nature neuroscience

Activity flow over resting-state networks shapes cognitive task activations

Michael W Cole[1], Takuya Ito[1], Danielle S Bassett[2,3] & Douglas H Schultz[1]

Resting-state functional connectivity (FC) has helped reveal the intrinsic network organization of the human brain, yet its relevance to cognitive task activations has been unclear. Uncertainty remains despite evidence that resting-state FC patterns are highly similar to cognitive task activation patterns. Identifying the distributed processes that shape localized cognitive task activations may help reveal why resting-state FC is so strongly related to cognitive task activations. We found that estimating task-evoked activity flow (the spread of activation amplitudes) over resting-state FC networks allowed prediction of cognitive task activations in a large-scale neural network model. Applying this insight to empirical functional MRI data, we found that cognitive task activations can be predicted in held-out brain regions (and held-out individuals) via estimated activity flow over resting-state FC networks. This suggests that task-evoked activity flow over intrinsic networks is a large-scale mechanism explaining the relevance of resting-state FC to cognitive task activations.

The neural basis of cognition has been primarily investigated in terms of task-evoked activation level changes. Over the past decade a separate focus on spontaneous (non-task-evoked) activity has challenged cognitive neuroscientists' focus on task-evoked activations[1–3]. Due to the lack of experimental control of spontaneous brain activity there has been a strong emphasis on discovering correlations (rather than activation level changes) among activity time series, an approach termed resting-state FC. Thus the theoretical framework and methodological approaches associated with cognitive task activations and resting-state FC are highly distinct, leading to a bifurcation in investigations of brain function.

Notably, this bifurcation mirrors the classic 'localized' versus 'distributed' neural processing debate[4–7], such that the relationship between localized cognitive task activations and distributed FC is also relevant to this broader theoretical divide in neuroscience. Here we sought to identify where the human brain lies with respect to these two extremes. We focused in particular on the role of intrinsic functional networks (as estimated by resting-state FC) in distributed processing. There is evidence that resting-state FC patterns are similar to cognitive task activation patterns[8–10], but we sought here to quantify this relationship using a large-scale mechanistic construct that may help explain why this relationship exists. Critically, we recently found that the FC architectures across a variety of tasks were highly similar (80% shared variance) to the resting-state FC architecture[11]. This suggests that the functional network architecture identified using resting-state FC is present during task performance and could plausibly reflect the routes by which activity flows during cognitive task performance. However, it remains unclear whether and how these FC patterns relate to cognitive task activation amplitudes—such as task-evoked blood oxygen level dependent (BOLD) functional MRI (fMRI) signal increases—and therefore how they relate to cognition.

We sought to answer these questions by testing whether estimated activity flow over resting-state FC networks can accurately predict cognitive task activations in held-out regions. Activity flow (often termed 'information flow') is the spreading of activation amplitudes between brain locations, such as task-evoked activations spreading from visual cortex to motor cortex in a visual–motor task. Decades of findings in local circuits and simulations have suggested that connectivity and activations are strongly interrelated neurophysiological variables, with activity flow as a key linking variable[12–14]. However, little is known about how FC and cognitive task activations relate at the large-scale network level, for example, as measured with fMRI. Beginning to fill this gap, several recent studies used abstract statistical models to predict cognitive task activations based on individual differences in large-scale connectivity[9,15,16]. We sought to build on these findings to identify why these predictions were possible. This involved testing the plausibility of a (large-scale) mechanistic relationship between connectivity and cognitive task activations in terms of the concept of activity flow.

Local circuit-level studies have suggested that task-evoked activation at a given location is primarily determined by activity flow from other neurons[17,18]. Activity flow is carried (via axons) by action potentials modulated by synaptic strengths (among other modulators). Thus activity flow is a mechanism that emerges from several more basic mechanisms. One can conceptualize activity flow as relating activations (action potentials and associated local field potentials) and functional pathways (their tendency to influence one another via, for example, synaptic strengths). Applied to large-scale measures of the human brain, we hypothesized that aggregate activation amplitudes

[1]Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, New Jersey, USA. [2]Department of Bioengineering, University of Pennsylvania, Philadelphia, Pennsylvania, USA. [3]Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, Pennsylvania, USA. Correspondence should be addressed to M.W.C. (mwcole@mwcole.net).

Received 18 April; accepted 7 September; published online 10 October 2016; doi:10.1038/nn.4406

(for example, as measured by BOLD fMRI signal) flow among brain regions via functional pathways (possibly reflecting, in part, aggregate synaptic strengths) described by FC. Thus we conceptualized activity flow as a linking variable between large-scale FC and cognitive task activations that could be used to demonstrate (and quantify) the functional relevance of these two measures to one another.

We tested the plausibility of this hypothesis by constructing activity flow mappings using FC and task activations. This involved predicting the cognitive task activation level at one location based on the FC-weighted sums of the activations at other locations (i.e., the sum of activity-flow estimates; **Fig. 1a**). We then repeated this process separately for each brain region, an approach similar to leave-one-out cross-validation from machine learning[19–21]. This resulted in a whole-brain activation pattern prediction, which could be compared with a given task's actual fMRI activation pattern. A successful prediction (i.e., high correspondence between predicted and actual activation patterns) would indicate the plausibility of resting-state FC pathways in shaping the empirically observed activation pattern. Further, successful prediction across a variety of tasks and subjects would indicate the general plausibility of the activity flow framework at the large-scale network level, suggesting that resting-state FC is relevant to cognitive task activations due to its role in shaping task-evoked activity flow among brain regions.

There are several reasons why activity-flow-based prediction of cognitive task activations is not guaranteed to work. For instance, cognitive task activations may be largely shaped by task-evoked network reconfigurations[18,22,23], making prediction of cognitive task activations by resting-state FC ineffective. Additionally, localized processing independent of other brain regions could be a major driver of cognitive task activations in any given brain region, such that activity flow is largely irrelevant to localized cognitive task activations. Indeed, many cognitive task activations have been interpreted under this assumption[7], such as task-evoked activations within dorsolateral prefrontal cortex during working memory maintenance[24]. Even with strong evidence that activity flow shapes activations at the local circuit level[12–14], this is not guaranteed at the large-scale network level since local processing (for example, within-region activity flow) is likely to be at least partially independent of the large-scale activity flow into a region. The spatiotemporal resolution of fMRI signals is several orders of magnitude lower than neuron-level events, leaving room for extensive local processing to occur independently of large-scale activity flow into and out of a given region.

The activity flow mapping approach is based on the local circuit-level findings described above. However, like recent models of activity-spreading dynamics[25,26], it is not meant to be a realistic simulation of neuronal dynamics but rather a tool for quantifying (and making inferences about) brain activity relationships. We see the present study as a precursor to more complex approaches that incorporate biophysical models of neuronal communication[27,28] to improve activation pattern predictions further. Here we sought to make as few assumptions as possible regarding the biophysical basis of the FC–activation relationship by using the simplest activity flow mapping approach possible: the FC-weighted sum of activations. This allowed us to make straightforward inferences regarding the relationship between FC and cognitive task activations, which future work can refine using more elaborate models of neuronal communication.

We began by validating the activity flow mapping procedure with a simple computational model of large-scale neural interactions. We then applied the activity flow mapping approach to empirical fMRI data acquired as healthy adult human participants (n = 100) rested and performed a variety of tasks. We used activity flow mapping to test our primary hypothesis: cognitive task activations can be predicted in held-out brain regions (and held-out individuals) via estimated activity flow over resting-state FC networks. This would suggest that task-evoked activity flow over intrinsic networks (i.e., the spread of activation amplitudes between regions) acts as a large-scale mechanism helping to explain the functional relevance of resting-state FC to cognitive task activations.

RESULTS

Computational validation and identification of factors contributing to cognitive task activations

Previous research has shown a statistical relationship between resting-state FC and cognitive task activations[8,9] but not why this relationship exists. We recently found that resting-state FC patterns are present during cognitive task performance (80% shared variance in whole-brain FC patterns between rest and task)[11]. This suggests that resting-state FC might describe activity flow among brain regions even during task performance. Here we tested this possibility in the context of task-evoked activation amplitudes, using activity flow among brain regions as a linking variable between resting-state FC and task-evoked activations. This involved modeling activity flow as task activation amplitudes (standard fMRI general linear model estimates) multiplied by FC strengths (standard Pearson correlations) between brain regions (**Fig. 1a**). Standard measures were used to maximally relate to the existing resting-state FC and cognitive task activation literatures. We hypothesized that this would allow us to predict cognitive task activations in held-out brain regions based on resting-state FC patterns.

We began by validating this activity flow mapping procedure with a simple computational model of large-scale neural interactions. The model was kept simple to reduce the number of assumptions regarding underlying biophysical detail (Online Methods). Interactions among 300 brain regions were simulated along with task-evoked activations (**Fig. 1b**). Knowing the ground truth connectivity and activations in the model (since we defined these parameters ourselves) allowed us to validate the activity flow mapping procedure.

We constructed the model to have three structural network communities, with the first community split into two 'functional' communities via modulation of synaptic strengths (**Fig. 1c**). This was of particular interest here given the potential for resting-state FC fMRI (unlike, for example, diffusion-weighted MRI) to detect the aggregate effects of synaptic strengths that are known to modify activity flow over structural (axonal) connections in local circuits[12]. Note that although we focused on synaptic strengths, there are other modulators of FC, such as large-scale changes in neurotransmitter concentrations. We then ran the model with spontaneous activity (Gaussian random values) in each unit while simulating fMRI data collection (Online Methods). We then computed Pearson correlations among all of the time series to produce simulated resting-state FC data (**Fig. 1d**).

