Title
A dynamic gradient architecture generates brain activity states

Authors
Brown, Jesse A
Lee, Alex J
Pasquini, Lorenzo
et al.

Peer reviewed

A dynamic gradient architecture generates brain activity states

Jesse A. Brown[*],

Alex J. Lee,

Lorenzo Pasquini,

William W. Seeley

Memory and Aging Center, Department of Neurology, University of California, San Francisco, CA, USA

Abstract

The human brain exhibits a diverse yet constrained range of activity states. While these states can be faithfully represented in a low-dimensional latent space, our understanding of the constitutive functional anatomy is still evolving. Here we applied dimensionality reduction to task-free and task fMRI data to address whether latent dimensions reflect intrinsic systems and if so, how these systems may interact to generate different activity states. We find that each dimension represents a dynamic activity gradient, including a primary unipolar sensory-association gradient underlying the global signal. The gradients appear stable across individuals and cognitive states, while recapitulating key functional connectivity properties including anticorrelation, modularity, and regional hubness. We then use dynamical systems modeling to show that gradients causally interact via state-specific coupling parameters to create distinct brain activity patterns. Together, these findings indicate that a set of dynamic, intrinsic spatial gradients interact to determine the repertoire of possible brain activity states.

Keywords

Gradients; Functional connectivity; Dynamical systems; Dimensionality reduction; Global signal

[*]Corresponding author. jesse.brown@ucsf.edu (J.A. Brown).

1.   Introduction

The brain is capable of reaching a diverse but constrained range of activity states (Allen et al., 2014; Betzel et al., 2016; Vidaurre et al., 2017; Pasquini et al., 2020; Lynn et al., 2021). Whole-brain activity undergoes relatively minor reconfigurations between task-free and task-engaged conditions, shaped primarily by an intrinsic network architecture (Cole et al., 2014; Gratton et al., 2018). This suggests that a low-dimensional set of neuroanatomical systems may be involved in generating the repertoire of states (Glomb et al., 2019; Saggar et al., 2018; Shine et al., 2019). Recent work indicates that these dimensions reflect gradients of continuous spatial variation in regional functional connectivity (Haak et al., 2018; Margulies et al., 2016; Zhang et al., 2019) and cytoarchitecture (Burt et al., 2018; Paquola et al., 2019; Wang, 2020). These gradients capture both the major sensorimotor-to-association axis of cortical organization (Sydnor et al., 2021) and finer-grained delineations between subregions supporting distinct cognitive processes like in the anterior temporal lobe (Faber et al., 2020). While gradients appear to reflect regional heterogeneity in circuit properties that sculpt brain activity dynamics (Kong et al., 2021), it remains unclear how spatially fixed functional gradients can dynamically coordinate to generate distinct brain activity states.

This study addresses two key questions about spatial gradients. First, do functional gradients represent distinct anatomical systems, akin to different ascending projection pathways, or do they instead reflect epiphenomenal anatomical features lacking causal influence? Second, can multiple gradients interact to generate state-specific activity patterns? We address these questions by using dimensionality reduction to derive an intrinsic latent space where dimensions represent distinct, spatially fixed gradients whose amplitudes dynamically fluctuate. Our analysis strategy has two major components (Fig. 1). First, we derive the latent space and validate that it recapitulates key functional connectomic properties including the brain-wide global signal (Fox et al., 2009), anticorrelation between large-scale networks (Fox et al., 2005), the presence of discrete functional modules and hub regions (Sporns and Betzel, 2016), and correspondence with spatial gene expression patterns (Richiardi et al., 2015). Second, we use the latent space to investigate how gradients dynamically interact to generate specific activity and connectivity states, either in different individuals in the task-free state or in specific cognitive tasks from the Human Connectome Project (Barch et al., 2013).

Our dimensionality reduction approach differs from the conventional method for deriving functional connectivity gradients (Hong et al., 2020; Vos de Wael et al., 2020). By performing principal component analysis (PCA) directly on BOLD timeseries data, rather than on the functional connectivity matrix, we find both previously described functional gradients and importantly, a novel principal activity gradient. This gradient appears to underlie the global signal of brain activity and best matches the primary spatial gene expression component. We show that this expanded set of dynamic activity gradients represent intrinsic systems that are stable across individuals and task-free or task-engaged scanning conditions. We then use a dynamical systems model and perform simulations to demonstrate how this set of gradients can transiently couple into state-specific configurations, suggesting a novel brain mechanism for dynamically generating diverse activity and functional connectivity states.

2.   Materials and methods

2.1.   Subjects and data

200 unrelated subjects were selected from the Human Connectome Project (HCP) 1200 Subjects Data Release with available resting (task-free) and task fMRI data from a 3T MRI scanner (https://db.humanconnectome.org/data/projects/HCP_1200). Informed consent was obtained for each individual by the HCP consortium. The HCP complied with all relevant ethical regulations. This study agreed to the Open Access Data Use Terms (https://www.humanconnectome.org/study/hcp-young-adult/document/wu-minn-hcp-consortium-open-access-data-use-terms) and was exempt from the UCSF IRB because investigators could not readily ascertain the identities of the individuals to whom the data belonged. Task-free state scans were 14.4 minutes long with a repetition time (TR) of 720 ms, resulting in 1200 fMRI volumes per scan. We divided these subjects into 100-subject discovery and validation datasets (56 female/44 male in the discovery dataset, mean age=28.9 years; 50 female/50 male in the validation dataset, mean age= 28.6 years). For task-free state fMRI data, we used the left-right phase encoded, minimally preprocessed scans with motion correction and FIX-ICA denoising. More detailed scanning parameters and preprocessing procedures have been described in detail elsewhere (Glasser et al., 2013). We used FSL (https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/) and AFNI (https://afni.nimh.nih.gov/) for additional fMRI preprocessing. The first 5 volumes for each fMRI scan were dropped to allow scanner stabilization. Scans were spatially downsampled to 2 mm^3 spatial resolution in MNI152 standard space based on the nonlinear 6th generation template distributed with FSL. Scans were bandpass filtered in the 0.008-0.15Hz frequency range and then normalized in each voxel across time to have zero mean and a standard deviation of one. fMRI volumes were temporally concatenated for all subjects in the discovery or validation datasets (119500 volumes each).

2.2.   Latent space and spatial gradient derivation

We derived a brain activity latent space by performing PCA on the data matrix of 119500 task-free fMRI volumes by 163049 gray matter voxels from cortex, subcortex, and cerebellum (Fig. 2A). This data matrix required a large amount of computer memory (145GB RAM), limiting the sample size to 100 subjects. In this latent space, each principal component represented a temporally orthogonal latent dimension and the entire space collectively encapsulated the full range of brain activity states present in this dataset. The voxel-wise component loadings (eigenvectors scaled by their corresponding eigenvalues) represented each voxel's weight on that component, and the spatial map for each component is henceforth referred to as a gradient. The PCA scores represented the slope of each gradient at each timepoint, which we refer to as the gradient slope timeseries (Fig. 2B and Video 1). Trajectory plots were used to illustrate the temporal evolution of gradient slopes for two latent dimensions simultaneously.

A primary aim of this study was to assess if this latent space reflected intrinsic spatial gradients that were consistent across individuals and task conditions. We approached this by projecting each of the other fMRI datasets in this study – task-free validation dataset scans and task discovery/validation dataset scans – into the task-free discovery latent space.

The only case where we instead derived independent dataset-specific latent spaces was when assessing gradient spatial consistency across subjects and states.

