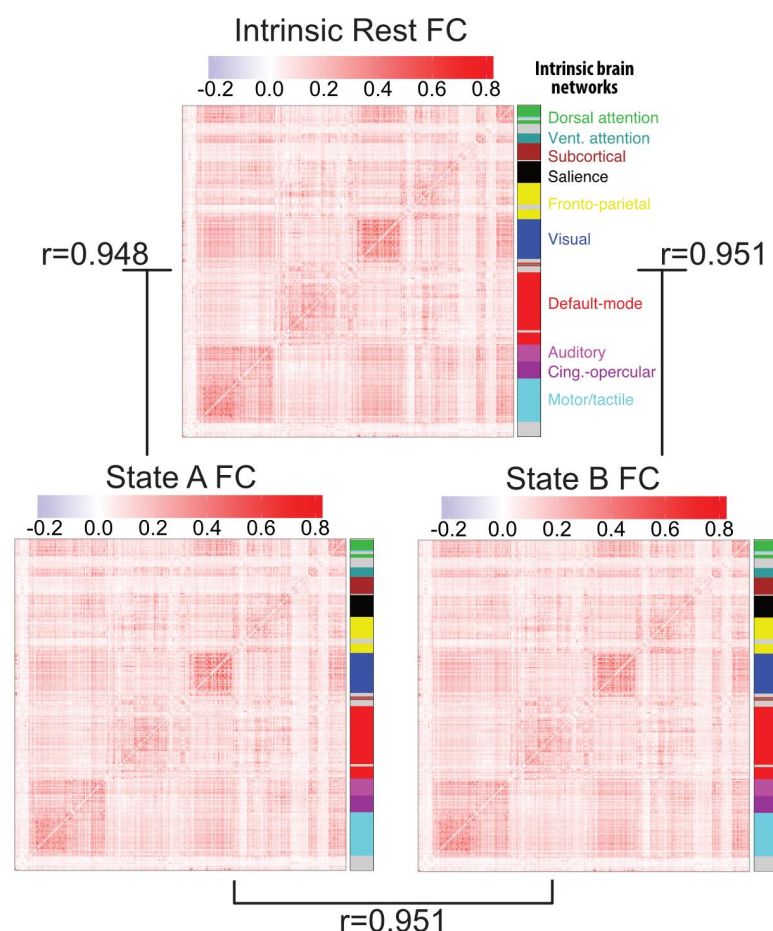


Figure 2 - Independence of activation state dynamics from FC state dynamics. The FC matrices for resting-state (over the entire rest run), State A (just the time points labeled state A), and State B (just the time points labeled state B) are pairwise correlated with each other for each subject. Averaged FC matrices and correlation values across subjects are shown here. These results suggest that the whole-brain multivariate activation states investigated here are independent of FC dynamics.

Functional relevance of the two states

Identifying a task-positive State A when using “task-negative” resting-state data suggests that subjects may have been performing covert tasks during “rest” periods. To test this hypothesis, we compared both State A and State B with rest and task data to better understand the functional significance of the two brain states. Specifically, we explored the temporal characteristics of the two states and how they related to ongoing task demands by correlating the two states’ prototypical activation patterns (Figure 1B) with every individual time point’s whole-brain activation pattern. This was done for all subjects individually in both resting-state and in each of the seven tasks (Figure 3B shows correlation results for the reasoning task). No across-subject temporal patterns were observed in the correlations with resting-state data (Figure 3A), as expected. On average, subjects were in State A only 39% of the time, and in State B 61% of the time. In contrast, for each of the seven tasks (Table 1), we found that every subject was in State A most of the time during task blocks and in State B during the inter-task rest periods (Figure 2B), with the exception of the language task (see Discussion). On average

(including the language task), subjects were in State A 54% (State B 46%) of the time during task blocks and in State A 41% (State B 59%) of the time during inter-task rest periods. Additionally, subjects tended to stay in State A longer for tasks that were likely more challenging. For instance, subjects were in State A 59% (State B 41%) for the reasoning task, but only in State A 49% (State B 51%) for the language task. This suggests that State A is not only relevant to rest data (the data used to derive it), but also likely to be the same state required for performing various active task demands.

Table 1 - Percent time in each state for each task. Percent of time in each state during task blocks and inter-task rest periods for each task.

Task Type	Task Blocks		Inter-task Rest Blocks	
	State A	State B	State A	State B
Reasoning	59%	41%	36%	64%
Emotion	52%	48%	40%	60%
Working Memory	53%	47%	39%	61%
Gambling	57%	43%	38%	62%
Language	49%	51%	50%	50%
Social	55%	45%	41%	59%
Motor	53%	47%	40%	60%
<i>Average</i>	<i>54%</i>	<i>46%</i>	<i>41%</i>	<i>59%</i>