We next simulated task-evoked activations by injecting stimulation (2 min of stimulation in 3 blocks) into five neighboring regions at a time. Six 'tasks' were simulated by changing the stimulated regions (**Fig. 1e**). We simulated fMRI data collection as with the 'rest' data, followed by application of a standard fMRI general linear model to obtain activation amplitude estimates for each simulated region (**Fig. 1e**). We then implemented the activity flow mapping algorithm, assessing its ability to predict task activations in held-out regions based on resting-state FC. We found that activity flow mapping was successful in recovering the original task-evoked activation pattern (across-task average r = 0.56, $P < 0.00001$; cross-task average Spearman's rank correlation ρ = 0.51, $P < 0.00001$).

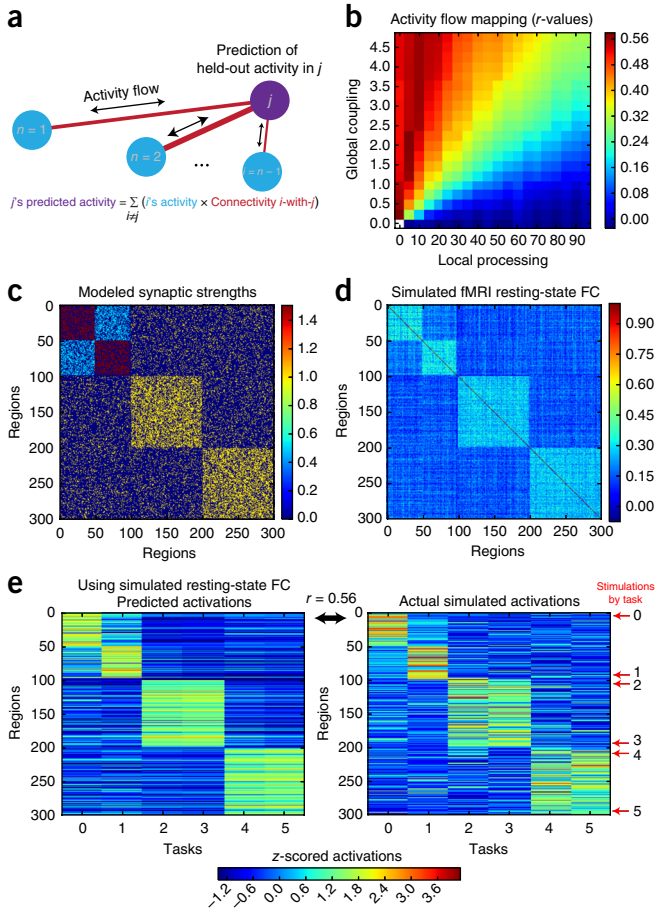

Figure 1 Activity flow mapping over resting-state FC networks allows prediction of held-out task activations. (**a**) Diagram showing our approach—linking resting-state FC (red lines) to task activations (circles) to assess the relevance of FC to cognitive task activations—to predicting a single region's activation amplitude for a single task. The to-be-predicted region's (purple circle) activation amplitude is held out from the prediction calculation. (**b**) Whole-brain predicted-to-actual Pearson correlations (*r*-values) for distinct model parameters are shown. The success of activity flow mapping depended on the relative degree of local (within-region, recurrent) vs. distributed (cross-region) processing. (**c**) Three structural connectivity graph communities (blocks along diagonal) were created, with the first split into two communities via synaptic strength modifications. (**d**) Resting-state FC (Pearson correlation) was computed based on simulated time series using the computational model, revealing a strong correspondence with the underlying synaptic strengths. (Note that other factors not modeled here likely also influence resting-state FC.) Global coupling and local processing parameters were set to 1.0 for this example. (**e**) Simulated task-evoked activations were produced by stimulating (red arrows) groups of 5 nearby units in 6 separate simulated 'tasks'. Activity flow mapping produced above-chance recovery (mean cross-task Pearson correlation $r = 0.56$, $t_{298} = 11.7$, $P < 0.00001$) of the actual activations using the resting-state FC matrix shown in **d**.

To ensure robustness of this result we repeated the entire simulation procedure 4,000 times (Online Methods). Over these iterations we varied a global coupling parameter to assess the role of aggregate synaptic strengths (and/or other potential modulators of neuronal communication), along with varying a local processing parameter to assess the role of nondistributed (local) activity. Global coupling was defined as a constant that linearly scaled all synaptic strengths, while local processing was defined as a constant that linearly scaled all self-connection (recurrent connection) strengths. We found that activity flow mapping worked to the extent that global coupling was high and local processing was low (**Fig. 1b**). The sensitivity of these results to the local–distributed processing relationship suggested that empirical assessment of activity flow mapping with real fMRI data would be nontrivial, in the sense that it would only be effective if the empirical data displayed certain properties. Further, these results suggested that activity flow mapping could provide evidence regarding the relative distributed versus localized processing that occurs in the human brain during cognitive task performance.

Activity flow mapping with empirical fMRI data

We next applied the activity flow mapping approach to empirical fMRI data, testing the hypothesis that cognitive task activations could be predicted in held-out brain regions via estimated activity flow over resting-state FC networks. This involved applying activity flow mapping to a Human Connectome Project data set involving rest and seven highly distinct tasks[29]. These tasks included an 'emotional' processing task, a 'gambling' incentive processing task, an auditory 'language' task involving stories and math problems, a simple 'motor' task (involving moving fingers, toes, and tongue), a 'social' cognition task, a 'relational reasoning' task, and an 'N-back' working memory task. A standard set of functionally defined brain regions[22,30] (**Fig. 2a**) was used along with standard Pearson correlation-based FC measures, calculated pairwise across all regions (**Fig. 2b**). The predicted activation pattern matrix was highly similar to the actual activation pattern matrix: cross-task average $r = 0.48$, $t_{99} = 39.29$, $P < 0.00001$. The *r*-values were similar for each of the seven tasks individually: $r = 0.42$ (emotional), $r = 0.49$ (gambling), $r = 0.46$ (language), $r = 0.53$ (motor), $r = 0.49$ (reasoning), $r = 0.50$ (social) and $r = 0.45$ (N-back).

These correlations were higher when comparisons were computed after averaging the predicted and actual activation patterns across subjects (cross-task average $r = 0.66$, 44% variance explained), likely due to an improved signal-to-noise ratio from aggregating more data before comparison (**Fig. 2c**). This was true of each of the seven tasks individually: $r = 0.67$ (emotional), $r = 0.66$ (gambling), $r = 0.66$ (language), $r = 0.65$ (motor), $r = 0.68$ (reasoning), $r = 0.65$ (social) and $r = 0.65$ (N-back). Note that these average-then-compare results likely better reflect the true effect sizes (due to better signal-to-noise ratios), while the compare-then-average results better demonstrate the consistency of the effects across subjects.

Results were similar when using global signal regression during preprocessing (average $r = 0.50$ across the seven tasks, $t_{99} = 45.32$, $P < 0.00001$). Further, it should be noted that all seven tasks used block designs and that future research investigating the efficacy of activity flow mapping with event-related relative to block designs will be important. Also note that similarity was high between tasks in the actual activation patterns relative to rest (**Supplementary Fig. 1a**), consistent with previous meta-analyses[31–34] and suggesting the existence of a 'task-general' activation pattern. We therefore conceptualized a given task-activation pattern as being composed of a task-general pattern and a task-specific pattern (**Supplementary Fig. 1b**). We applied activity flow mapping on isolated task-specific activations (Online Methods), allowing us to identify the role of activity flow in shaping task-specific activations (for example, motor network activations during the motor task).

These results demonstrate the plausibility of activity flow as a large-scale linking mechanism between resting-state FC and activations across a variety of distinct cognitive tasks. Further, these results suggest a strong role for large-scale distributed (rather than primarily local) processing in the human brain, establishing the relevance of resting-state FC to understanding cognitive task activations.