2.3. Gradient analysis and reproducibility

Gradient analyses were made more computationally tractable by spatially downsampling the voxel-wise gradient maps. Maps were averaged within 273 regions of interest by combining a parcellation of 210 cortical regions and 36 subcortical regions from the Brainnetome atlas (Fan et al., 2016) (http://www.brainnetome.org/) and 27 cerebellar regions from the SUIT atlas (Diedrichsen, 2006) (http://www.diedrichsenlab.org/imaging/suit.htm). When measuring the variance explained in each region by each gradient, gradients were derived independently by running PCA on the $[119500 \times 273]$ or $[119500 \times 210]$ region mean BOLD timeseries. We measured the variance explained by each principal component to determine the relative effect size of each gradient on the BOLD signal. The polarity of each gradient was calculated as the percentage of the 273 regions that had positive or negative mean weights.

Gradient reproducibility across individuals and conditions was assessed by comparing the discovery dataset spatial gradient maps to independently derived validation dataset gradient maps. Here, PCA was independently run on the 119500 volumes from the 100 validation subjects. Spatial gradient maps from the discovery and validation datasets were downsampled to 273-region average weight vectors which were then correlated with each other for the first 12 components. Procrustes alignment was used to align the $[273 \times 12]$ validation gradient weight matrix to the discovery data, after which spatial correlation was recalculated (Vos de Wael et al., 2020). This same procedure was used to compare gradient maps derived independently from the task fMRI discovery data (116100 volumes) to the task-free gradients. We assessed gradient spatial reliability by comparing the spatial gradient patterns derived from PCA of voxel-wise or region-wise timeseries using either: 1) 273 regions from the Brainnetome and SUIT atlases, or 2) 232 regions, combining 200 cortical regions from the Schaefer 17 network atlas and 32 subcortical regions from the Melbourne atlas at scale 2 (Schaefer et al., 2018; Tian et al., 2020).

2.4. Timeseries analysis

Regional BOLD activity at a given timepoint could be calculated by multiplying the region's weight on each gradient by the slope of the corresponding gradient (see Video 1). The formula for regional BOLD signal in latent space is:

$$BOLD_{i,t} = \sum_{k=1}^{d} \beta_{i,k} S_k$$

where $\beta_{i,k}$ is the current region's weight on a given gradient (PCA loading), S is the current timepoint's gradient slope (PCA score) on a given dimension, k is the dimension number out of d total dimensions, i is the current region, and t is the current timepoint. The value of decomposing a region's BOLD signal into its piecewise contributions from each gradient is to determine if certain gradients play an outsize role in causing a region to have particularly

high or low BOLD signal. We next discuss how functional connectivity relates to regional gradient weights and gradient slopes. The BOLD signal functional connectivity between two regions is typically measured by the correlation, which is the covariance divided by the product of the region's standard deviations. In the PCA-derived latent space, the covariance between a pair of regions is:

$$cov_{ij} = \frac{\sum_{t=1}^{n} \begin{bmatrix} S_1 S_1 & \cdots & S_1 S_k \\ \vdots & \ddots & \vdots \\ S_k S_1 & \cdots & S_k S_k \end{bmatrix} \circ \begin{bmatrix} \beta_{i,1}\beta_{j,1} & \cdots & \beta_{i,1}\beta_{j,k} \\ \vdots & \ddots & \vdots \\ \beta_{i,k}\beta_{j,1} & \cdots & \beta_{i,k}\beta_{j,k} \end{bmatrix}}{n}$$

Where $S_{k,t}$ is the gradient slope for latent dimension k at timepoint t, and β is the gradient weight for region i or j on gradient k. This decomposition of the BOLD covariance into the interactions between pairs of gradients is also valuable when determining if certain gradients make an especially strong contribution to functional connectivity. Higher-order regional interactions such as edge functional connectivity (Faskowitz et al., 2020) can also be decomposed in this fashion. Computing the correlation from the covariance requires the standard deviation, which is the square root of the variance. The variance for a region with a demeaned timeseries is calculated by multiplying the region's weight on each gradient by the slope of the corresponding gradient at each timepoint, taking the square, and dividing by the number of timepoints:

$$\sigma_i^2 = \frac{\sum_{t=1}^{n} \sum_{k=1}^{d} (\beta_{i,k} S_k)^2}{n}$$

$$\sigma_j^2 = \frac{\sum_{t=1}^{n} \sum_{k=1}^{d} (\beta_{j,k} S_k)^2}{n}$$

Finally, the correlation between region i-region j is determined by:

$$corr_{ij} = \frac{cov_{ij}}{\sigma_i \sigma_j}$$

2.5. Task fMRI analysis

Preprocessed task fMRI data for the four tasks from the HCP were analyzed (working memory, motor, language, emotion) (https://db.humanconnectome.org/data/projects/HCP_1200). We used left-right phase encoded scans. Several postprocessing steps were applied to the task fMRI data, which had not undergone FIX-ICA denoising. The first five scans were removed. Nuisance parameters were estimated for the six motion parameters, temporal derivatives, squares, the white matter timeseries extracted using a mask of the highest probability cortical white matter according to the FSL tissue prior mask, and the CSF timeseries extracted using a mask in the central portion of the lateral ventricles (Satterthwaite et al., 2013). Scans were bandpass filtered in the 0.008-0.15Hz frequency range to match the task-free state data. The 26 nuisance parameters were also bandpass

filtered and regressed out the data. Global signal regression was not performed for the main dataset and was only used for a supplementary analysis. Scans were then voxel-wise normalized across time to a mean of zero and a standard deviation of one. These images were projected into the task-free PCA space from the discovery dataset.

When determining task-specific activation patterns, the following regressors were used for each task: for working memory, the 2-back and 0-back conditions, merging blocks from the faces, places, tools and body parts stimulus blocks (Barch et al., 2013); for the motor task, the active condition combining the right hand, left hand, right foot, left foot, and tongue blocks, and the fixation condition; for language, the story and math conditions; and for emotion, the faces and shapes conditions. Task condition block regressors were convolved with a hemodynamic response function using the 'spm_get_bf' function in SPM12 (https://www.fil.ion.ucl.ac.uk/spm/software/spm12/).

2.6. Genetic spatial correlation

We compared each gradient map to Allen Human Brain spatial gene expression patterns using the 'abagen' package (https://github.com/rmarkello/abagen) (Arnatkevicǐ bt et al., 2019; Hawrylycz et al., 2012). We used default options including for donors (all), tolerance (2 mm), collapsing across probes (diff_stability), and intensity-based filtering threshold (0.5). Expression data was available for 261 out 273 Brainnetome and SUIT regions from both hemispheres for 15655 genes. All non-cortical regions were eliminated because of substantial differences in subcortical expression values, which would hamper brain-wide spatial correlation estimates, leaving 202 regions. Data-driven filtering was used to remove regions with outlying expression values. Using K-means clustering we identified an outlying cluster with 6 regions, which were removed to give the final $[15655 \times 196]$ matrix of expression values. For each of the first six gradients in the discovery or validation datasets, we calculated the spatial Pearson correlation between each 196-region gene expression vector and the 196-region gradient weight vector. We defined the statistical significance threshold based on two criteria. First, the spatial correlation had to survive a Bonferroni corrected threshold of $p < 0.05$ / 15655 genes / 6 gradients = 5.32×10^{-7} for at least one gradient in both the discovery and validation datasets. Second, the spatial correlation strength had to significantly exceed that of 5000 simulated surrogate "null" gradient maps with preserved spatial autocorrelation (Burt et al., 2020). These surrogate gradient maps were estimated using BrainSMASH (https://brainsmash.readthedocs.io/en/latest/). We used the Euclidean distance between the 273 regions. The surrogate p-value was determined by measuring the fraction of trials out of 5000 in which a surrogate gradient map had a stronger correlation to a given gene expression map than did the true gradient map. A false discovery rate-corrected $p < 0.05$ in both the discovery and validation datasets was required for determining significance. Finally, the principal component of spatial gene expression was derived by performing PCA on the $[15655 \times 196]$ regional expression matrix and obtaining the PCA scores for each of the 196 regions on the first component.