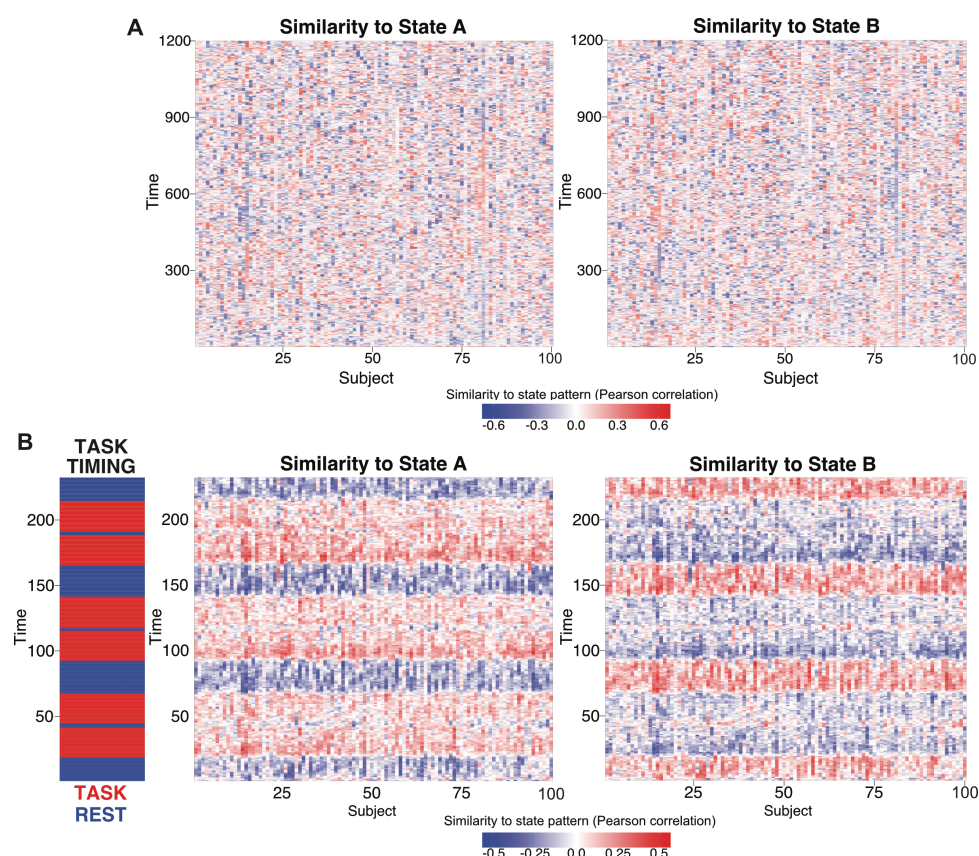


Figure 3 - Functional relevance of the two main states. The two brain states observed at the group level (see Figure 1B) are correlated with each subjects' rest and task data (separately for each time point). No discernable temporal organization of state A and state B is observed in resting-state data across all 100 subjects (A). The two subjects removed during the state identification analysis were included here for completeness. Evidence of state A during rest may suggest subjects are performing covert tasks during rest. Clear structure is seen when correlated with activity during the reasoning task (B). All subjects consistently enter state A when entering task blocks (in red), and enter state B when entering inter-task rest periods (in blue).

Hierarchical organization of brain states

The previous analyses suggest that State A and State B may be general and multi-functional states. Further, the vast fMRI literature indicates that each active task (cognitive state) has a unique activation pattern, despite an overarching "task-positive" activation pattern for externally oriented and cognitive control tasks (52–54). We therefore hypothesized brain states have a hierarchical organization, with state A and state B at the top level and more specific task states at lower levels (Figure 4). The two top states were broken down into more states by reducing the density of the adjacency matrix before applying Infomap clustering, from 100% to 20% density at decreasing increments of 20%. Reducing the density of the adjacency matrix removes weaker graph edges (lower correlations) between the group state prototypes. The task-positive state A divides into two states early in the hierarchy (60% density), shown by the red links between the levels of the hierarchy (Figure 4). The task negative state B splits into multiple states at the lowest level of the hierarchy (20% density), shown with blue links between the levels of the hierarchy.

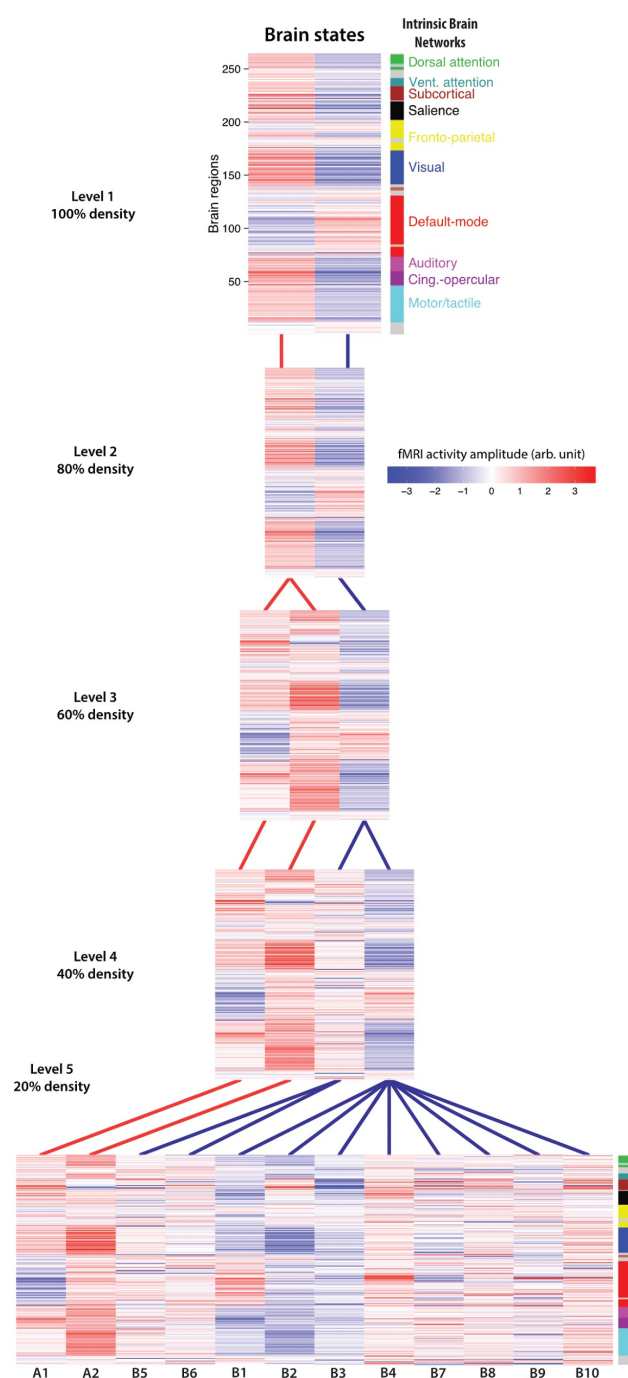


Figure 4 – Relationship between levels of the brain state hierarchy. The top level of the hierarchy (i.e., with no thresholding) is depicted at top, with additional states lower in the hierarchy with more stringent thresholding. Each level is defined by removing graph edges between the group prototype states (e.g., 80% density means removing the weakest 20% of edges) prior to running the community detection algorithm. The lines depict which state in the higher level that the lower-level state is most associated with. Specifically, a line indicates which higher-level state is composed of the largest percentage of the same group-level prototypes as the lower level state. Red links are associated with State A (task-positive) at the top level while blue links are associated with State B (task-negative) at the top level. The state labels below each of the states at 20% density (level 5) corresponds with the state labels in Figure 7.