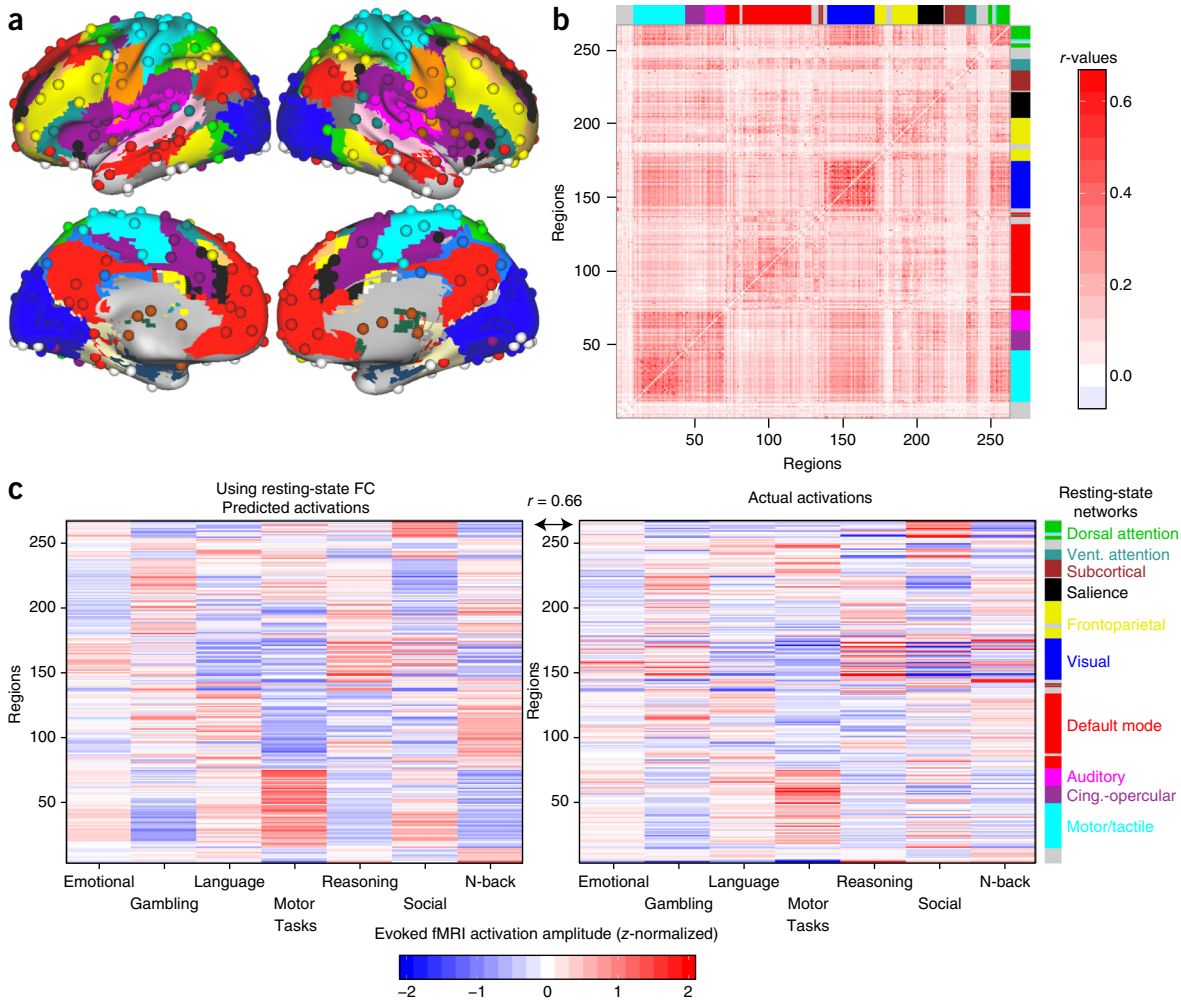

Figure 2 Activity flow mapping predicts cognitive task activations with empirical fMRI data. (**a**) We used a standard set of functionally defined regions of interest with associated resting-state FC network assignments[30]. (**b**) The cross-subject average resting-state FC matrix (Pearson correlations) among the 264 regions shown in **a**. Results were similar when using global signal regression during preprocessing (Online Methods). Colors in **a** and **b** indicate networks as named in **c**. (**c**) Pearson correlation-based resting-state FC was used to predict activation patterns across the 7 tasks (mean activity amplitude of each region for each task). Cross-task activation similarities were removed to emphasize task-specific activations (**Supplementary Fig. 1**). The high correspondence between predicted (left) and actual (right) activation patterns (average prediction accuracies $r = 0.66$, based on cross-subject mean predicted and actual activations) suggests resting-state FC shapes activity flow in task contexts.

Improving activity flow mapping predictions using multiple regression

We used Pearson correlations to this point due to their prominent role in the resting-state FC fMRI literature. However, multiple regression is a standard measure for making predictions of a single variable based on many other variables, which was the goal of the activity flow mapping approach. We reasoned that this might produce better predictions since, relative to correlation, multiple regression would reflect more direct FC relationships between regions (accounting for, for example, signals passing through a third region). Further, unlike correlation (which is an abstract statistical measure), multiple regression would scale FC values to the same units as activity in each to-be-predicted region during rest, likely producing more accurately scaled predictions. We therefore adapted the activity flow mapping approach to use multiple regression in place of Pearson correlation. This involved calculating resting-state FC using a standard linear regression model (i.e., a general linear model) for each region, with all other regions as predictor variables. Each regression coefficient in the resulting FC matrix represents how much a given source region's activity must be scaled (statistically controlling for all other source regions) to match the activity amplitude of a given target region during resting state.

Using this new FC matrix substantially improved activity flow mapping predictions: cross-task average $r = 0.69$, $t_{99} = 46.18$, $P < 0.00001$. The r-values were similar for each of the seven tasks individually: $r = 0.61$ (emotional), $r = 0.72$ (gambling), $r = 0.64$ (language), $r = 0.73$ (motor), $r = 0.72$ (reasoning), $r = 0.75$ (social) and $r = 0.64$ (N-back). The average correlation was higher when comparisons were computed after averaging predicted and actual activation patterns across subjects (cross-task average $r = 0.91$; 83% of variance), likely due to improved signal-to-noise ratios (**Fig. 3**). This was true of each of the seven tasks individually: $r = 0.86$ (emotional), $r = 0.93$ (gambling), $r = 0.88$ (language), $r = 0.93$ (motor), $r = 0.94$ (reasoning), $r = 0.92$ (social) and $r = 0.86$ (N-back).

These results demonstrate the utility of using multiple regression rather than Pearson correlation in the context of activity flow mapping. Further, the high correlations obtained further support the possibility that activity flow over intrinsic networks (as estimated by resting-state FC) strongly shapes cognitive task activations.

Additional validation of the activity flow mapping approach

The activity flow mapping approach assumes that the observed prediction accuracies are dependent on the particular organization of the FC network architecture. To test this assumption we randomly permuted FC patterns across regions (10,000 permutations; Online Methods). We found that the original result was highly dependent on each region's particular FC pattern. Of the 10,000 permutations, the highest Pearson correlation *r*-value between predicted and actual activity was $r = 0.024$. This indicates that the nonparametric permutation test *P*-value for the original multiple-regression FC result ($r = 0.69$) was $P < 0.0001$. **Supplementary Figure 2a** depicts a prediction based on an example permutation, while **Supplementary Figure 2b** visually illustrates the null distribution created for the permutation test. Similar results were obtained with Pearson correlation FC (highest value from 10,000 permutations: $r = -0.009$).

Another assumption of the activity flow mapping approach is that—outside the hypothesized activity flow mechanism—resting-state FC and task activations are largely independent. Such independence is likely because the resting-state fMRI data were collected during separate runs from the task fMRI data. However, there may be vasculature-based or other fMRI-related confounds consistent across both the resting-state and task runs that link resting-state FC and task activations. If such confounds exist (for example, effects of differential signal-to-noise ratio across regions), they would likely result in region-to-region correlations in activation amplitudes across rest and task. We therefore tested for correlations between task-specific activation amplitudes and resting-state amplitude of low-frequency fluctuation values, a standard measure of resting-state activation amplitudes. We found no region-to-region correlation between resting-state and task-specific activation amplitudes: cross-task average $r = -0.005$, $t_{99} = -1.30$, $P = 0.20$. This suggested independence of the resting-state and task activation amplitudes at least as far as task-specific activations are concerned.

Another possibility is that regions with overall stronger resting-state FC tend to have higher task-activation amplitudes. While not strictly incompatible with the proposed activity-flow network mechanism, this could create a situation in which resting-state FC could predict activation amplitudes without the activity flow mapping procedure. We calculated the overall (sum) resting-state multiple-regression FC for each region and tested for a correlation with task-activation amplitudes. Note that this is equivalent to running the activity flow mapping procedure with resting-state FC values alone (no task-activation amplitudes). We found that there was only a very small (but significant) negative correlation between summed resting-state FC and task-specific activation amplitudes: average $r = -0.01$, $t_{99} = -4.46$, $P = 0.00002$. This demonstrated that increased resting-state FC did not correspond with increased task-specific activations, meaning this effect could not have driven the observed activity flow mapping results. Together these results further support the assumptions underlying the activity flow mapping approach.

Voxelwise activity flow with empirical fMRI data

We next sought to test whether resting-state FC describes the routes of task-evoked activity flow at a finer-grained scale, using voxels instead of regions. We also performed this analysis to gain a more general assessment of the accuracy of the activity-flow mapping approach (for example, without assuming a set of a priori defined brain regions). Note that we excluded FC with all voxels within the same region as (and voxels within 9 mm of) the to-be-predicted voxel to reduce the chance of spatial autocorrelations[35] contributing to prediction accuracies (Online Methods).

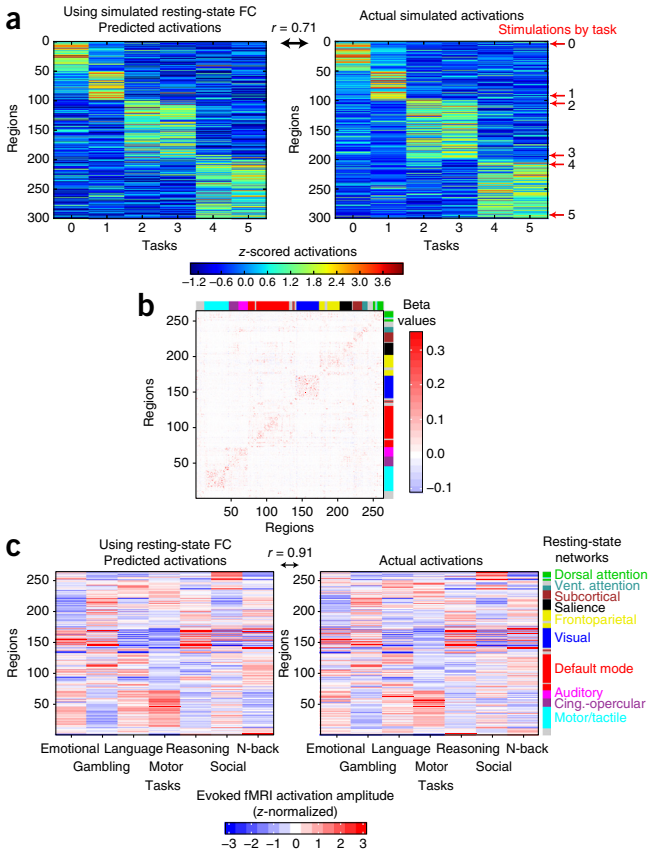

Figure 3 Using multiple regression to estimate resting-state FC increases prediction accuracy. (**a**) We applied standard multiple linear regression to estimate each region's FC in the same simulated data shown in **Figure 1**. This increased prediction accuracy from $r = 0.56$ (with Pearson correlation FC) to $r = 0.71$ in this example. (**b**) Multiple-regression FC matrix from the real resting-state fMRI data. The cross-subject average regression coefficient matrix is shown. Some community structure was apparent, despite the increased sparseness relative to when Pearson correlation was used (**Fig. 2b**). Colors indicate resting-state networks as given in **c**. (**c**) Prediction accuracy was also increased with real fMRI data: from an average of $r = 0.66$ (using Pearson correlation FC) to an average of $r = 0.91$ (with multiple-regression FC); a 39% increase in linear variance explained.