2.7. Latent space trajectory analysis

Here we attempt to demonstrate a link between the BOLD timeseries and latent space representations of regional BOLD activity, so that regional BOLD activity and covariance

can be visualized as the magnitude and direction of a latent space trajectory. We focus this analysis on a two-dimensional latent space but this can be extended to *n* dimensions. It can be shown that the formula for regional BOLD activity at a given timepoint is mathematically equivalent to a latent activity trajectory's current position in latent space. Given that regional BOLD activity is equal to a region's weight on two gradients multiplied by the current slopes of those gradients (Fig. 2 and Video 1), the region's BOLD signal at each timepoint can be visualized for two dimensions of latent space as an "activity field" (Results, Gradient basis for functional modularity and hubness). A region's BOLD activity at a specific timepoint is mathematically equivalent to the (*x,y*) position of the latent trajectory in the regional activity field. In this field, an optimal trajectory direction exists that will maximize this region's BOLD activity relative to all other regions. We next extend this approach to visualize the BOLD covariance of two regions, i.e. the functional connectivity. When considering two regions, their covariance at a given timepoint – their instantaneous covariance – is simply the product of their BOLD signals at that timepoint. This can be represented in latent space as the element-wise product of the two regional activity fields. In this case, the region A-region B BOLD covariance is mathematically equivalent to the (*x,y*) position of the latent trajectory in the regional "covariation field". Here, the optimal trajectory direction will maximize the covariance of these two regions relative to all other pairs of regions. The covariance across a scan can be measured by taking the sum of the trajectory locations for each timepoint on the covariation field, and dividing by the total number of timepoints. The correlation coefficient can be derived by dividing the covariance by the product of the standard deviations for each regional BOLD signal, where standard deviation in latent space is equivalent to the mean trajectory width along a given dimension. This approach was subsequently used to investigate functional modularity and regional hub properties.

2.8. Functional modularity and regional hubness

In the two-dimensional latent space, we computed the optimal trajectory direction for each region, which had angles between $0°–360°$. Based on our hypothesis that the angular distribution of regions in this space would correspond to functional modularity, we next assigned regions to modules using the conventional approach (Rubinov and Sporns, 2010). The task-free functional connectivity matrix was derived for the discovery dataset by 1) reconstructing each region's timeseries based on Gradient 2 and Gradient 3 (see Timeseries analysis) and 2) computing the pairwise Pearson correlations. Graph theory analyses were run using the Brain Connectivity Toolbox (BCT; https://sites.google.com/site/bctnet/). Networks were thresholded to keep the strongest 10% of edges, similar to other studies of functional network modularity and hub properties (Power et al., 2011; Bertolero et al., 2015). Modularity was determined using the Louvain algorithm (Blondel et al., 2008) as implemented in BCT, running 1000 iterations with the default parameters, choosing the partition that maximized the Q value. We examined two network properties related to modularity and hubness: within-module degree Z-score and participation coefficient (Guimerà and Nunes Amaral, 2005). Within-module Z-score is commonly used to identify provincial regions, which connect primarily to regions in the same module, while participation coefficient is used to designate connector regions that link to regions in multiple modules (Power et al., 2013). We defined regions in 85[th] percentile of participation

coefficient as connector regions and regions in the 85th percentile of within-module degree Z-score as provincial regions, similar to previous approaches (Gordon et al., 2018). 82/273 regions were identified as provincial or connector regions and the remainder were labeled as 'other'. Each region's neighbors were identified as the top 10 most highly functionally correlated regions.

We statistically compared the angular position of different regions using the Watson-Williams test (Berens, 2009). This test is a circular analogue of the two-sample t-test and assesses whether the mean directions of two groups are identical or not. The mean angle to neighboring regions was compared amongst provincial regions, connector regions, and the remaining non-provincial/connector regions.

2.9. Differential equation modeling

We modeled the gradient temporal dynamics with differential equations describing the continuous influence of each activity gradient on one another, using a data-driven strategy to estimate the coupling parameters between gradients (Brunton et al., 2016). We first focused on the task-free fMRI data. For each fMRI timepoint, we used the gradient slope timeseries (g) to estimate the first and second temporal derivatives of the slope timeseries (g', measuring the rate of the change in the gradient-associated BOLD signal, and g''). The derivatives were calculated using the 'gradient' function in MATLAB. We then used a linear regression model for each of the first six gradients to estimate gradient slope second derivative as a function of all gradients' slopes and slope first derivatives. The parameter estimates (which we henceforth refer to as coupling parameters) for the 13 terms from each regression - the slope and slope first derivative terms for each gradient and an intercept - were then used to define a system of six differential equations:

$$G1'' = \beta_{1,0} + \beta_{G1,1}G1 + \beta_{G1',1}G1' + \beta_{G2,1}G2 + \beta_{G2',1}G2' + ... + \beta_{G6,1}G6 + \beta_{G6',1}G6'$$
$$G2'' = \beta_{2,0} + \beta_{G1,2}G1 + \beta_{G1',2}G1' + \beta_{G2,2}G2 + \beta_{G2',2}G2' + ... + \beta_{G6,2}G9 + \beta_{G6',2}G6'$$
$$...$$
$$G6'' = \beta_{6,0} + \beta_{G1,6}G1 + \beta_{G1',6}G1' + \beta_{G2,6}G2 + \beta_{G2',6}G2' + ... + \beta_{G6,6}G6 + \beta_{G6',6}G6'$$

Gradient timeseries simulations were performed using the system of differential equations for gradients 1-6. This system of equations was solved numerically using the 'ode45' function in MATLAB. We performed three sets of simulations: 1) task-free, using group level coupling parameters, 2) task-free, using individual subject coupling parameters, and 3) task-specific, using group level task coupling parameters. To initiate the task-free group level simulations, we randomly selected 1000 timepoints from the task-free state scans in the discovery dataset. For each of these timepoints, each gradient's slope and first derivative were used as the initial conditions to solve the differential equations. We then simulated the gradient slope timeseries for 200 future timepoints, equivalent to 2.4 minutes. The [6 × 200] gradient timeseries were matrix multiplied by the [273 × 6] gradient weights to obtain [273 × 200] region timeseries, from which [273 × 273] functional connectivity matrices were computed. These 1000 matrices were averaged to produce the simulated functional connectivity (FC) matrix. For the actual data, the functional connectivity matrices were derived from the six gradients' timeseries.

To assess individual differences in task-free functional connectivity, each subject's task-free coupling parameters were individually calculated from their $[6 \times 1195]$ gradient timeseries. Here we initialized each simulation with identical initial conditions to hold this factor constant across subjects, then simulated the gradient slope timeseries for 200 timepoints. With the simulated gradient timeseries, we derived the region timeseries and the resulting functional connectivity matrix. We assessed the simulation accuracy in two ways: first, by comparing the simulated FC matrices to each subject's actual FC matrix, derived from the six-dimensional gradient timeseries, using Pearson correlation; and second, by comparing the simulated FC-derived modular community membership to each subject's actual modularity using the Rand index (Rand, 1971).