To better understand the relationships between the individual prototypes that contribute to the state hierarchy, we visualized all 412 individual state prototypes in a force-directed graph (Figure 5). This helped characterize the state space underlying the states in the hierarchy. Each node in the graph represents an individual's brain state prototype. Edges connecting the nodes were weighted based on Pearson correlation scores between each pair of nodes, with stronger weights pulling nodes closer together. The correlation strengths are depicted on the graph with solid lines ($p < 0.01$) and dashed lines ($p < 0.05$). Negative correlations were excluded as possible edges on the graph. The nodes are colored based on the Infomap clustering assignments for the 100% density level (Figure 5A) and the 20% density level (Figure 5B). At the top level of the hierarchy, the centers of the two states consist of tightly clustered and strongly connected prototypes. However, at the 20% density level, the nodes located further away from the centers of State A and State B break off to form several clusters of prototypes, each constituting a unique brain state. Five states with the largest clusters are emphasized in Figure 5B. Two of these clusters (A1, A2) are from state A and three of the clusters (B1, B2, B3) are from state B.

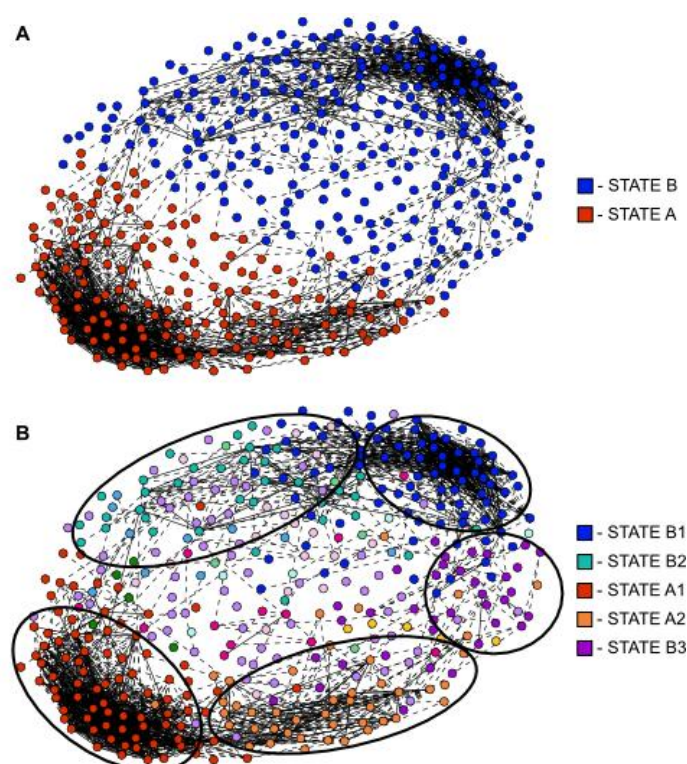


Figure 5 – Force directed graph visualization of brain states. All 412 group-level prototype brain states are depicted as nodes in a force-directed graph. More spatially similar states are closer together, while less similar states are forced apart. Solid lines indicate a significant spatial correlation between states at $p < 0.01$, while a dashed line indicates a significant correlation at $p < 0.05$. Panel **A** indicates the top-level solution (100% density), with red being the task-positive State A and blue being the task-negative State B. Panel **B** shows the lowest level of the hierarchy (20% density), with each color indicating a separate graph community. Large communities are highlighted with circles. Note that states appear to break off from the primary State A and State B clusters when they are mixtures of those two states (i.e., they were between the large State A and State B clusters).

Hierarchical organization in task fMRI

We next analyzed task-state fMRI brain states to test the hypothesis that task activations have a similar hierarchical structure as identified using resting-state fMRI data. Importantly, we could use a more supervised approach here, since the experimental cognitive manipulations

indicate (roughly) which mental state each subject is in. This involved assuming the experimental cognitive state corresponded to the brain activation state, then working backward to see if the activations were hierarchical in nature.

We first calculated the mean brain state prototypes for each of the seven tasks, shown in Figure 6A. This was done using a standard fMRI general linear model, with a canonical hemodynamic response function (given that all 7 tasks used block designs) and inter-block rest periods as baseline. The extremely high similarity across the seven distinct cognitive states observed ($r=0.9984$ on average) suggests that the highly similar brain states identified with resting-state fMRI may map to distinct cognitive states as well.

Next, we used principal component analysis (PCA; the first principal component) to isolate a dominant task-general activation state (Figure 6B middle), which was quite similar to the State A pattern identified using resting-state data ($r=0.56$, $p<0.0001$) but anti-correlated with State B ($r=-0.54$, $p<0.0001$). The task-general activation state was then removed from each of the seven mean task states using linear regression to obtain task-specific activation states (Figure 7B right). Each of the task-specific activation states included active regions associated with each cognitive task. For example (Figure 6B), after subtracting the task-general activation state from the motor task state, the remaining task-specific activation state included predominantly motor region activity and some attention region activity. Similar results were observed for the remaining six tasks as well (Figure 6C). The presence of a strongly dominant task-general state across diverse task demands, along with subordinate task-specific activation states across those tasks, indicates a hierarchical organization of the state space that is similar to what was observed with resting-state fMRI.