We found that whole-brain voxelwise activation patterns were predicted well above chance: cross-task average $r = 0.63$, $t_{99} = 40.68$, $P < 0.00001$. The *r*-values were similar for each of the seven tasks individually: $r = 0.54$ (emotional), $r = 0.65$ (gambling), $r = 0.57$ (language), $r = 0.67$ (motor), $r = 0.66$ (reasoning), $r = 0.72$ (social) and $r = 0.58$ (N-back). The average correlation was higher when comparisons were computed after averaging the predicted and actual activation patterns across subjects (cross-task average $r = 0.92$; 85% of variance; **Fig. 4** and **Supplementary Fig. 4**). This was likely due to an improved signal-to-noise ratio from averaging more data. Results were similar for each of the seven tasks individually: $r = 0.84$ (emotional), $r = 0.93$ (gambling), $r = 0.90$ (language), $r = 0.94$ (motor), $r = 0.95$ (reasoning), $r = 0.94$ (social) and $r = 0.86$ (N-back). While these predictions are highly accurate, even with 90% of the variance ($r = 0.95$; $r^2 = 0.90$) being explained for the reasoning task (**Fig. 4b**), there appear to be meaningful differences in the remaining 10% of variance, such as a lack of primary motor cortex activation for the reasoning task. This suggests that, despite the strong performance of resting-state FC-based activity-flow mapping, there may be important roles for task-evoked FC and/or local within-voxel processing that does not flow to other brain regions.

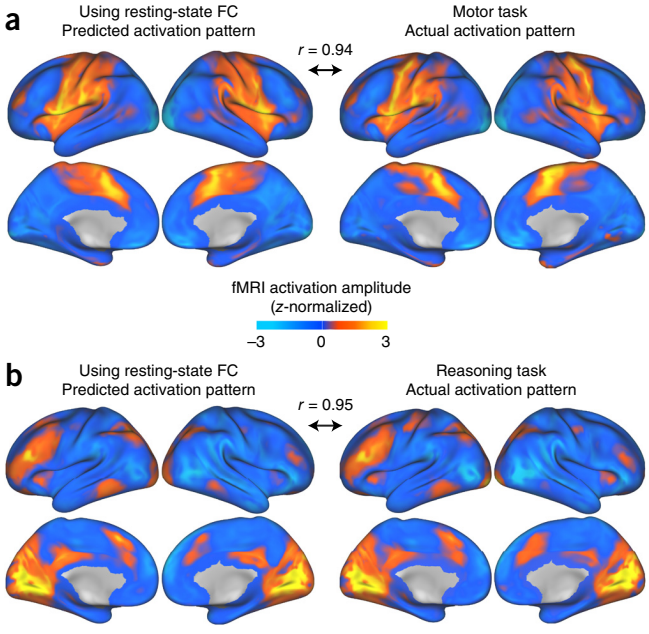

Figure 4 Predicting voxelwise activation patterns. We used multiple-regression-based resting-state FC with a voxelwise activity flow mapping approach, again finding above-chance prediction accuracy across the 7 tasks. Activity flow mapping with the (**a**) motor and reasoning (**b**) relational tasks are illustrated. Voxels within the same region and within 9 mm of that region were excluded from prediction calculations to reduce the influence of spatial autocorrelations (Online Methods). Also, due to fewer time points than predictors, only a subset of the data could be used to compute voxelwise multiple-regression FC (Online Methods). See **Supplementary Figure 4** for predicted and actual activation maps for all 7 tasks.

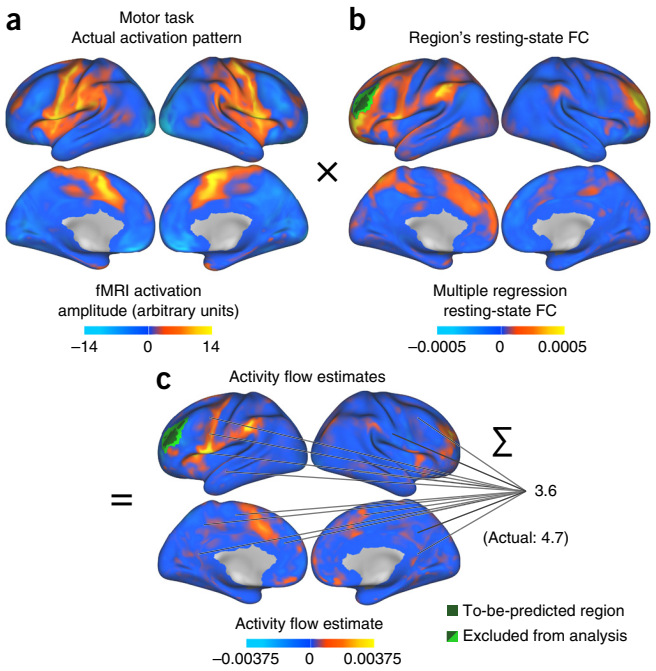

Figure 5 Illustration of activity flow mapping of single region. Voxelwise prediction of a single region is illustrated. (**a**) The region of interest was defined based on positive activation during the motor task. The same procedure was used for prediction of voxels (**Fig. 4**). (**b**) Multiple-regression-based resting-state FC using the region of interest's average time series as a seed. The average across all 100 subjects is shown. (**c**) Activity flow estimates to and from the region of interest during the motor task. Small values indicate that the prediction is based on a highly distributed activation pattern, with each activity flow estimate contributing only a small amount (full range: –0.0095 to 0.015). The (infrequent) occurrence of negative FC multiplied by negative activations leading to positive activity flow estimates may reflect disinhibition or inhibitory activity flow from (rather than to) the to-be-predicted region. While not exact, the prediction is in the same range as the region of interest's actual activation amplitude during the motor task. The whole-brain motor task activation amplitude range was –23.6 to 16.3.

Note that, unlike in the region-wise analysis above, it was statistically impossible to include all predictors (voxels) for all to-be-predicted voxels. This is because multiple regression requires more data points than predictors. We used principal components regression to get around this limitation (Online Methods). However, because not all resting-state fMRI time series variance was included in the predictions, these may be underestimates of the voxelwise predictions possible with more data. The voxelwise activity flow mapping procedure is illustrated in **Figure 5**, as applied to predicting activation in a single region during the motor task. An anterior prefrontal cortex region in the cingulo-opercular network (**Fig. 2a**) was chosen for illustration since neither anterior prefrontal cortex nor the cingulo-opercular network are typically considered in the literature in the context of such simple motor tasks[36,37]. The FC and activity flow with primary motor cortex in **Figure 5** provide an example of the additional insight that can be gained using the activity-flow approach. Overall, these results further demonstrate the plausibility of the activity-flow network mechanism in shaping cognitive task activations.

Quantification of prediction accuracies by network

We found that overall prediction accuracy was high across all tasks, but we wanted to also quantify prediction accuracy for each network separately. We used the same approach as our main results (task-specific activations predicted using multiple-regression FC, calculated after averaging across subjects), but predicted-to-actual correlations were calculated for each network separately. We found that correlations for all networks were high ($r > 0.8$ on average), though there was some variability. The cross-task average predicted-to-actual correlations for each network were: $r = 0.89$ (motor–tactile (hand)), $r = 0.91$ (motor–tactile (mouth)), $r = 0.95$ (cingulo-opercular), $r = 0.93$ (auditory), $r = 0.92$ (default mode), $r = 0.98$ (memory retrieval), $r = 0.93$ (visual), $r = 0.93$ (frontoparietal), $r = 0.93$ (salience), $r = 0.67$ (subcortical), $r = 0.94$ (ventral attention), $r = 0.96$ (dorsal attention) and $r = 0.84$ (cerebellum). These correlations were also calculated for each task separately (**Supplementary Fig. 3a**). Note that networks showing poor accuracy for individual tasks were accurate overall when predictions across all tasks were considered (**Supplementary Fig. 3b**).

These results illustrate that effects were similar across cortical networks but tended to be lower for subcortical regions. It may be illuminating for future studies to investigate this as a possible difference between cortical and subcortical activity flow mechanisms. Note, however, that the 32-channel MRI head coil and multiband fMRI sequence used here are thought to reduce signal-to-noise for subcortical regions relative to cortical regions[29], possibly leading to the observed effect.

Predicting individual differences in cognitive task activations

We next tested whether activity flow mapping can be used to predict individual differences in cognitive task activations based on individual differences in resting-state FC. A recent study was able to do this using an abstract statistical model trained to directly associate

(within small patches of cortex) resting-state FC values with cognitive task activations[9]. We postulated that if activity flow is a large-scale mechanism linking resting-state FC to cognitive task activations, then activity flow mapping would also produce above-chance prediction of held-out individual subjects. Notably, unlike in the previous study, activity flow mapping does not involve training an abstract statistical model associating resting-state FC with task activations, thus potentially demonstrating a more direct relationship between resting-state FC and cognitive task activations.

In addition to holding out each region one at a time, for this analysis we also held out activations from each subject one at a time. This allowed us to use the held-out individual's resting-state FC, in combination with other subjects' mean task activations, to predict the held-out individual's cognitive task activations (Online Methods). The predicted task-specific activation patterns were again above chance on average: $r = 0.45$, $t_{99} = 25.15$, $P < 0.00001$. The r-values were similar for each of the seven tasks individually: $r = 0.23$ (emotional), $r = 0.48$ (gambling), $r = 0.37$ (language), $r = 0.54$ (motor), $r = 0.52$ (reasoning), $r = 0.58$ (social) and $r = 0.40$ (N-back). These results demonstrate that resting-state FC describes individualized routes of activity flow, which shape individual differences in cognitive task activations.