When simulating task fMRI gradient timeseries, we first labeled volumes in each task as occurring during the active task condition (Working memory: 2-back; Motor: active; Language: story; Emotion: faces) or baseline task condition (Working memory: 0-back; Motor: fixation; Language: math; Emotion: shapes) when the HRF-convolved task waveform was greater than 0.5. Group-level coupling parameters were calculated for each condition from the discovery dataset scans. Simulations for each of these eight conditions were performed using the same set of initial conditions from the 1000 randomly selected timepoints from the task-free scans. The simulated gradient timeseries for each condition were used to derive simulated regional BOLD timeseries. Simulated task activation maps were computed by subtracting the average region BOLD value for the active task condition minus the baseline task condition. FC matrices were computed for each condition. The resultant activation maps and functional connectivity matrices were compared to the actual activation maps and matrices from the validation dataset. We compared the simulated dynamic activity for each active task condition by subtracting the simulated task-free FC matrix from each simulated task FC matrix . This resulted in four FC matrices: Working memory 2-back vs. task-free, Motor active vs. task-free, Language story vs. task-free, and Emotion faces vs. task-free. Each of these contrasts was compared to the actual FC subtraction matrices, e.g. actual Working memory 2-back vs. actual task-free. We finally estimated the specificity of the simulated activity patterns by computing the partial correlation of the simulated and actual FC matrices for each task, controlling for the actual FC in the other three task conditions.

3. Results

3.1. Low-dimensional brain activity latent space

We first evaluated whether the brain activity latent space was able to explain the majority of the variation in BOLD activity with a small number of dimensions. When reconstructing each BOLD timeseries for 210 cortical brain regions, the first six components explained 64.8% of BOLD activity variance and 59.5% when including subcortical and cerebellar regions (Fig. 3A). Component 1 explained 35.8% of the brain-wide BOLD signal variance, 4.6 times more than the 7.8% variance explained by component 2. Across individual subjects, the first six components explained an average of 63.7% of variance in cortical regions (mean=63.7% ± 8.1%, range: 39.4% - 79.8%). For comparison, we derived functional connectivity components (gradients) from the group-level functional connectivity

matrix using PCA and found that they explained more variance with fewer components (48% with one component, 89% with six components) than the BOLD timeseries PCA components, similar to previous studies (Zhang et al., 2019; Hong et al., 2020). This was likely due to a combination of three factors: many fewer observations (273 vs. 119,500), static versus dynamic data, and group-averaged versus individual data. Here, our subsequent analysis focused on the first six BOLD activity latent dimensions and their associated gradient maps, based on the criterion that each of these components explained at least 3% of the total variance and in line with other studies of low-dimensional brain activity and functional connectivity (Shine et al., 2019; Margulies et al., 2016).

The polarity of regional "positive" versus "negative" weights on a gradient was arbitrary, signifying only whether regions were correlated or anti-correlated with each other on that gradient. The spatial weights for Gradient 1 were positive across 98.9% of the gray matter (Fig. 3B)" albeit with topographically varied weights that were highest in the primary somatomotor, visual, and auditory cortices. The gradient slope timeseries associated with this dimension had a near-perfect correlation with the global gray matter signal ($r = 0.99$, $p < 0.00001$), a major influence on the estimated strength of functional connectivity (Fox et al., 2009). Gradient 1 was the only "unipolar" gradient, as no other gradient had a ratio of more than 1.43 times as many voxels (59%) with positive versus negative weights. We ensured that this unipolar component was not a statistical artifact by voxel-wise demeaning BOLD activity across time before PCA. Crucially, Gradient 1 was not apparent when global signal regression was performed before BOLD timeseries PCA (Supplementary Fig. S1A).

The activity gradients demonstrated both similarities and key differences with the set of functional connectivity gradients described by Margulies and colleagues (referred to as FCG1-FCG5) (Margulies et al., 2016). The strongest spatial matches between the known functional connectivity gradients and the activity gradients reported here were: FCG1/ Gradient 2, $r = 0.93$; FCG2/Gradient 4, $r = 0.83$; FCG3/Gradient 3, $r = 0.93$; and FCG4/ Gradient 5, $r = 0.79$ (all $p < 0.00001$). Gradient 1′s strongest spatial match was to FCG1, though with a substantially weaker spatial correlation ($r = 0.61$) and without FCG1′s bipolar distribution of regional gradient weights. Thus, the activity gradients had strong correspondence with canonical functional connectivity gradients, with Gradient 1 being an important exception.

Gradient 2′s strong match with FCG1, the principal functional connectivity gradient (Margulies et al., 2016), indicated that this sensory-to-cognitive system is equally apparent when performing dimensionality reduction on continuous BOLD timeseries data. When measuring the mean weight in seven canonical functional brain networks (Yeo et al., 2011), Gradient 2′s most positive weights were in areas of the default mode network and frontoparietal network, while the most negative weights were in somatomotor, visual, and ventral attention networks (Fig. 3B). Dimension 3 resembled a task-positive (frontoparietal) to task-negative (default mode) gradient. The subsequent dimensions included oppositions between specific sensory modalities like the visual and somatomotor networks (Gradient 4), differential involvement of sub-components of larger super-systems like the default mode network (Gradients 2, 3, and 6) (Andrews-Hanna et al., 2010), and lateralized activity gradients (Gradient 6, Gradient 7 and 11 in Supplementary Fig. S2).

Gradient spatial patterns were also assessed in the cerebellum and subcortical regions. The cerebellar aspect of Gradient 2 and Gradient 3 (Supplementary Fig. S3) had strong correspondence with the two primary cerebellar functional connectivity gradients defined by Guell and colleagues (Guell 1/Gradient 2, r = 0.90; Guell 2/Gradient 3, r = 0.92; all $p < 0.00001$) (Guell et al., 2018). Gradient 2 spanned from cerebellar areas involved in transmodal (Crus II) to primary (lobule VIIIa) function, while Gradient 3 separated task-positive (lobule IX) and task-negative (lobule VIIb) associated cerebellar areas. The subcortical spatial patterns for Gradient 3 and Gradient 5 had moderate similarity to the subcortical functional connectivity gradient pattern used to derive the Melbourne subcortex atlas of the thalamus, basal ganglia, hippocampus, and amygdala (Tian et al., 2020) (Melbourne 1/Gradient 3, r = 0.69; Melbourne 1/Gradient 5, r = 0.72; all $p < 0.00001$). Thus, while the activity gradients explained less BOLD signal variance in subcortical/cerebellar regions than in cortical regions, regional gradient profiles were sufficient to demarcate known functional subdivisions.

We determined gradient reproducibility based on three criteria. First, maps were highly reproducible in the validation dataset. For the first three components from the discovery dataset, the spatial correlation with the corresponding gradient in the validation dataset was high (r = 0.996, r = 0.994, r = 0.987, all $p < 0.00001$). The first 12 gradient maps in the discovery dataset had a median spatial correlation of r = 0.99 (median $p < 0.00001$) to their best matched gradient in the validation dataset after Procrustes alignment (Fig. 4B). Second, we tested the evidence that the spatial gradient patterns were "intrinsic" and stable across cognitive states. The first 12 gradient maps in the task fMRI-derived latent space had a median spatial correlation of r = 0.89 (median $p < 0.00001$) to their best matched gradient in the task-free latent space after Procrustes alignment (Fig. 4C). When deriving gradient maps from each of the four tasks separately, the median spatial correlation was r = 0.80 (median $p < 0.00001$) to the best matched task-free gradients after Procrustes alignment, with the best match always being the correspondingly numbered gradient (Supplementary Fig. S4). Third, we assessed the spatial reliability of activity gradient maps when derived from voxel-wise or region-wise timeseries data. Gradient spatial patterns had high reliability for the first four components (r = 0.999, r = 0.992, r = 0.989, r = 0.969, all $p < 0.00001$) and moderate reliability for subsequent components (Supplementary Fig. S1B). A supplementary analysis with an alternative regional atlas found even higher correspondence between gradients derived from voxel-wise or region-wise BOLD timeseries data (Supplementary Fig. S1C). Overall, we found that by performing dimensionality reduction on BOLD time-series data, we could detect a set of activity gradients including both the canonical functional connectivity gradients and the unipolar Gradient 1 which explained most of the variance in brain activity. The gradients' spatial consistency across individuals, task-free or task-engaged states, and voxel or region-based timeseries suggests that they reflect intrinsic anatomical systems of brain functional organization.