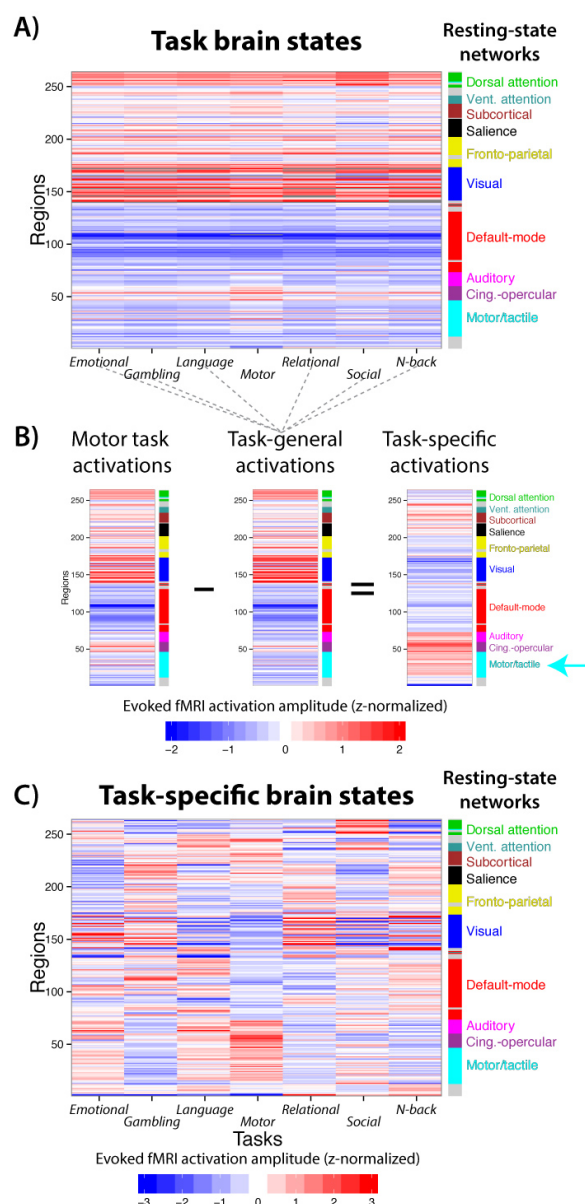


Figure 6 – Hierarchical organization of task activation brain states. **A)** The mean whole-brain activation patterns for each state are shown. These brain states were highly similar to each other on average ($r=0.9984$). The similarity of the brain states underlying these quite distinct mental states suggests the similarities observed across the resting-state brain states likely correspond to distinct mental states as well. **B)** We hypothesized that task activation states are also organized hierarchically. We used PCA to derive a task-general activation pattern – a higher-order brain state – which was correlated with State A ($r=0.56$, $p<0.0001$) but anti-correlated with State B ($r=-0.54$, $p<0.0001$). A lower-order task-specific brain state was revealed when the higher-order task-general brain state was subtracted from each task’s activation pattern. The motor task is shown for illustration. Motor/tactile network activations are not readily apparent in the original activation pattern, but are revealed when the task-general pattern is subtracted. Note that such task-specific patterns are typically observed using specific contrasts (e.g., left vs. right hand motor responses), whereas we use a task-general contrast here. **C)** Task-specific activation patterns for all 7 tasks are shown (average across-task $r=-0.01$). These results suggest important functional distinctions relevant to cognition and behavior are hierarchically embedded within the observed rest states despite dominance of State A and State B.

Decoding brain states using Neurosynth

We next sought to take advantage of the vast fMRI task activation literature to decode the identified resting-state brain states, identifying likely cognitive processes occurring during rest. This also allows us to test the hypothesis that State A and State B are likely to be general states governing basic modes of cognition, while the lower level states (12 states at 20% density) are likely more functionally specific. To examine the specific cognitive functions of each state at the 100% and 20% density levels, each of the 14 states (state A, state B, and the 12 states at 20%) were decoded using the Neurosynth decoder function and visualized as word clouds. The word clouds for state A and state B (Figure 7A) were consistent with our previous interpretation of a “task-positive” state and a “task-negative” state, respectively. Additionally, the cognitive terms for each of the lower level states (Figure 7B) are quite distinct from one another, despite certain terms repeating across states (“healthy” for state B2, B5). State A branched into two lower-level states that shared similar cognitive terms. In contrast, State B is subdivided into ten lower-level states that include more diverse cognitive terms spanning multiple functional domains not seen in State B word cloud. One possible explanation for the variability in cognitive terms observed could be that State B is also an internally-focused state for passive homeostatic functions and covert task executions, which could contribute to a wider spread of equally significant cognitive terms in the lower-level State B word clouds (see Discussion).

We validated the Neurosynth findings by running a permutation test for each of the brain states (see Methods for details). We found that all of the cognitive terms of interest depicted in State A and State B word clouds were significantly different when compared against the permutation tests’ null distributions ($p < 0.001$) for all cognitive terms (see Supporting Information Figure 2). Additionally, all cognitive terms for State A1, A2, B1, and B2 were significantly different ($p < 0.05$) when compared to the null distributions. Note that the permutation tests controlled for multiple comparisons, since all 3400 term comparisons were included in each permutation (see Methods). The remaining States B3-B10 (Figure 7B, in grey) did not survive the permutation tests. Together, these results support the accuracy of Neurosynth decoding of spontaneous activation states and suggest that the identified brain states correspond with distinct mental/cognitive states.

We identified a pair of competing anti-correlated states based on spontaneous activation patterns. State A appears to be a “task-positive” state, with high activation amplitudes in common task-active areas such as sensorimotor and cognitive control networks (28, 54, 60, 61). State B appears to be a “task-negative” state, with high activation patterns mainly in the default-mode network (29, 62). Breaking individual brain states into more clusters, we observed 12 distinct brain states that are more functionally specific than State A and B, as indicated by associated terms from the cognitive neuroscience literature. Qualitatively, the activation pattern for each brain state includes highly activated and/or de-activated brain regions that are affiliated with various functional network definitions (Figure 4 bottom), further supporting the specificity of each brain state. However, detecting the same states across unconstrained rest and task at lower tiers of the hierarchy would be unlikely, given that any two subjects are unlikely to enter the exact same mental state spontaneously. Additionally, reverse inferences are more likely to be problematic for lower-level (more specific) states, since reverse inferences are often used inaccurately when making overly broad generalizations about specific activations (63). Therefore, we focused predominantly on State A and State B, which were both present for 97 of the 98 subjects.

Relation to Resting-state FC

In this study, we emphasized the use of whole-brain activation patterns over alternative FC-based approaches (5, 64) for several reasons. Mainly, by focusing on activation amplitudes, we were able to directly compare the identified patterns with the large fMRI activation amplitude literature (15, 56). Also, characterizing activations allowed us to directly test whether subjects were in active task (externally oriented) brain states during “rest” periods, which the results support (Figure 7). Importantly, we found that FC architecture remains relatively unchanged across State A and State B despite significant differences in the underlying activation patterns (Figure 2). This suggests the activation states identified here are only weakly (if at all) related to FC states (5, 64). However, the observed activation states appear to be related to static resting-state FC networks (based on estimating FC across entire resting-state runs). For instance, the State B activation pattern is highly similar to the DMN resting-state network, though a portion of the FPN is also present. Several studies (19, 23) have investigated the relationship between static resting-state networks and co-activation patterns more directly. These studies showed that co-activation patterns can resemble common resting-state networks identified using traditional FC-based approaches, but some differences in the co-activation patterns remained undetected in FC-based analyses. Together, these observations support the use of activation patterns to investigate brain state organization.