It is possible that activation predictions in the held-out individuals were above chance due to the general similarity of activations across subjects, rather than due to prediction of individual differences. Consistent with this, cross-subject cognitive task activation pattern similarity was $r = 0.20$ on average. We therefore used regression to isolate individual differences in resting-state FC and actual activation patterns (Online Methods), reducing cross-subject cognitive task activation pattern similarity to $r = -0.01$ on average. This revealed that individual differences in cognitive task activations could be predicted based on individual differences in resting-state FC: cross-task average $r = 0.59$, $t_{99} = 25.15$, $P < 0.00001$. The r-values were similar for each of the seven tasks individually: $r = 0.66$ (emotional), $r = 0.56$ (gambling), $r = 0.59$ (language), $r = 0.54$ (motor), $r = 0.65$ (reasoning), $r = 0.58$ (social) and $r = 0.51$ (N-back). These results suggest that resting-state FC describes individualized routes of activity flow that shape individual differences in cognitive task activations.

DISCUSSION

Recent studies have shown a strong statistical relationship between resting-state FC and cognitive task activations. This was shown using meta-analytic data from thousands of fMRI experiments[8] and in individual subjects performing specific tasks[9]. However, it has remained unclear how or why this relationship exists. Understanding this relationship in a more mechanistic manner would provide critical insight into the relevance of resting-state FC for cognitive task activations. This would also provide insight into the factors that shape cognitive task activations, a central goal of cognitive neuroscience. Based on our recent work showing that resting-state FC patterns are present during task performance[11], we expected that resting-state FC may describe the routes of activity flow even during task performance. Building on this, we tested the possibility that activity flow is a linking (large-scale) mechanism between resting-state FC and cognitive task activations, potentially explaining the statistical relationship previously observed between these two constructs.

We quantified activity flow as the FC-weighted sum of activations in other brain regions (**Fig. 1a**). Using empirical fMRI data, we found that estimating activity flow across resting-state FC networks allowed prediction of cognitive task activations (**Fig. 2**). This was true when holding out each brain region (or voxel) but also when holding out each individual subject. This demonstrated that individual differences in intrinsic network activity flow can help explain individual differences in cognitive task activations. This may have application in the future for predicting and understanding cognitive task activations in patients who cannot perform a given task (for example, due to lack of consciousness or cognitive disability) or who perform the task poorly. This addresses a key issue in the study of cognitive disability: We wish to investigate patients with cognitive disabilities using the tasks they have difficulty with, but by definition they will be performing those tasks differently than healthy control subjects. This leads to 'performance confounds', in which any observed change in cognitive task activations could be either a cause or a consequence of the disrupted cognitive task performance. Use of activity flow mapping (and related approaches[9,15,16,38–40]) may allow us to get around such confounds, since we can now understand individual differences in cognitive task activations in terms of connectivity variables estimated independently of task performance.

Several recent studies also sought to identify the relationship between individual subject connectivity and cognitive task (as well as brain-stimulation-based[41]) activations[9,15,16]. These studies found that functional and structural connectivity patterns, when combined with a statistical model fit to separate data, could be used to predict individual differences in cognitive task activations. These studies provided further evidence of a relationship between large-scale connectivity and cognitive activations. Unlike these studies, we used a large-scale mechanistic construct—activity flow—to link connectivity and cognitive task activity without use of a statistical model trained to relate connectivity to activations. This allowed us to infer a more direct relationship between connectivity and task activations and to link this relationship to a potential underlying large-scale mechanism (activity flow, likely shaped in part by aggregate synaptic connectivity strengths; **Fig. 1d**). Linking to mechanistic constructs (large-scale or otherwise) is important for theoretical advances in neuroscience. In this case, linking to activity flow supports an explanation for the statistical relationships observed in these previous studies. Future studies may build on these findings with manipulations of FC and activity flow to make more causal inferences about these constructs. Further, it may be useful to investigate the relationship between these constructs using more direct measures of neural activity such as multiunit recording or magnetoencephalography, given that (while strong[42,43]) the link between neural activity and BOLD fMRI is indirect.

We began by using a simple FC measure (Pearson correlation) to model activity flow. We did this primarily to make minimal assumptions regarding the true nature of brain interactions and because Pearson correlations are widely used for FC estimation in the literature. It is noteworthy that we observed such high accuracy in our predictions (over 40% of variance explained; **Fig. 2**) despite using FC estimates that lacked information about both the direction of influence and whether an influence between nodes was indirect (i.e., effective connectivity)[44]. We found that when we used multiple linear regression as an FC measure, activity flow mapping accuracies increased (**Fig. 3**). Unlike Pearson correlation, this measure isolates unique influences between regions. Just as we found that using multiple-regression FC increased the activation-pattern prediction accuracy, so we expect that adding additional (or more accurate) information—perhaps using more sophisticated effective connectivity methods[45–47]—will improve prediction accuracy further. This would provide evidence for the importance of these factors in shaping cognitive task activations. More generally, this illustrates a benefit of the activity flow framework: the accuracy of predicted activation patterns can be evidence for the veracity of any connectivity properties of interest.

The large-scale aggregate activity-flow construct tested here was derived from well-known mechanisms for activity flow at the local circuit level. The key local mechanism allowing for activity flow among neurons is the propagation of action potentials along axons. The demonstration of large-scale activity flow here is nontrivial given that action potentials occur at the level of individual neurons and on the order of tens of milliseconds, while BOLD fMRI signals involve hundreds of thousands of neurons over seconds. This suggests that aggregating many instances of local activity flow via action potentials results in self-similar[48] large-scale properties. This is consistent with the computational model results (**Fig. 1**), which demonstrated that activity flow occurring on the order of milliseconds can nonetheless be estimated accurately using simulated fMRI signals. However, it will be important to investigate activity flow at other spatiotemporal scales. Most crucially, it will be important to increase temporal resolution to observe the time-lagged propagation of signals between brain regions, allowing for clear directional activity-flow estimates. Note that the primary reason we did not take temporal lags into account is the uncertainty of precise BOLD response timings relative to underlying neural response timings across brain regions and voxels[47,49]. Further progress characterizing BOLD signal timing and advances in spatial coverage and localization for high-temporal-resolution methods will likely lead to improved estimates of activity flow.

It is important to consider our approach in the context of other modeling frameworks. The activity flow approach is analogous to a model of task activation at time $t + 1$ based on the product of the activation at time t and the connectivity matrix. As such, this is a single-time-step prediction based on the simple linear model used in previous studies[50], with one important change: the connectivity matrix we used is based on FC rather than structural connectivity. It will be interesting in future to extend our model to include considerations of nonlinear dynamics, such as those implemented in the Virtual Brain project[28].

To conclude, it is well established that there is a strong statistical relationship between resting-state FC and cognitive task activations[8], yet the reason for this relationship has remained unclear. We provided evidence for a large-scale mechanism involving activity flow over intrinsic networks (described by resting-state FC) shaping cognitive task activations. This suggests that observed cognitive task activations should not be interpreted simply in terms of localized processing but should also consider distributed processing in the form of activity flow across intrinsic networks. Further, these results suggest strong relevance of resting-state FC for the task activations that produce cognition. We expect that these insights and the activity flow mapping procedure introduced here will facilitate future investigation into the functional relevance of resting-state FC, the factors that influence cognitive task activations and the balance of large-scale distributed versus localized processing in the human brain.

Data availability. The MRI data set analyzed as part of the current study is available in the Human Connectome Project's ConnectomeDB repository (https://db.humanconnectome.org) under the identifiers 'WU-Minn HCP Data' and '100 Unrelated Subjects'.

METHODS

Methods, including statements of data availability and any associated accession codes and references, are available in the online version of the paper.

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

We thank B. Biswal, M. Dixon, T. Braver, S. Petersen and J. Power for helpful conversations during preparation of this manuscript. Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: D. Van Essen and K. 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. M.W.C. was supported by the US National Institutes of Health under award K99-R00 MH096801. D.S.B. acknowledges support from the John D. and Catherine T. MacArthur Foundation, the Army Research Laboratory and the Army Research Office through contract numbers W911NF-10-2-0022 and W911NF-14-1-0679, the National Institute of Mental Health (2-R01-DC-009209-11), the National Institute of Child Health and Human Development (1R01HD086888-01), the Office of Naval Research and the National Science Foundation (#BCS-1441502, #BCS-1430087, and #PHY-1554488). The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

AUTHOR CONTRIBUTIONS

M.W.C. conceived of the study, developed the activity flow mapping algorithm, developed the computational model, developed the multiple-regression functional connectivity approach, performed the analyses and wrote the manuscript. T.I. developed the computational model and assisted with writing the manuscript. D.S.B. provided feedback on the activity flow mapping algorithm and assisted with writing the manuscript. D.H.S. assisted with writing the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