3.2. Correspondence with spatial gene expression patterns

To assess the biological composition of each activity gradient, we compared each gradient's spatial similarity with genetic expression maps using the 15,655 genes from the Allen Human Brain Atlas (Methods, Genetic spatial correlation). While numerous relationships

between structural or functional gradients and spatial gene expression patterns have been established (Burt et al., 2018; Fulcher et al., 2019; Vogel et al., 2020; Shafiei et al., 2020; Huntenburg et al., 2021), we sought to determine how strongly each of the activity gradients defined here corresponded to gene expression patterns. Across the first six gradients, 1263 genes showed significant spatial correlations with at least one gradient, with correlation coefficients ranging between r = 0.35-0.77 (Supplementary Fig. S5A). The most striking correspondences were with Gradient 1, for which 1007 genes were significantly correlated. The most significant positive relationships between gene expression and Gradient 1 were *SEMA7A* (discovery/validation mean r = 0.73), *SCN1B* (r = 0.72), *ASB13* (r = 0.72), *ANK1* (r = 0.71), and *SCN1A* (r = 0.71), while the most negative relationships were *KCNG1* (r = −0.77), *ASCL2* (r = −0.74), *ANKRD6* (r = −0.71), and *PYDC1* (r = −0.71). Because such a large number of genes exhibited strong positive or negative spatial correlations with Gradient 1, we tested for a potential link to the primary spatial component of genetic expression variability, which is known to stratify sensory and association areas (Burt et al., 2018). There was a significant spatial correlation between Gradient 1 and the principal spatial component of gene expression (r = 0.70; Supplementary Fig. S5B), for which each of the aforementioned individual genes were strongly loaded (all loading absolute Z scores > 3). This indicated that the predominant sources of variability in BOLD activity and spatial gene expression are strongly linked. Substantially fewer genes were significantly correlated with the remaining gradients. For gradients 2-6, the number of significantly correlated genes was: Gradient 2: 4, Gradient 3: 0, Gradient 4: 242, Gradient 5: 0, Gradient 6: 10 (Supplementary Fig. S5A and **Supplementary Data**). The strongest gene/gradient relationships for Gradient 4 were *NTNG1* (r = 0.61), *VPS51* (r = 0.59), *TPI1* (r = −0.59), *AKAP8* (r = −0.59), and *PMPCB* (r = −0.58). Among the strongest other gene/gradient relationships was *CARTPT* on Gradient 2 (r = 0.55). Top genes associated with activity gradients were frequently linked to functional and structural properties likely to influence brain-wide activity patterns including sodium and potassium channels (e.g. *SCN1A*, *SCN1B*, and *KCNG1* on Gradient 1), excitation/inhibition balance (Brackenbury et al., 2013), thalamocortical connectivity (Carcea et al., 2014), and anatomical morphogenesis (Tsunemoto et al., 2018).

3.3. Gradient basis for functional modularity and hubness

We next characterized how interactions between multiple gradients reflect two key functional connectivity properties – modularity and hubness – and how individual differences in multi-gradient interactions give rise to different activity and connectivity patterns. For illustrative purposes, we focused on two example regions, the anterior cingulate cortex (ACC) and middle frontal gyrus (MFG), which belong to dissociable functional networks (Seeley et al., 2007). We focused on Gradient 2 and 3 (Fig. 5A), the first two bipolar gradients, to demonstrate how modularity reflects the spacing of regions along multiple gradients. We found that when brain activity was rendered as a latent space trajectory (Methods, Latent space trajectory analysis and Functional modularity and regional hubness), specific latent space trajectory directions maximized either MFG BOLD activity, ACC BOLD activity, or MFG-ACC functional connectivity (Fig. 5B). The ACC's activity was maximized by a mostly downward trajectory (Fig. 5B, **top**), requiring engagement of Gradient 3, while the MFG's activity was maximized by a down-right

trajectory (Fig. 5B, **bottom**) engaging both Gradients 2 and 3. MFG-ACC functional connectivity was maximized by a trajectory that bisected the angle for maximizing ACC or MFG activity. We confirmed that the individual subjects with the most positive or negative ACC-MFG functional connectivity had latent trajectories that were most aligned with or most orthogonal to the optimal covariation angle (Fig. 5B). This illustrates that specific combinations of engagement in each gradient, which are reflected by specific latent trajectory directions, are needed to maximize activity in a given region or the connectivity between two regions.

The existence of continuous spatial activity gradients may appear to be at odds with the presence of discrete modular brain networks, a major principle of brain functional organization (Sporns and Betzel, 2016). We therefore attempted to reconcile the gradient and modular perspectives by testing the hypothesis that non-uniform spacing of regions along a gradient would recapitulate modular boundaries. We determined the optimal latent trajectory direction in the two-dimensional latent space for maximizing BOLD activity in each brain region. The trajectory angles fully spanned the 360° of the latent space (Fig. 5C). Regions clustered with their contralateral homologues and other regions belonging to the same functional connectivity network, while regions that were diametrically opposed belonged to canonically anti-correlated networks. Based on this observation, we expected that functional connectivity modules derived from the functional connectome would correspond well with different angular ranges in latent space. We found that region module membership from the functional connectome corresponded exactly to the sequence of regions as grouped by optimal activity angle (Fig. 5C, **middle**). Consequently, provincial regions and connector regions were found to have different spacing to their neighboring regions. Provincial regions were significantly closer to their most strongly connected neighbors than other non-connector regions (angle for provincial regions: mean=1.7°±1.5°; other regions: mean=3.1° ±2.8°; $p < 0.00001$; Fig. 5D). Other non-connector regions in turn showed smaller angles to their neighbors than connector regions (connector regions: mean=3.6°±2.8°; $p = 2.3 \times 10^{-4}$). Thus, the presence of modularity and hub regions appears to be consistent with the non-uniform spacing of regions along gradients. Different latent trajectories reflect specific between-gradient coupling patterns which drive the differential activity of specific regions and modules.

3.4. Gradient coupling generates different brain activity states

Based on the observation that specific gradient engagement patterns reflect the activity and functional connectivity of specific regions and networks, we hypothesized that a generative model should be able to simulate gradient dynamic activity during different FC states. This motivated our use of a dynamical systems model that uses differential equations to describe how each gradient's activity (i.e. slope) changes as a function of its current activity level and influence from other gradients' activity levels (Fig. 6A) (Breakspear, 2017). Specifically, we used a coupled system of second-order differential equations to model the activity of gradients 1-6. Using this model, we tested whether specific gradient coupling parameters were sufficient to produce either task-free or task-specific FC patterns. We first estimated the coupling parameters – a measurement of how much influence gradients had on one another – in the task-free data. Each gradient's second temporal derivative (g'') had a strong

negative relationship with its own slope (*g*, mean β=−0.06±0.006, mean t=−242.6±14.8, $p < 0.00001$), as is characteristic of an oscillating signal. In addition, all six gradients were significantly influenced by the slope of Gradient 1 (mean absolute t=17.4, $p < 0.00001$; Fig. 6B, column 2) and the first derivative of Gradients 1-4 (*g'*, mean absolute t=19.4/18.8/17.7/23.5, $p < 0.00001$; Fig. 6B, columns 3/5/7/9). This demonstrated that a given gradient's rate of change is dependent on both its own level of engagement and the rate of change in other specific gradients. Gradient 1 had the strongest collective influence on the ongoing activity of each gradient.