Despite evidence that the present results are independent of FC dynamics, these results appear to be related to observed anti-correlation between DMN and task-positive network time series according to (static) resting-state FC estimates (53). That result has been recently questioned, however, given that it is dependent on the global signal regression preprocessing step (39). Importantly, we did not use global signal regression here, such that DMN time series were not anti-correlated with task-positive network time series (see Figure 2). This suggests the observed spatial (not temporal, which is used with FC) anti-correlation between DMN-dominated and task-positive-dominated states are at least somewhat independent of FC-based results. This may have been possible because spatial anti-correlation does not imply temporal anti-correlation (and vice versa). For example, two networks can be activated above baseline (high spatial correlation across time points) while being temporally anti-correlated at higher frequencies. It remains unclear exactly how spatial activation dynamics relate to time series dynamics, however, such that it will be important for future studies to fully characterize the relationship between spatial and temporal brain activity correlations in the future.

Brain Activation States Common Across Task and Rest

By using a data-driven approach across dozens of subjects we were able to obtain a comprehensive characterization of brain states across rest and task. The results confirmed our hypothesis that “task-positive” brain states occur regularly during rest periods. We specifically found that subjects spent the majority of rest in State B (61% of the time) but that a significant portion of rest was also spent in State A (39% of the time.) This ratio is reversed (as expected) during the seven tasks: subjects were in State A more often (54% of the time on average) during task periods. The only exception was the language task, where subjects were in State A only 49% of the time during task periods. One potential explanation is that the language task includes two trial types: a math task and a story task. It is possible that during the story task, subjects enter a self-reflective cognitive state consistent with State B, either due to the introspective nature of stories or due to a lack of active task demands, which may bias the percentages in favor of State B over State A during task periods.

A popular account of resting periods is that mind wandering is the primary mental phenomenon occurring during those times. Past studies have linked increased activity in the DMN, present in the State B activation pattern, with mind wandering (65, 66). However, it is likely that increased DMN activity is not the only neural mechanism underlying mind wandering. The significant presence of State A suggests that the brain is frequently performing active tasks, even during “rest”. A few studies have suggested personally relevant planning as one such task, which involves activity in the autobiographical memory system (67) and executive control systems (68, 69). Alternatively, common self-reported experiences of engaging in problem solving during rest (26) may also require increased activity in areas such as the FPN. Additionally, one might expect periodic activations in motor and sensory systems that are largely involved with passive attending to sensorimotor events (30). These “task-positive” networks are present in the State A activation pattern, suggesting State A might be important in a full explanation of mind wandering. However, future work will be required to assess the specific tasks that are performed during unconstrained resting-state mind wandering. Also of note, State B was present for a substantial portion of time during task performance. These results suggest there is a balance between these two states – with only moderate shifting from this homeostatic baseline – regardless of outward behavioral state.

Limitations

The present study involves several limitations that will be important for future studies to address. For instance, with fMRI the choice of a representative “baseline” for activation analyses is a complex issue. For standard task analyses, this is often circumvented using inter-task rest periods as the baseline to compare across task conditions. This choice has several issues, especially given that some regions are more active during rest than task (70). Yet even this imperfect baseline choice is unavailable when investigating activations during resting states. Building on common practice in the fMRI FC literature, we used each region’s (or voxel’s) time series mean as baseline. This is equivalent to using the spatial average across all whole-brain images in an fMRI run as the baseline, which removes brain anatomy and other potentially spurious similarities between brain images at individual time points. However, differences in the choice of baseline may influence the amplitude of the activations we observe, which in turn may affect the interpretability of activation patterns. Notably, this would be a problem for methods with a true baseline as well (e.g., positron emission tomography), given that without temporal de-meaning the spatial activation patterns could mostly reflect anatomy (and so would always be highly similar to each other). However, the majority of our analyses and findings are rooted in correlation-based methods, which are designed to remove the influence of baseline (or linear scaling) shifts in signals. Thus, our results may be less influenced by baseline issues than other alternative approaches. While the choice of a valid baseline remains an important problem to explore in fMRI research, its effect is likely minimized by the analytical approaches used in this study.

The dimensionality of the feature space, downscaling from approximately 60,000 grey-matter voxels to 264 brain regions, is another potential limiting factor for this study. While signal-to-noise is likely increased with a reduction in dimensionality (through averaging related signals within brain regions), meaningful information at the voxel level can also be lost. We focused on the scale of brain regions to take advantage of higher signal-to-noise, and also for computational tractability. Importantly, we found that the conclusions drawn at the region level generalized to the voxel level, based on mapping the time points assigned to each brain state to the voxel level (Figure 1C). It will be important, however, for future studies to perform our dMVPA approach (or related methods) directly on voxel patterns to assess potentially more detailed (but potentially more noisy) activation pattern brain states.

Conclusion

We used a novel dMVPA approach combining insights from multivariate activation methods as well as graph theoretical methods to identify and characterize activation pattern brain states during resting-state and task-state fMRI. This provided three primary benefits. First, relative to related M/EEG dMVPA approaches, the increased spatial accuracy of fMRI facilitated identification of activation patterns. Second, use of both resting-state and task-state data across 98 subjects allowed for an especially comprehensive identification of human brain activation states. Third, a focus on fMRI activation amplitude patterns allowed us to begin to interpret even spontaneous activations in terms of specific functions, given the vast task fMRI activation literature and associated cognitive functions. Together these benefits revealed a hierarchical organization of brain states with shared features across rest periods and task performance. It will be important for future studies to build on these results, improving our understanding of the many brain states that are entered spontaneously and as a result of social interactions (e.g., cognitive task instructions) along with their relation to ongoing brain network dynamics.