1. Fox, M.D. & Raichle, M.E. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* **8**, 700–711 (2007).
2. Power, J.D., Schlaggar, B.L. & Petersen, S.E. Studying brain organization via spontaneous fMRI signal. *Neuron* **84**, 681–696 (2014).
3. Biswal, B.B. *et al.* Toward discovery science of human brain function. *Proc. Natl. Acad. Sci.* **107**, 4734–4739 (2010).
4. Saxe, R., Brett, M. & Kanwisher, N. Divide and conquer: a defense of functional localizers. *Neuroimage* **30**, 1088–1096, discussion 1097–1099 (2006).
5. Haxby, J.V. *et al.* Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* **293**, 2425–2430 (2001).
6. Duncan, J. The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends Cogn. Sci.* **14**, 172–179 (2010).
7. Kanwisher, N. Functional specificity in the human brain: a window into the functional architecture of the mind. *Proc. Natl. Acad. Sci. USA* **107**, 11163–11170 (2010).
8. Smith, S.M. *et al.* Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. USA* **106**, 13040–13045 (2009).
9. Tavor, I. *et al.* Task-free MRI predicts individual differences in brain activity during task performance. *Science* **352**, 216–220 (2016).
10. Biswal, B., Yetkin, F.Z., Haughton, V.M. & Hyde, J.S. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* **34**, 537–541 (1995).
11. Cole, M.W., Bassett, D.S., Power, J.D., Braver, T.S. & Petersen, S.E. Intrinsic and task-evoked network architectures of the human brain. *Neuron* **83**, 238–251 (2014).
12. Jessell, T.M. & Kandel, E.R. Synaptic transmission: a bidirectional and self-modifiable form of cell-cell communication. *Cell* **72** (Suppl.), 1–30 (1993).
13. Laughlin, S.B. & Sejnowski, T.J. Communication in neuronal networks. *Science* **301**, 1870–1874 (2003).
14. Hodgkin, A.L. & Huxley, A.F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol. (Lond.)* **117**, 500–544 (1952).
15. Saygin, Z.M. *et al.* Anatomical connectivity patterns predict face selectivity in the fusiform gyrus. *Nat. Neurosci.* **15**, 321–327 (2011).
16. Osher, D.E. *et al.* Structural connectivity fingerprints predict cortical selectivity for multiple visual categories across cortex. *Cereb. Cortex* **26**, 1668–1683 (2016).
17. Smith, V.A., Yu, J., Smulders, T.V., Hartemink, A.J. & Jarvis, E.D. Computational inference of neural information flow networks. *PLoS Comput. Biol.* **2**, e161 (2006).
18. Fries, P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn. Sci.* **9**, 474–480 (2005).
19. Norman, K.A., Polyn, S.M., Detre, G.J. & Haxby, J.V. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* **10**, 424–430 (2006).
20. Cole, M.W., Etzel, J.A., Zacks, J.M., Schneider, W. & Braver, T.S. Rapid transfer of abstract rules to novel contexts in human lateral prefrontal cortex. *Front. Hum. Neurosci.* **5**, 142 (2011).

21. Haynes, J.-D. A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives. *Neuron* **87**, 257–270 (2015).
22. Cole, M.W. *et al.* Multi-task connectivity reveals flexible hubs for adaptive task control. *Nat. Neurosci.* **16**, 1348–1355 (2013).
23. Medaglia, J.D., Lynall, M.-E. & Bassett, D.S. Cognitive network neuroscience. *J. Cogn. Neurosci.* **27**, 1471–1491 (2015).
24. Curtis, C.E. & D'Esposito, M. Persistent activity in the prefrontal cortex during working memory. *Trends Cogn. Sci.* **7**, 415–423 (2003).
25. Mišić, B. *et al.* Cooperative and competitive spreading dynamics on the human connectome. *Neuron* **86**, 1518–1529 (2015).
26. Gu, S. *et al.* Controllability of structural brain networks. *Nat. Commun.* **6**, 8414 (2015).
27. Deco, G. *et al.* Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations. *J. Neurosci.* **33**, 11239–11252 (2013).
28. Ritter, P., Schirner, M., McIntosh, A.R. & Jirsa, V.K. The virtual brain integrates computational modeling and multimodal neuroimaging. *Brain Connect.* **3**, 121–145 (2013).
29. Barch, D.M. *et al.* Function in the human connectome: task-fMRI and individual differences in behavior. *Neuroimage* **80**, 169–189 (2013).
30. Power, J.D. *et al.* Functional network organization of the human brain. *Neuron* **72**, 665–678 (2011).
31. Duncan, J. & Owen, A.M. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci.* **23**, 475–483 (2000).
32. Chein, J.M. & Schneider, W. Neuroimaging studies of practice-related change: fMRI and meta-analytic evidence of a domain-general control network for learning. *Brain Res. Cogn. Brain Res.* **25**, 607–623 (2005).
33. Cabeza, R. & Nyberg, L. Imaging cognition II: an empirical review of 275 PET and fMRI studies. *J. Cogn. Neurosci.* **12**, 1–47 (2000).
34. Shulman, G.L. *et al.* Common blood flow changes across visual tasks: I. increases in subcortical structures and cerebellum but not in nonvisual cortex. *J. Cogn. Neurosci.* **9**, 624–647 (1997).
35. Kiebel, S.J., Poline, J.B., Friston, K.J., Holmes, A.P. & Worsley, K.J. Robust smoothness estimation in statistical parametric maps using standardized residuals from the general linear model. *Neuroimage* **10**, 756–766 (1999).
36. Power, J.D. & Petersen, S.E. Control-related systems in the human brain. *Curr. Opin. Neurobiol.* **23**, 223–228 (2013).
37. Badre, D. Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends Cogn. Sci.* **12**, 193–200 (2008).
38. Kannurpatti, S.S., Rypma, B. & Biswal, B.B. Prediction of task-related BOLD fMRI with amplitude signatures of resting-state fMRI. *Front. Syst. Neurosci.* **6**, 7 (2012).
39. Mennes, M. *et al.* Inter-individual differences in resting-state functional connectivity predict task-induced BOLD activity. *Neuroimage* **50**, 1690–1701 (2010).
40. Heinzle, J., Kahnt, T. & Haynes, J.-D. Topographically specific functional connectivity between visual field maps in the human brain. *Neuroimage* **56**, 1426–1436 (2011).
41. Fox, M.D. *et al.* Resting-state networks link invasive and noninvasive brain stimulation across diverse psychiatric and neurological diseases. *Proc. Natl. Acad. Sci. USA* **111**, E4367–E4375 (2014).
42. Lee, J.H. *et al.* Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. *Nature* **465**, 788–792 (2010).
43. Siero, J.C.W. *et al.* BOLD matches neuronal activity at the mm scale: a combined 7T fMRI and ECoG study in human sensorimotor cortex. *Neuroimage* **101**, 177–184 (2014).
44. Friston, K.J. Functional and effective connectivity: a review. *Brain Connect.* **1**, 13–36 (2011).
45. Ramsey, J.D. *et al.* Six problems for causal inference from fMRI. *Neuroimage* **49**, 1545–1558 (2010).
46. Friston, K.J., Harrison, L. & Penny, W. Dynamic causal modelling. *Neuroimage* **19**, 1273–1302 (2003).
47. Smith, S.M. *et al.* Network modelling methods for FMRI. *Neuroimage* **54**, 875–891 (2011).
48. Mandelbrot, B. How long is the coast of Britain? Statistical self-similarity and fractional dimension. *Science* **156**, 636–638 (1967).
49. Handwerker, D.A., Ollinger, J.M. & D'Esposito, M. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage* **21**, 1639–1651 (2004).
50. Galán, R.F., Ermentrout, G.B. & Urban, N.N. Optimal time scale for spike-time reliability: theory, simulations, and experiments. *J. Neurophysiol.* **99**, 277–283 (2008).

ONLINE METHODS

Activity flow mapping. We developed a method to quantify the relationship between FC and task-activation patterns (**Fig. 1a**). This involved estimating net input to each target region by multiplying each other brain region's task-related activation amplitude (analogous to the amount of neural activity) by its FC with the target region (analogous to aggregate synaptic strength):

$$P_j = \sum_{i \neq j \in V} A_i F_{ij}$$

where P_j is the predicted mean activation for region j in a given task, A_i is the actual mean activation for region i in a given task (a beta value estimated using a general linear model), i indexes all brain regions (vector V) with the exception of region j, and F_{ij} is the FC estimate between region i and region j (the Fisher z-transformed Pearson correlation or multiple regression estimate of the regions' time-series). This algorithm results in a vector predicting the pattern of mean activations across regions for a given task. Note that when FC is used (rather than directed/effective connectivity) this approach estimates total bidirectional (and/or indirect) activity flow.

Activation amplitudes were z-normalized for each task separately via subtracting each activation amplitude from the cross-region mean and dividing by the cross-region s.d. This facilitated a focus on the activation patterns (rather than absolute activation levels) across tasks. Prediction accuracy was assessed using Pearson correlation between the predicted activation values and the actual activation values (i.e., the actual activations for each region and task). This was done for each task separately and (unless noted otherwise) each subject separately. Each correlation value was Fisher's z-transformed before averaging, then converted back to a Pearson correlation for reporting purposes. Statistical significance tests were conducted using t-tests (two-sided, one-sample) of Fisher's z-transformed Pearson correlations, facilitating the ability to infer generalization of results across subjects (rather than just on cross-subject mean patterns). The group distributions of these Fisher's z-transformed Pearson correlations were approximately normally distributed. When P-values were computed based on non-normally distributed data we also reported a P-value based on the Spearman's rank correlation. Note that predicted activation patterns and actual activation patterns were averaged across subjects before comparison for a subset of analyses. This was done primarily to increase the signal-to-noise ratio via averaging of more data, likely providing more accurate effect size estimates (i.e., percent variance explained).

Computational modeling. We used a simple computational model of large-scale neural interaction to help validate key aspects of activity flow mapping. We sought the simplest computational model possible to reduce the number of biophysical assumptions and improve the likely generality of our results.

The model consists of 300 abstract units, each representing a brain region. The units interact via a standard spiking rate code passed via predefined structural (and synaptic) connectivity[51]. Activity at a given node is determined using a standard sigmoid function on the mean of the input activities. Note that the sigmoid function introduces a nonlinearity to the interactions among units that is similar to aggregate nonlinearity from neuronal action potentials[52]. Specifically, the model used the following equation to determine activity in a given unit at a given time step:

$$\tau_i \frac{dx_i}{dt} = -x_i + f_i\left(\sum_{j=1}^{n} w_{ji} x_j + \text{bias}_i\right) \; i = [1..n]$$

where w_{ji} refers to the synaptic weight from region i to j and x_j refers to the activity level at region j. Bias_i is the bias of region i, but for this model this is set to 0. τ_i is the time constant for region i, and is set to 1 time step for all regions. f_i is a standard sigmoid function.