In the next step, we used these coupling parameters in our dynamical systems model to simulate the timeseries for gradients 1-6 (Fig. 6A). The dynamical model accuracy was evaluated with two analyses. First, we compared the real and simulated latent space trajectories from time *t* to *t*+1 for Gradient 1 and Gradient 2. We found that the simulated trajectories reproduced the two most salient aspects of the actual latent trajectories: counterclockwise flow and larger magnitude changes further from the origin (Fig. 6C). This captured the typical sequence of activity flow observed in the actual BOLD data: 1) a global signal peak, 2) activity increases in cortical association areas and decreases in primary cortical areas, 3) a global signal trough, and 4) activity increases in primary cortical areas and decreases in cortical association areas. Importantly, the success in predicting the magnitude of change in one gradient based on the level of engagement of another gradient suggests that the gradients exhibit mutual causal influence, in the statistical sense of causality where causes precede and help predict their effects (Seth et al., 2015).

In the second analysis, we used individual coupling parameters derived from each subject's task-free scan to assess whether a simulation with those parameters was sufficient to "reconstruct" that subject's actual FC pattern. As an example, we compared the simulations for two subjects (Fig. 7A). The different coupling parameters yielded distinctly different FC matrices that accurately matched the subject's actual FC patterns. Among the modularity differences that the model captured was a more widespread default mode network in Subject 2, indicating that different gradient mixtures were sufficient to produce unique network topographies. Overall, each subject's actual FC was significantly more similar to their own simulated FC than to simulated FC from the remaining subjects (t=18.93, $p = 3.12 \times 10^{-77}$; Fig. 7B), as was also the case for modularity (t=8.05, $p = 1.02 \times 10^{-15}$). This provided evidence that a dynamical model with a set of fixed spatial gradients and specific coupling parameters can generate diverse patterns of whole-brain functional connectivity and modularity.

Finally, we hypothesized that specific gradient coupling modes would be required to induce task-specific activation and functional connectivity patterns. We tested this by performing task-specific simulations of fMRI activity timeseries, using differential equations with task-specific coupling parameters. We found that each simulated task activation contrast map was significantly spatially correlated with the actual task activity contrast maps derived with a general linear model (working memory r = 0.75; motor r = 0.67; language r = 0.86; emotional faces r = 0.55; all $p < 0.00001$). The simulated functional connectivity matrix generated from state-matched coupling parameters was always most similar to the actual functional connectivity matrix from the matched task condition (partial correlations:

working memory, simulated vs. actual functional connectivity matrix r = 0.59; motor r = 0.72; language r = 0.83; emotional faces r = 0.62; all p < 0.00001; Fig. 8A/B). This set of experiments showed that differential equations with state-specific coupling parameters captured the most salient aspects of state-specific brain activity including the spatial activation patterns and the dynamics of the task activity timeseries that result in distinct functional connectivity patterns.

4. Discussion

This study contributes to understanding the functional-anatomical basis of low-dimensional brain activity. By performing dimensionality reduction on fMRI BOLD timeseries data rather than a static functional connectivity matrix, we derived a set of spatial activity gradients that dynamically oscillate in time. This approach revealed that the primary dimension of brain activity represents a novel unipolar spatial gradient ("Gradient 1") most strongly weighted in unimodal cortical areas. This gradient explained five times more variance than any other gradient, appeared to underlie the global signal, had the strongest correspondence with the principal spatial component of gene expression, and exerted the strongest influence on the dynamic activity of the other gradients. We then provided evidence that these activity gradients appear to be intrinsic. The same spatial gradients were identified in a validation dataset and in task fMRI data. Furthermore, we were able to generate plausible task-specific activation patterns with simulations where these intrinsic gradients adopted task-specific configurations. Finally, we used a dynamical systems model to understand how gradient interactions can generate different brain activity states. We found that coupling parameters determine the level of influence that gradients have on one another. A dynamical model with state-specific coupling parameters yielded accurate simulations of brain activity and functional connectivity in different individuals during the task-free state and during different tasks.

We performed several confirmatory analyses showing that this set of gradients are consistent with core functional connectomic phenomena. First, the global signal was apparent, driven by unipolar Gradient 1, as was strong anticorrelation between "task-positive" and "task-negative" brain networks (Fox et al., 2005), supported by the subsequent bipolar gradients. Second, these activity gradients had specific correspondences with known functional connectivity gradients (Margulies et al., 2016). Third, regions were spaced non-uniformly along gradients, resulting in functional modularity and the presence of hub regions (Sporns and Betzel, 2016). Fourth, gradients had strong spatial correlations with spatial gene expression patterns (Richiardi et al., 2015). Thus, gradients and functional networks may be considered as alternative representations of the same underlying brain activity.

4.1. Global gradients dynamically fluctuate

The primary activity gradient had a unipolar spatial pattern and an activity timeseries that mirrored the global signal of brain activity. The global signal is a source of ongoing controversy in fMRI literature, appearing to have a neuronal basis (Turchi et al., 2018), relating to individual differences in vigilance (Liu et al., 2017) and behavior (Uddin, 2020), but also associating with respiration or head motion (Chang and Glover, 2009; Power et al.,

2014). We used ICA-FIX denoised HCP fMRI data to minimize the impact of non-neural signals including respiration (Power et al., 2017; Glasser et al., 2019). We found that this gradient had strongest involvement of unimodal visual, somatomotor, and auditory areas, consistent with previous reports that the global signal has a heterogenous spatial topography (Liu et al., 2018b; Yang et al., 2017; Li et al., 2019; Orban et al., 2020). The unipolar nature of Gradient 1 is a likely reason why it has not previously been detected with conventional gradient discovery methods. Gradient 1 acts like a rising and falling tide, driving positive correlation between all areas. For this reason, only methods which derive latent dimensions of variability in *activity* will detect this gradient, unlike methods which identify latent dimensions of *functional connectivity*. Temporal dimensionality reduction of BOLD timeseries data has previously been applied using both PCA (Carbonell et al., 2011; Hong et al., 2020) and ICA (Smith et al., 2012; Glasser et al., 2018), though this approach notably differs from how spatial gradients are typically derived from fMRI data, by applying dimensionality reduction on the functional connectivity matrix (Vos de Wael et al., 2020). PCA and related dimensionality reduction methods find components that account for maximal variance in the dataset across observations. When applied to timeseries data, PCA can detect a global unipolar spatial factor which causes all regions to have high or low BOLD signal at different timepoints. By contrast, in a functional connectivity matrix where the time dimension is collapsed, all regions would share a common pattern of connectivity reflecting the influence of Gradient 1. This lack of region-to-region variability in connectivity would render this factor invisible to the variance-maximizing PCA algorithm. Importantly, Gradient 1 was not detected when performing PCA on BOLD timeseries data after global signal regression. Thus, processing methods should be carefully considered when deriving gradients from either BOLD activity timeseries or functional connectivity data.

Gradient 1 had the strongest spatial correlation with the principal spatial component of cortical gene expression, supporting its biological distinctiveness. The principal genetic expression component is known to have a strong relationship with the cortical myelination pattern that demarcates the borders between sensory and association areas (Burt et al., 2018). In the current study, the genes most positively associated with Gradient 1 included *SEMA7A* and *SCN1B*, which have known roles in excitation/inhibition balance and seizure disorders (Brackenbury et al., 2013; Carcea et al., 2014). Using the dynamical systems model, we found that Gradient 1′s activity level had the strongest influence on each other gradient's ongoing activity. This indicates that Gradient 1 had the largest causal influence in a statistical sense, arguing for its primacy in sculpting brain activity dynamics. Collectively, the primary gradient appears to represent a system with a distinct functional, anatomical, and gene expression fingerprint which may modulate global neuronal excitation/inhibition balance (Wang, 2020).