ACKNOWLEDGEMENTS

We would like to thank Patrick Shafto, Bharat Biswal, and Bart Krekelberg for helpful conversations during preparation of this manuscript. Data were provided by the Human Connectome Project, Washington University-Minnesota Consortium (principal investigators David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH institutes and centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. Our work was supported by NIH award K99-R00 MH096801 (M.W.C.). The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

SUPPORTING INFORMATION

Permutation test for statistical significance of the anti-correlation between State A and State B

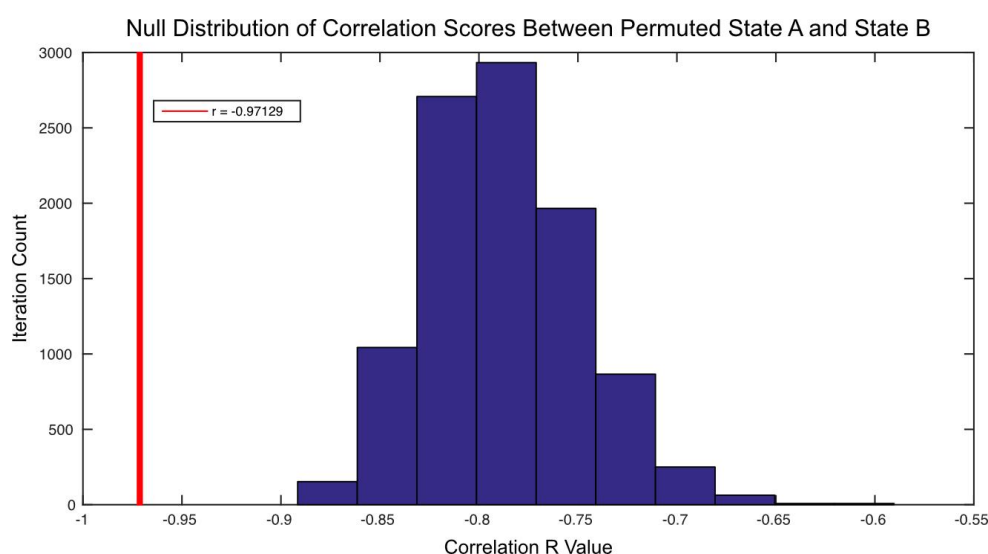
Methods

We sought to confirm the statistical significance of the observed anti-correlation by performing a permutation test with 10,000 iterations. For every iteration cycle, each subject's Infomap clustering solution was re-mapped into a two-state label solution based on the group-level clustering solution. The two-state labeling solution was then permuted for each timepoint prior to averaging all timepoints with the same label to create a "permuted" State A and State B. The

permuted State A and State B activation patterns were then correlated. The 10,000 r-values across all iterations formed the null distribution for comparing the observed anti-correlation between State A and State B.

Results

The null distribution derived from the permutation test followed a relatively normal distribution. Pseudo-randomly dividing the resting state data into two states tended to favor anti-correlated states, as seen in the negative mean of the null distribution. However, the observed anti-correlation State A and State B is statistically significant relative to this null distribution ($p < 0.0001$), supporting the distinctness of State A and State B.



Supporting Information Figure 1 - Permutation test for State A vs. State B anti-correlation. A total of 10000 permutations were performed to create the plotted null distribution. The red vertical line represents the observed anti-correlation r-value between State A and State B ($r = -0.97$). The negative null distribution suggests that pseudo-randomly splitting the data into two “states” tends to create anti-correlated activation patterns. However, the permutation test indicates that State A and State B were more distinct from one another than they would be by chance ($p < 0.0001$).

Permutation test for statistical significance of Neurosynth decoding for all states

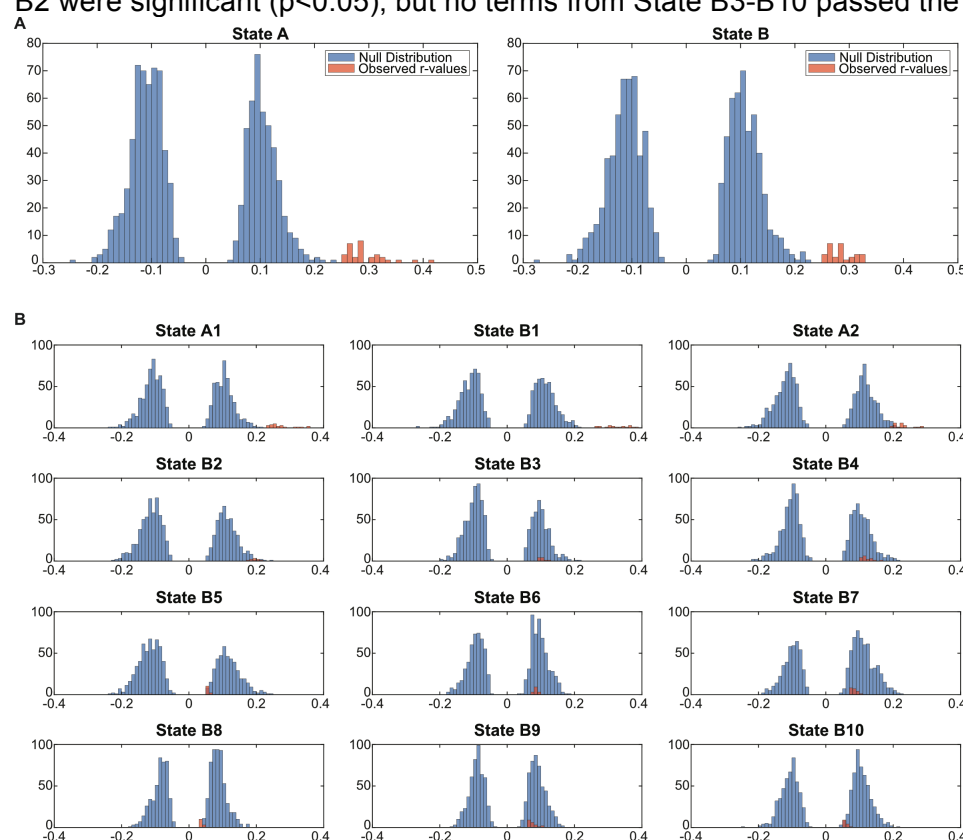
Methods

We validated the Neurosynth results by running permutation tests (1000 iterations) for each brain state (State A, State B, and the 12 states at 20% density tier of the hierarchy). For each permutation iteration, we permuted the order of the brain state temporal windows (consecutive time points of the same state), dictated by the Infomap clustering algorithm, for each individual subject. This method preserved the temporal properties of the brain states while still permuting the data being used to average and obtain the brain state prototypes. These newly permuted individual brain state prototypes were then averaged based on the group-level clustering solution to generate the permuted brain states (State A, State B, 12 states). These permuted states were passed through the Neurosynth decoder, comparing across all 3400 cognitive terms for each iteration of the 1000 permutations. The null distribution was formed by taking the

strongest correlation value (regardless of signage) across the 3400 comparisons for each of the 1000 permutations. Using this null distribution (size $n=1000$), we compared the observed correlation values for each cognitive term of interest in the word clouds (Figure 7) to test for statistical significance. Note that this approach allowed us to control for multiple comparisons, since all possible comparisons were computed for this permutation test (51).