The model's network connectivity was constructed by first defining a random set of structural connections (15% density), then creating three graph communities/subnetworks by randomly connecting each node to ten other nodes within the same community. Structural connections were defined as nonzero connection weights (all set to the same value of 1.0), while synaptic connections were modifications on the initial connection weight. Normally distributed random synaptic weights were added to all structural connections, scaled to be quite small (mean of 0 and s.d. of 0.001). Finally, synaptic weights were used to split the first structural connectivity community into two 'functional' communities. Specifically, the synaptic weights were increased (multiplied by 1.5) within the first half of the first graph community, while synaptic weights for the second half of that community were also increased (multiplied by 1.5). Also, synaptic weights between these communities were reduced (multiplied by 0.5). These modifications were designed to test the impact of synaptic weights on simulated activity flow.

Spontaneous activity for each node was added as normally distributed random values (mean = 0, s.d. = 1) every time step (100 ms). An autocorrelation factor of 0.10 was used to maintain some activity across multiple time steps. We simulated 20,000 time steps using purely spontaneous activity (resting state data) and, separately, using spontaneous with task-evoked activity (task data). Task-evoked activity was implemented as increased activity (normally distributed random values centered at 1 with an s.d. of 0.5) added linearly to ongoing spontaneous activity. Activity consisted of three blocks of 2,000 timepoints each, with each block separated by 3,000 time points. Each task was simulated by adding task-evoked activity to six separate groups of five regions simultaneously (two per structural graph community). fMRI data collection was simulated by convolving the simulated time series with the SPM canonical hemodynamic response function, then downsampling to a standard TR of 2 s. All analyses of the simulated fMRI data were identical to the analyses conducted on the empirical fMRI data.

We defined a global coupling parameter as a scalar multiplier on all synaptic strengths, and a local processing parameter as a scalar multiplier of all self-connection strengths. Self-connections increase the influence of a region's activity on itself in the next time-step, separating variance in its activity from the activity of other regions. For the parameter sweep (**Fig. 1b**) we used 20 global coupling parameters (from 0 to 5, using 0.25 increments) and 20 local processing parameters (from 0 to 100, using 5.0 increments), each averaged across 10 'subjects' (separate iterations with random initial structural/synaptic connectivity matrices). This totaled 4,000 simulations. Modeling was carried out using Python (version 2.7).

Data collection. Data were collected as part of the Washington University–Minnesota Consortium Human Connectome Project (HCP)[53]. Human participants were recruited from Washington University (St. Louis, MO) and the surrounding area. All participants gave informed consent consistent with policies approved by the Washington University Institutional Review Board. The data used were from the "500 Subjects" HCP release. The "100 Unrelated Subjects" ($n = 100$) subset of this data set was used, given that a subset of unrelated individuals is more appropriate for statistical analyses intended to represent the general population. Details regarding randomization can be found in the relevant HCP paper[53]. Based on our primary statistical tests (one-sample t-tests, $\alpha = 0.05$) and assuming a moderate Cohen's d effect size of 0.5, $n = 100$ provides 99.86% power[54] (higher than the standard criterion of 80%). The average age of the participants was 29 years (ranging from 22 to 36), and 54% were female. Whole-brain echo-planar imaging acquisitions 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 = 2,290 Hz/Px, in-plane FOV = 208 × 180 mm, 72 slices, 2.0 mm isotropic voxels, with a multiband acceleration factor of 8 (ref. 55). Data were collected over 2 d. On each day 28 min of rest (eyes open with fixation) fMRI data across two runs were collected (56 min total), followed by 30 min of task fMRI data collection (60 min total). Each of the seven tasks was completed over two consecutive fMRI runs. Resting-state data collection details for this data set can be found elsewhere[56], as can task data details[29].

Task paradigms. The data set was collected as part of the Human Connectome Project, and included rest and a set of seven tasks[29]. These tasks included seven distinct domains: emotion, reward learning, language, motor, relational reasoning, social cognition and working memory. Briefly, the emotion task involved matching fearful or angry faces to a target face. The reward learning task involved a gambling task with monetary rewards and losses. The language task involved auditory stimuli consisting of narrative stories and math problems, along with questions to be answered regarding the prior auditory stimuli. The motor task involved movement of the hands, tongue and feet. The relational reasoning task involved higher-order cognitive reasoning regarding relations among features of presented shape stimuli. The social cognition (theory of mind) task used short

doi:10.1038/nn.4406

video clips of moving shapes that interacted in some way or moved randomly, with subjects making decisions about whether the shapes had social interactions. The working memory task involved a visual N-back task, in which subjects indicate a match of the current image to either a constant target image or two images previous.

Data preprocessing. Preprocessing consisted of standard resting-state functional connectivity preprocessing (typically performed with resting state data), with several modifications given that analyses were also performed on task data. Resting-state and task data were preprocessed identically in order to facilitate comparisons between them.

Spatial normalization to a standard template, motion correction, and intensity normalization were already implemented as part of the Human Connectome Project in a minimally processed version of the data set described elsewhere[57]. With the volume (rather than the surface) version of the minimally preprocessed data, we used AFNI[58] to additionally remove nuisance time series (motion, ventricle and white matter signals, along with their derivatives) using linear regression, remove the linear trend for each run and spatially smooth the data. The data were smoothed using a non-Gaussian filter (nearest neighbor averaging) at 4 mm to reduce the chance of introducing circularity in the activity flow mapping procedure (see below). Unlike some standard resting-state FC preprocessing pipelines, global signal was not included as a nuisance covariate (given current controversy over this procedure[59]). Note that activity flow mapping results were similar after global signal regression (GSR; see the section "Activity flow mapping with global signal regression" below). We did not apply a low-pass temporal filter, given the likely presence of task signals at higher frequencies than the relatively slow resting-state fluctuations and our desire to preprocess the rest and task data similarly. Freesurfer[60] was used to identify ventricle, white matter and gray matter anatomical structures for each participant.

For the main analyses, data were sampled from a set of 264 brain regions (rather than individual voxels) in order to make inferences at the region and systems level (**Fig. 2a**). We used an independently identified set of putative functional brain regions[30] rather than anatomically defined sets of regions in order to reduce the chance of combining signal from multiple functional areas[61]. These brain regions were identified using a combination of resting-state FC parcellation[62] and task-based neuroimaging meta-analysis[30]. Data were summarized for each region by averaging signal in all voxels falling inside each region. Analyses were carried out with Matlab 2014b (Mathworks) and R 3.1.2 (The R Foundation for Statistical Computing).

FC estimation. The initial analyses estimated FC using Pearson correlations between time series (averaging across voxels within each region) from all pairs of brain regions. All computations used Fisher's *z*-transformed values, which were reconverted to *r*-values for reporting purposes.

We used standard multiple linear regression (the regstats function in Matlab) as an alternative to Pearson correlation. This involved computing a linear model for each to-be-predicted region separately. Resting-state fMRI time series from all other regions were used as predictors of the to-be-predicted region's resting-state fMRI time series. The resulting betas, which were directional from the predictor regions to the predicted region, were then used as FC estimates in the activity flow mapping algorithm. Note that beta estimate directionality reflects optimal linear scaling of the source time series to best match the target time series (based on resting-state fMRI data), not necessarily the direction of activity flow.

Task-activation level estimation. The activation amplitudes were estimated using a standard general linear model. The SPM canonical hemodynamic response function was used for general linear model estimation, given that all tasks involved block designs.

Activity flow mapping permutation testing. We used permutation testing to help validate the activity flow mapping approach and provide an additional means of inferring statistical significance. The permutation test was constructed so as to facilitate a conservative statistical inference, wherein only the hypothesized essential aspect of the analysis was permuted. This involved keeping all aspects of the analysis the same except for random permuting (without replacement) which region's FC was used on each iteration. In other words, the entire set of FC strengths for the to-be-predicted region was swapped with the entire set of FC strengths for another region chosen uniformly at random from the set of all regions. This permutation process was run 10,000 times (with resting-state FC), resulting in a null distribution of *r*-values.

Voxelwise activity flow mapping. We made relatively minimal changes to the regional activity flow modeling procedure when applying it in a voxelwise manner. First, we excluded all voxels within the same functional region (defined as local voxels with similar resting-state FC patterns) as the to-be-predicted voxel in order to reduce the influence of potentially trivial within-region activity flow upon prediction accuracies. Second, we excluded all voxels within 9 mm of the same functional region as the to-be-predicted voxel to reduce the chance of spatial autocorrelations contributing to prediction accuracies[35]. Non-Gaussian smoothing was also used (averaging neighboring voxels) to further reduce spatial autocorrelation. The recently developed Gordon cortical area parcellation[63] was used because (unlike the Power brain area parcellation used for the other analyses) it includes a voxelwise version amenable to our processing pipeline and because its development involved similar principles as the Power brain area parcellation. Note that we used the Power brain area parcellation for the region-wise analyses because it is better established and may have more accurate network assignments than the Gordon parcellation (an issue not relevant to the voxelwise analyses). This conclusion is based on greater similarity of network assignments to independently derived network assignments by Yeo *et al.*[64]. Unlike the region-wise analyses, the voxelwise analyses were restricted to the voxels included in the Gordon parcellation (i.e., cortex). The 2-mm cubic voxels were downsampled to 3-mm cubic voxels (using linear interpolation) to increase computational tractability. Finally, the voxelwise activity-flow predictions were calculated for each subject independently, and the resulting prediction (and actual) maps were subsequently averaged across subjects before actual-to-predicted comparison. Results are also reported with predictions compared to actual activation patterns for each subject separately. We used Connectome Workbench software (v1.0) for visualization. Statistical maps were smoothed on the surface with two standard (in Connectome Workbench) iterations before visualization.