Gradient 2 strongly resembled the principal macroscale gradient of functional connectivity (Margulies et al., 2016). This gradient represents a sensory-to-cognitive axis with the default mode network at one extreme and somatomotor and visual areas at the other. By deriving this gradient using BOLD timeseries dimensionality reduction, we were able to quantify its dynamic causal interactions with other gradients. We found that the amount and direction of change in Gradient 2's slope depended on both its own current slope and Gradient 1's

slope, along with additional influence from the other gradients. An important aim for future research is to identify circuit-level mechanisms which drive dynamic gradient fluctuations and their mutual interactions. Subcortical neuromodulatory ascending projection systems for dopamine, noradrenaline, serotonin, and acetylcholine are all well positioned to exert widespread influence over brain activity (van den Brink et al., 2016; Shine et al., 2018; Shafiei et al., 2019; Grandjean et al., 2019; Munn et al., 2021). Gradient polarity may be mediated by reciprocal interactions between pairs of thalamic nuclei. When thalamocortical neurons in a specific nucleus are activated, they can drive activity in both their cortical targets and neurons in the thalamic reticular nucleus (TRN) via axon collaterals. Activation of these GABAergic TRN cells can in turn inhibit thalamocortical neurons in a paired nucleus, reducing activity in associated cortical target regions (Crabtree, 2018). This motif may enable thalamic coordination of widespread cortical functional connectivity (Hwang et al., 2017; Buckner and DiNicola, 2019).

Subsequent gradients also had strong correspondence with known functional connectivity gradients. Gradient 3 separated task-positive versus task-negative regions while Gradient 4 had visual and somatomotor regions at opposing poles. Previous studies have shown that similar gradients have variable expression by different individuals and correlate with specific ongoing thought patterns (Mckeown et al., 2020; Karapanagiotidis et al., 2020). Brain activity states with more engagement of the task-negative side of the task-negative/task-positive gradient correlated with more thoughts about the past (Karapanagiotidis et al., 2020), while the visual-sensorimotor gradient had greater engagement in individuals reporting more thoughts about problem-solving (Mckeown et al., 2020). When viewing gradients from the perspective of fluctuating dynamic systems, three different processes may play a role in generating specific activity and connectivity patterns. First, a transient gradient imbalance could preferentially activate regions at one pole of a gradient, producing more occurrences of the resultant brain activity state. Second, greater fluctuation amplitude of a given gradient could produce more segregation (anticorrelation) between regions at opposite ends. Third, specific coupling patterns between multiple gradients could achieve segregation or integration of multiple different networks on demand, as needed for the type of task being performed (Cohen and D'Esposito, 2016). Individual differences in gradient imbalance, amplitude, and coupling could be trait or state-related and are an interesting topic for future studies.

The first 12 activity gradients were spatially consistent across individuals and task-free or task-engaged states, suggesting that they represent intrinsic activity systems. This is consistent with a previous study showing consistent spatial topography of functional connectivity gradients across task states and levels of vigilance (Cross et al., 2021). Seeking additional support for the claim that gradients reflect intrinsic systems, we showed that plausible task activation maps resulted when gradients derived from the task-free based latent space formed specific dynamic combinations during task fMRI scans. This task-specific coordination of intrinsic activity gradients offers a new perspective on the consistency of the brain's functional architecture across rest and task (Smith et al., 2009).

4.2. Gradient coupling patterns can generate brain activity states

Using a dynamical systems model with specific gradient coupling parameters, we were able to replicate the spatial sequence of brain activity flow in actual task-free fMRI data. We then used this gradient-based model and found that we could accurately simulate different activity and functional connectivity states by using state-specific coupling parameters. The model captured individual functional connectivity differences in the task-free state and specific task activation and connectivity patterns supporting working memory, motor movements, language comprehension, and visual processing of emotional stimuli (Barch et al., 2013). Our findings build on previous work showing that distinct low-dimensional trajectories underlie task-specific brain activity states (Shine et al., 2019; Hawco et al., 2021). These results support the hypothesis that the shape of the low-dimensional activity manifold is determined by gradients interacting according to specific coupling parameters, which will result in different coactivation patterns (Liu et al., 2018a), state transition probabilities (Vidaurre et al., 2017), and sequences of activity flow (Cole et al., 2016; Yousefi and Keilholz, 2021; Raut et al., 2021). One noteworthy aspect of our modeling approach is that while the PCA-derived gradient timeseries are temporally orthogonal, this does not imply that gradients cannot interact with each other. Ordinary differential equations can accurately capture non-linear relationships between orthogonal variables in a low-dimensional system (Brunton et al., 2016). One implication of this is that even if the global signal or the Gradient 1 timeseries are regressed out of the data, Gradient 1′s influence on the other gradients' timeseries is likely to remain as a residual effect. Here our central modeling assumption was that gradients can exert distinct influence on one another. Future work can test this assumption with causal interventions and investigate how gradient coupling may be affected by individual trait and state differences or neurological conditions.

5. Conclusion

Here we derived a brain activity latent space by applying PCA to task-free BOLD fMRI timeseries data from 100 typical adults in the Human Connectome Project. We found that each dimension in this latent space represented a spatial activity gradient that dynamically fluctuates in time. The gradient maps corresponded to known functional connectivity gradients with the key exception of Gradient 1, a unipolar sensory-association gradient that explained the most variance in the BOLD signal, had a near-perfect correlation with the global signal, showed the strongest spatial correlation with the principal component of gene expression, and exerted the strongest influence on the dynamic activity of the other gradients. This set of activity gradients appeared to be intrinsic based on their spatial consistency in the same subjects during task-based fMRI scans, their reproducibility in different individuals in a validation dataset, and their ability to produce task-specific activity patterns. We demonstrated the equivalency of network-based and gradient-based representations of functional connectivity, showing that the non-uniform spacing of regions along multiple gradients reflects modularity and regional hub properties. We finally used a dynamical systems model to show that specific gradient coupling parameters were sufficient to generate state-specific global activity and functional connectivity patterns, both in different individuals in the task-free state and in working memory, motor, language, and emotion tasks.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

Data and code availability

Original data was obtained from the Human Connectome Project (1U54MH091657, PIs Van Essen and Ugurbil) and the Allen Human Brain Atlas (http://human.brain-map.org/). All code (latent space derivation, dynamical system modeling, and gene expression correlation) and processed data (gradient maps/region weights, gradient timeseries, and region gene expression values) are available at https://github.com/jbrown81/gradients.

References

Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD, 2014. Tracking whole-brain connectivity dynamics in the resting state. Cereb. Cortex 24, 663–676. doi:10.1093/cercor/bhs352. [PubMed: 23146964]

Andrews-Hanna JR, Reidler JS, Sepulcre J, Poulin R, Buckner RL, 2010. Functional-anatomic fractionation of the brain's default network. Neuron 65, 550–562. doi:10.1016/j.neuron.2010.02.005. [PubMed: 20188659]

Arnatkevicǐ bt A, Fulcher BD, Fornito A, 2019. A practical guide to linking brain-wide gene expression and neuroimaging data. NeuroImage 189, 353–367. doi:10.1016/j.neuroimage.2019.01.011. [PubMed: 30648605]

Barch DM, Burgess GC, Harms MP, Petersen SE, Schlaggar BL, Corbetta M, Glasser MF, Curtiss S, Dixit S, Feldt C, Nolan D, Bryant E, Hartley T, Footer O, Bjork JM, Poldrack R, Smith S, Johansen-Berg H, Snyder AZ, Van Essen DC, 2013. Function in the human connectome: Task-fMRi and individual differences in behavior. NeuroImage, Mapp. Connectome 80, 169–189. doi:10.1016/j.neuroimage.2013.05.033.