Results

The null distribution for each state were formed by taking the highest correlation score across 3400 comparisons for every permutation cycle ($n=1000$). Each cognitive term of interest included in the respective word cloud (Figure 7) were tested against this null distribution for significance. Across all panels in Supporting Information Figure 2, the blue histograms represents the null distribution ($n=1000$) derived from the 1000 permutations and the red histograms illustrates all observed correlations scores for each cognitive term. All terms for State A and State B were significant ($p<0.001$). Additionally, all terms for State A1, A2, B1, and B2 were significant ($p<0.05$), but no terms from State B3-B10 passed the permutation test.



Supporting Information Figure 2 - Permutation test for Neurosynth decoding of State A and State B. The order of the state time windows were permuted per subject before averaging to obtain “permuted” prototypical states. A total of 1000 permutations per state were performed prior to Neurosynth decoding. The maximum (absolute value) r-value across the 3400 terms was selected for each permutation. The observed r-value for each cognitive term of interest was then compared to the null distribution formed from the 1000 permutations. A) Results for State A and State B are shown. Note that the blue histogram represents the null distribution and the red histogram represents correlation values for all terms included in the word cloud. All cognitive terms were significant ($p<0.001$) for both states. B) Results for the 12 states are displayed in the same manner as (A). All terms for A1, A2, B1, and B2 were significant ($p<0.05$) while no terms

from B3-B10 passed the permutation test. Note that multiple comparisons were corrected as part of the permutation test procedure (see Methods).

REFERENCES

1. Haynes J-D (2015) A Primer on Pattern-Based Approaches to fMRI: Principles, Pitfalls, and Perspectives. *Neuron* 87(2):257–270.
2. Haxby JV, Connolly AC, Guntupalli JS (2014) Decoding neural representational spaces using multivariate pattern analysis. *Annu Rev Neurosci* 37:435–456.
3. Fox MD, Raichle ME (2007) Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat Rev Neurosci* 8(9):700–711.
4. Smith SM, et al. (2011) Network modelling methods for FMRI. *Neuroimage* 54(2):875–891.
5. Hutchison RM, et al. (2013) Dynamic functional connectivity: Promises, issues, and interpretations. *Neuroimage* 80(C):360–378.
6. King J-R, Dehaene S (2014) Characterizing the dynamics of mental representations: the temporal generalization method. *Trends Cogn Sci* 18(4):203–210.
7. Yuan H, Zotev V, Phillips R, Drevets WC, Bodurka J (2012) Spatiotemporal dynamics of the brain at rest—exploring EEG microstates as electrophysiological signatures of BOLD resting state networks. *Neuroimage* 60(4):2062–2072.
8. Betzel RF, et al. (2012) Synchronization dynamics and evidence for a repertoire of network states in resting EEG. *Front Comput Neurosci* 6:74.
9. Anderson JR, Fincham JM, Schneider DW, Yang J (2012) Using brain imaging to track problem solving in a complex state space. *Neuroimage* 60(1):633–643.
10. Junejo I (2010) Learning Self-Similarities for Action Recognition Using Conditional Random Fields. *Bayesian Network*, ed Rebai A (Sciyo). doi:10.5772/46965.
11. Furusawa C, Kaneko K (2012) A dynamical-systems view of stem cell biology. *Science* 338(6104):215–217.
12. Newman MEJ, Girvan M (2004) Finding and evaluating community structure in networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 69(2 Pt 2):026113.
13. Smith SM, et al. (2009) Correspondence of the brain’s functional architecture during activation and rest. *Proceedings of the National Academy of Sciences* 106(31):13040–13045.
14. Smith SM, et al. (2012) Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences* 109(8):3131–3136.

15. Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods* 8(8):665–670.
16. Laird AR, et al. (2009) ALE Meta-Analysis Workflows Via the Brainmap Database: Progress Towards A Probabilistic Functional Brain Atlas. *Front Neuroinform* 3:23.
17. Bzdok D, et al. (2016) Formal Models of the Network Co-occurrence Underlying Mental Operations. *PLoS Comput Biol* 12(6):e1004994.
18. Liu X, Duyn JH (2013) Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proc Natl Acad Sci U S A* 110(11):4392–4397.
19. Liu X, Xiao L, Catie C, Duyn JH (2013) Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. *Front Syst Neurosci* 7. doi:10.3389/fnsys.2013.00101.
20. Khanna A, Pascual-Leone A, Michel CM, Farzan F (2015) Microstates in resting-state EEG: Current status and future directions. *Neurosci Biobehav Rev* 49:105–113.
21. Grech R, et al. (2008) Review on solving the inverse problem in EEG source analysis. *J Neuroeng Rehabil* 5(1):25.
22. Schoffelen J-M, Gross J (2009) Source connectivity analysis with MEG and EEG. *Hum Brain Mapp* 30(6):1857–1865.
23. Chen JE, Chang C, Greicius MD, Glover GH (2015) Introducing co-activation pattern metrics to quantify spontaneous brain network dynamics. *Neuroimage* 111:476–488.
24. 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(1):253–258.
25. Buckner RL, Vincent JL (2007) Unrest at rest: default activity and spontaneous network correlations. *Neuroimage* 37(4):1091–6; discussion 1097–9.
26. Smallwood J, Schooler JW (2006) The restless mind. *Psychol Bull* 132(6):946–958.
27. Anticevic A, et al. The role of default network deactivation in cognition and disease. 16(12):584–592.
28. Cole MW, Schneider W (2007) The cognitive control network: Integrated cortical regions with dissociable functions. *Neuroimage* 37(1):343–360.
29. Raichle ME, et al. (2001) A default mode of brain function. *Proc Natl Acad Sci U S A* 98(2):676–682.
30. Fox MD, Corbetta M, Snyder AZ, Vincent JL, Raichle ME (2006) Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proc Natl Acad Sci U S A* 103(26):10046–10051.
31. Van Essen DC, et al. (2013) The WU-Minn Human Connectome Project: an overview. *Neuroimage* 80:62–79.
32. Barch DM, et al. (2013) Function in the human connectome: task-fMRI and individual