For the multiple-regression-based voxelwise activity-flow approach, there were many more predictors (voxels) than time points. Thus, unlike the region-wise analyses, this made it impossible to compute FC estimates using all available predictors. Instead we used a standard statistical approach for performing multiple regression with many more predictors than data points: principal components regression[65]. Briefly, this involved extracting the time series for the first 1,200 principal components, performing the regression on each voxel using those components as predictors, then projecting the resulting beta values back into the original voxel space (from the principal component space). The principal components were calculated independently for each to-be-predicted region, with that region's voxels and voxels within 9 mm excluded to avoid circularity. We used the first 1,200 components (out of 4,800 resting-state fMRI time points) for computational tractability. Note that the same procedures were used for the **Figure 5** illustrative analysis, except that a single region of interest's activation level was predicted rather than a single voxel's activation level.

Task-specific activation patterns. Task-general activation patterns were defined as the first principal component across task activation patterns. Principal component analysis was used rather than averaging to reduce the chance that any individual task's activation pattern dominated the task-general pattern. This was computed separately for each subject and also for each task; the to-be-predicted (or compared) task's activations were excluded to remove circularity from the calculation. Results were virtually identical if all seven tasks were included in the task-general activation calculation. Task-specific activation patterns were defined as a given task's activation pattern after regressing out task-general activations (the first principal component across the other six tasks' activation patterns). The average pairwise similarity among task-specific activation patterns (i.e., after regressing out task-general activations) was $r = -0.1$.

Prediction of individualized task activations. Each subject's cognitive task activations were held out in a leave-one-subject-out approach. The held-out individual's resting-state FC, along with other subjects' task activations, was used to predict the held-out individual's cognitive task activations. Specifically, task activations were averaged across all subjects except the held-out subject and then the activity flow mapping procedure was applied along

with the held-out subject's resting-state FC. This allowed us to quantify the likely role of that individual's intrinsic connections (as estimated by resting-state FC) in shaping cognitive task activations.

In a separate analysis, we sought to further test the conclusion that resting-state FC describes individualized routes of activity flow that shape individual differences in cognitive task activations. This involved removing subject-general patterns from resting-state FC and each task's activations before implementing the activity flow mapping procedure. This better isolates the subject-specific FC and activation patterns, allowing us to better assess prediction accuracy of these patterns. Subject-general patterns were identified as the first principal component across subjects. These subject-general patterns were then regressed out of each subject's FC and activation patterns. This approach was similar to the task-specific pattern isolation approach described above. After removing the subject-general activation patterns, cross-subject activation similarity dropped from $r = 0.20$ on average to $r = -0.01$ on average. Cross-subject resting-state FC similarity dropped from $r = 0.15$ to $r = -0.009$ on average.

Statistics. All statistical inferences with empirical fMRI data that produced *P*-values were made using two-tailed one-sample *t*-tests relative to 0 ($n = 100$; degrees of freedom: 99) or, where indicated, permutation tests. Pearson correlation (r) was used as a measure of pattern similarity, with *P*-values only calculated for group-level inferences using two-sided one-sample *t*-tests on the Fisher's *z*-transformed *r*-values. All Fisher's *z*-transformed *r*-value distributions were confirmed to be approximately normally distributed using histograms and Q-Q plots. Cross-task average predicted-to-actual similarities were reported as the primary results, with single-task predicted-to-actual similarities reported to confirm the primary results were not driven by a subset of tasks. *P*-values were calculated for the cross-task average predicted-to-actual similarities, rather than for each task separately, to reduce the number of reported *P*-values and therefore the multiple comparisons needing to be corrected. The same statistical approach was used for the computational model analyses, with the exception of a single analysis with non-normally distributed comparisons that was also analyzed using Spearman's rank correlation (which does not assume a normal distribution).

Note that data collection and analysis were not performed blind to the conditions of the experiments. Other than selection of the "100 Unrelated Subjects" subset of the HCP data (see "Data Collection" above), no subjects or data points were excluded from analysis.

The average-then-compare *r*-values are reported as effect sizes of pattern similarity (squaring their values produces percent linear variance explained), rather than their associated *P*-values. Their *P*-values were not calculated due to the compared patterns possibly not being normally distributed (a requirement for Pearson correlation *P*-values but not effect sizes), as well as this not being a population (random effects) statistical inference. *P*-values below 0.00001 were reported as $P < 0.00001$ based on the convention that typical data analysis approaches likely do not have a level of precision consistent with such small values, such that reporting these small values would be misleading. Exact *P*-values can be calculated based on the reported *t*-values.

Correcting for multiple comparisons for each analysis was unnecessary (and not possible) due to lack of thresholding and calculation of only a single *P*-value per analysis. As an additional level of statistical caution, however, we used a conservative Bonferroni correction for multiple comparisons across all calculated *P*-values reported in this study (12 total). This revealed an uncorrected $P < 0.004$ threshold. All *P*-values reported as statistically significant were below this threshold, such that all significant *P*-values were statistically significant ($P < 0.05$) after correcting for multiple comparisons across all analyses.

Activity flow mapping with global signal regression. We chose to not use GSR for the primary analyses, due to controversy around this resting-state FC preprocessing step. Specifically, GSR is known to introduce anticorrelations into FC graphs[59]. However, there is evidence that some of the introduced anticorrelations are real[66], that GSR may increase the accuracy of some FC patterns[67] and that GSR reduces the impact of motion artifacts[68]. Therefore, we also applied activity flow mapping to data that had been preprocessed using GSR, testing whether the primary conclusions are unchanged when including this preprocessing step. We focused on the correlation-based FC results, given that multiple-regression FC already implicitly removes the global signal by controlling for signals in all other regions. Note that unlike GSR, however, multiple-regression FC does not (via averaging all time series into a global signal time series) regress out portions of regions' time series from themselves. Thus, multiple-regression FC may reduce introduction of negative FC (but see ref. 69).

As expected, we found that the primary conclusions were unchanged (and even slightly improved) when using GSR. Specifically, we found (as reported in "Results") that the predicted activation patterns were similar to the actual activation patterns with GSR: cross-task average $r = 0.50$, $t_{99} = 45.32$, $P < 0.00001$. This was slightly better than the results when not using GSR (average $r = 0.48$). Again, similar to the results obtained when not using GSR, these correlations were higher when comparisons were computed after averaging the predicted and actual activation patterns across subjects: cross-task average $r = 0.73$ (53% variance explained). This was again slightly better than the results when not using GSR (average $r = 0.66$). Overall these results support the conclusion that the choice of whether or not to use GSR does not substantially affect the outcome of activity flow mapping.

Code availability. The code used for activity flow mapping and multiple-regression FC is available on our lab website (http://www.colelab.org/#resources), as **Supplementary Software** and on GitHub (https://github.com/ColeLab/actflowmapping).

A **Supplementary Methods Checklist** is available.

51. Ermentrout, B. Phase-plane analysis of neural activity. in *The Handbook of Brain Theory and Neural Networks* (ed. Arbib, M.A.) 732–738 (MIT Press, 1998).
52. Hopfield, J.J. Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA* **81**, 3088–3092 (1984).
53. Van Essen, D.C. *et al.* The WU-Minn Human Connectome Project: an overview. *Neuroimage* **80**, 62–79 (2013).
54. Faul, F., Erdfelder, E., Buchner, A. & Lang, A.-G. Statistical power analyses using G*Power 3.1: tests for correlation and regression analyses. *Behav Res* **41**, 1149–1160 (2009).
55. Ugurbil, K. *et al.* Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. *Neuroimage* **80**, 80–104 (2013).
56. Smith, S.M. *et al.* Functional connectomics from resting-state fMRI. *Trends Cogn. Sci.* **17**, 666–682 (2013).
57. Glasser, M.F. *et al.* The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* **80**, 105–124 (2013).
58. Cox, R.W. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* **29**, 162–173 (1996).
59. Murphy, K., Birn, R.M., Handwerker, D.A., Jones, T.B. & Bandettini, P.A. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? *Neuroimage* **44**, 893–905 (2009).
60. Destrieux, C., Fischl, B., Dale, A. & Halgren, E. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *Neuroimage* **53**, 1–15 (2010).
61. Wig, G.S., Schlaggar, B.L. & Petersen, S.E. Concepts and principles in the analysis of brain networks. *Ann. NY Acad. Sci.* **1224**, 126–146 (2011).
62. Cohen, A.L. *et al.* Defining functional areas in individual human brains using resting functional connectivity MRI. *Neuroimage* **41**, 45–57 (2008).
63. Gordon, E.M. *et al.* Generation and evaluation of a cortical area parcellation from resting-state correlations. *Cereb. Cortex* **26**, 288–303 (2016).
64. Yeo, B.T.T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).
65. Jolliffe, I.T. A note on the use of principal components in regression. *Appl. Stat.* **31**, 300–303 (1982).
66. Chai, X.J., Castañón, A.N., Öngür, D. & Whitfield-Gabrieli, S. Anticorrelations in resting state networks without global signal regression. *Neuroimage* **59**, 1420–1428 (2012).
67. Fox, M.D., Zhang, D., Snyder, A.Z. & Raichle, M.E. The global signal and observed anticorrelated resting state brain networks. *J. Neurophysiol.* **101**, 3270–3283 (2009).
68. Power, J.D. *et al.* Methods to detect, characterize, and remove motion artifact in resting state fMRI. *Neuroimage* **84**, 320–341 (2014).
69. Smith, S.M. The future of FMRI connectivity. *Neuroimage* **62**, 1257–1266 (2012).