- differences in behavior. *Neuroimage* 80:169–189.
33. Uğurbil K, et al. (2013) Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. *Neuroimage* 80:80–104.
 34. Smith SM, et al. (2013) Functional connectomics from resting-state fMRI. *Trends Cogn Sci* 17(12):666–682.
 35. Glasser MF, et al. (2013) The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* 80:105–124.
 36. Cox RW (1996) AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res* 29(3):162–173.
 37. Fischl B, et al. (2002) Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33(3):341–355.
 38. Fischl B, et al. (2004) Sequence-independent segmentation of magnetic resonance images. *Neuroimage* 23 Suppl 1:S69–84.
 39. Murphy K, et al. (2009) The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced? *Neuroimage* 44(3):893–905.
 40. Power JD, et al. (2011) Functional network organization of the human brain. *Neuron* 72(4):665–678.
 41. Wig GS, Schlaggar BL, Petersen SE (2011) Concepts and principles in the analysis of brain networks. *Ann N Y Acad Sci* 1224:126–146.
 42. Cohen AL, et al. (2008) Defining functional areas in individual human brains using resting functional connectivity MRI. *Neuroimage* 41(1):45–57.
 43. Cha S-H (2007) Comprehensive survey on distance/similarity measures between probability density functions. *Cityscape* 1(2):1.
 44. Ahlgren P, Jarneving B, Rousseau R (2003) Requirements for a cocitation similarity measure, with special reference to Pearson's correlation coefficient. *J Am Soc Inf Sci* 54(6):550–560.
 45. Bohlin L, Ludvig B, Daniel E, Andrea L, Martin R (2014) Community Detection and Visualization of Networks with the Map Equation Framework. *Measuring Scholarly Impact*, pp 3–34.
 46. Rosvall M, Axelsson D, Bergstrom CT (2009) The map equation. *Eur Phys J Spec Top* 178(1):13–23.
 47. Rosvall M, Bergstrom CT (2011) Multilevel compression of random walks on networks reveals hierarchical organization in large integrated systems. *PLoS One* 6(4):e18209.
 48. Lancichinetti A, Fortunato S (2012) Consensus clustering in complex networks. *Sci Rep* 2:336.
 49. Rubinov M, Sporns O (2011) Weight-conserving characterization of complex functional

- brain networks. *Neuroimage* 56(4):2068–2079.
50. Blondel VD, Jean-Loup G, Renaud L, Etienne L (2008) Fast unfolding of communities in large networks. *J Stat Mech: Theory Exp* 2008(10):P10008.
 51. Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp* 15(1):1–25.
 52. Power JD, et al. (2011) Functional network organization of the human brain. *Neuron* 72(4):665–678.
 53. Fox MD, et al. (2005) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci U S A* 102(27):9673–9678.
 54. Corbetta M, Shulman GL (2002) CONTROL OF GOAL-DIRECTED AND STIMULUS-DRIVEN ATTENTION IN THE BRAIN. 3(March). doi:10.1038/nrn755.
 55. Uddin LQ, Kelly AM, Biswal BB, Castellanos FX, Milham MP (2009) Functional connectivity of default mode network components: correlation, anticorrelation, and causality. *Hum Brain Mapp* 30(2):625–637.
 56. Laird AR, et al. (2013) Networks of task co-activations. *Neuroimage* 80:505–514.
 57. Duncan J, Owen AM (2000) Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci* 23(10):475–483.
 58. Dosenbach NUF, Fair DA, Cohen AL, Schlaggar BL, Petersen SE (2008) A dual-networks architecture of top-down control. *Trends Cogn Sci* 12(3):99–105.
 59. Kriegeskorte N, Simmons WK, Bellgowan PSF, Baker CI (2009) Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci* 12(5):535–540.
 60. Dosenbach NUF, et al. (2007) Distinct brain networks for adaptive and stable task control in humans. *Proc Natl Acad Sci U S A* 104(26):11073–11078.
 61. Raichle ME (2010) Two views of brain function. *Trends Cogn Sci* 14(4):180–190.
 62. Buckner RL, Andrews-Hanna JR, Schacter DL (2008) The Brain's Default Network: Anatomy, Function, and Relevance to Disease. *Ann N Y Acad Sci* 1124(1):1–38.
 63. Poldrack RA, Wager TD (2006) Introduction to “Tools of the Trade.” *Soc Cogn Affect Neurosci* 1(1):72–72.
 64. Allen EA, et al. (2014) Tracking whole-brain connectivity dynamics in the resting state. *Cereb Cortex* 24(3):663–676.
 65. Mittner M, Matthias M, Hawkins GE, Wouter B, Forstmann BU (2016) A Neural Model of Mind Wandering. *Trends Cogn Sci*. doi:10.1016/j.tics.2016.06.004.
 66. Mason MF, et al. (2007) Wandering Minds: The Default Network and Stimulus-Independent Thought. *Science* 315(5810):393–395.
 67. Baird B, Smallwood J, Schooler JW (2011) Back to the future: autobiographical planning

and the functionality of mind-wandering. *Conscious Cogn* 20(4):1604–1611.

68. Christoff K, Gordon AM, Smallwood J, Smith R, Schooler JW (2009) Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proc Natl Acad Sci U S A* 106(21):8719–8724.
69. Smallwood J, Brown K, Baird B, Schooler JW (2012) Cooperation between the default mode network and the frontal-parietal network in the production of an internal train of thought. *Brain Res* 1428:60–70.
70. Stark CE, Squire LR (2001) When zero is not zero: the problem of ambiguous baseline conditions in fMRI. *Proc Natl Acad Sci U S A* 98(22):12760–12766.