Berens P, 2009. CircStat: A MATLAB toolbox for circular statistics. J. Stat.l Softw 31, 1–21. doi:10.18637/jss.v031.i10.

Bertolero MA, Yeo BTT, D'Esposito M, 2015. The modular and integrative functional architecture of the human brain. PNAS 112, E6798–E6807. doi:10.1073/pnas.1510619112. [PubMed: 26598686]

Betzel RF, Gu S, Medaglia JD, Pasqualetti F, Bassett DS, 2016. Optimally controlling the human connectome: the role of network topology. Sci. Rep 6, 30770. doi:10.1038/srep30770. [PubMed: 27468904]

Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E, 2008. Fast unfolding of communities in large networks. J. Stat. Mech.: Theory Exp 2008, P10008. doi:10.1088/1742-5468/2008/10/P10008.

Brackenbury WJ, Yuan Y, O'Malley HA, Parent JM, Isom LL, 2013. Abnormal neuronal patterning occurs during early postnatal brain development of Scn1b-null mice and precedes hyperexcitability. Proc Natl Acad Sci USA 110, 1089–1094. doi:10.1073/pnas.1208767110. [PubMed: 23277545]

Breakspear M, 2017. Dynamic models of large-scale brain activity. Nature Neurosci. 20, 340–352. doi:10.1038/nn.4497. [PubMed: 28230845]

Brunton SL, Proctor JL, Kutz JN, 2016. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. PNAS 113, 3932–3937. doi:10.1073/pnas.1517384113. [PubMed: 27035946]

Buckner RL, DiNicola LM, 2019. The brain's default network: updated anatomy, physiology and evolving insights. Nat. Rev. Neurosci 20, 593–608. doi:10.1038/s41583-019-0212-7. [PubMed: 31492945]

Burt JB, Demirta M, Eckner WJ, Navejar NM, Ji JL, Martin WJ, Bernacchia A, Anticevic A, Murray JD, 2018. Hierarchy of transcriptomic specialization across human cortex captured by structural neuroimaging topography. Nature Neurosci. 21, 1251–1259. doi:10.1038/s41593-018-0195-0. [PubMed: 30082915]

Burt JB, Helmer M, Shinn M, Anticevic A, Murray JD, 2020. Generative modeling of brain maps with spatial autocorrelation. NeuroImage 220, 117038. doi:10.1016/j.neuroimage.2020.117038. [PubMed: 32585343]

Carbonell F, Bellec P, Shmuel A, 2011. Global and system-specific resting-state fMRI fluctuations are uncorrelated: principal component analysis reveals anti-correlated networks. Brain Connect. 1, 496–510. doi:10.1089/brain.2011.0065. [PubMed: 22444074]

Carcea I, Patil SB, Robison AJ, Mesias R, Huntsman MM, Froemke RC, Buxbaum JD, Huntley GW, Benson DL, 2014. Maturation of cortical circuits requires Semaphorin 7A. PNAS 111, 13978–13983. doi:10.1073/pnas.1408680111. [PubMed: 25201975]

Chang C, Glover GH, 2009. Relationship between respiration, end-tidal CO2, and BOLD signals in resting-state fMRI. NeuroImage 47, 1381–1393. doi:10.1016/j.neuroimage.2009.04.048. [PubMed: 19393322]

Cohen JR, D'Esposito M, 2016. The segregation and integration of distinct brain networks and their relationship to cognition. J. Neurosci 36, 12083–12094. doi:10.1523/JNEUROSCI.2965-15.2016. [PubMed: 27903719]

Cole MW, Bassett DS, Power JD, Braver TS, Petersen SE, 2014. Intrinsic and task-evoked network architectures of the human brain. Neuron 83, 238–251. doi:10.1016/j.neuron.2014.05.014. [PubMed: 24991964]

Cole MW, Ito T, Bassett DS, Schultz DH, 2016. Activity flow over resting-state networks shapes cognitive task activations. Nat. Neurosci 19, 1718–1726. doi:10.1038/nn.4406. [PubMed: 27723746]

Crabtree JW, 2018. Functional diversity of thalamic reticular subnetworks. Front. Syst. Neurosci 12. doi:10.3389/fnsys.2018.00041.

Cross N, Paquola C, Pomares FB, Perrault AA, Jegou A, Nguyen A, Aydin U, Bernhardt BC, Grova C, Dang-Vu TT, 2021. Cortical gradients of functional connectivity are robust to state-dependent changes following sleep deprivation. NeuroImage 226, 117547. doi:10.1016/j.neuroimage.2020.117547. [PubMed: 33186718]

Diedrichsen J, 2006. A spatially unbiased atlas template of the human cerebellum. NeuroImage 33, 127–138. doi:10.1016/j.neuroimage.2006.05.056. [PubMed: 16904911]

Faber M, Prze dzik I, Fernández G, Haak KV, Beckmann CF, 2020. Overlapping connectivity gradients in the anterior temporal lobe underlie semantic cognition. 10.1101/2020.05.28.121137

Fan L, Li H, Zhuo J, Zhang Y, Wang J, Chen L, Yang Z, Chu C, Xie S, Laird AR, Fox PT, Eickhoff SB, Yu C, Jiang T, 2016. The human brainnetome atlas: a new brain atlas based on connectional architecture. Cereb. Cortex Bhw 157. doi:10.1093/cercor/bhw157.

Faskowitz J, Esfahlani FZ, Jo Y, Sporns O, Betzel RF, 2020. Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. Nature Neurosci. 1–11. doi:10.1038/s41593-020-00719-y. [PubMed: 31844312]

Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME, 2005. The human brain is intrinsically organized into dynamic, anticor-related functional networks. Proc. Natl. Acad. Sci. USA 102, 9673–9678. doi:10.1073/pnas.0504136102. [PubMed: 15976020]

Fox MD, Zhang D, Snyder AZ, Raichle ME, 2009. The global signal and observed anticorrelated resting state brain networks. J. Neurophysiol 101, 3270–3283. doi:10.1152/jn.90777.2008. [PubMed: 19339462]

Fulcher BD, Murray JD, Zerbi V, Wang X-J, 2019. Multimodal gradients across mouse cortex. Proc. Natl. Acad. Sci. USA 116, 4689–4695. doi:10.1073/pnas.1814144116. [PubMed: 30782826]

Glasser MF, Coalson TS, Bijsterbosch JD, Harrison SJ, Harms MP, Anticevic A, Van Essen DC, Smith SM, 2019. Classification of temporal ICA components for separating global noise from fMRI data: Reply to Power. NeuroImage 197, 435–438. doi:10.1016/j.neuroimage.2019.04.046. [PubMed: 31026516]

Glasser MF, Coalson TS, Bijsterbosch JD, Harrison SJ, Harms MP, Anticevic A, Van Essen DC, Smith SM, 2018. Using temporal ICA to selectively remove global noise while preserving global signal in functional MRI data. NeuroImage 181, 692–717. doi:10.1016/j.neuroimage.2018.04.076. [PubMed: 29753843]

Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, Xu J, Jbabdi S, Webster M, Polimeni JR, Van Essen DC, Jenkinson M, 2013. The minimal preprocessing pipelines for the Human Connectome Project. NeuroImage Mapp. Connectome 80, 105–124. doi:10.1016/j.neuroimage.2013.04.127.

Glomb K, Kringelbach ML, Deco G, Hagmann P, Pearson J, Atasoy S, 2019. Functional harmonics reveal multi-dimensional basis functions underlying cortical organization. bioRxiv, 699678 doi:10.1101/699678.

Gordon EM, Lynch CJ, Gratton C, Laumann TO, Gilmore AW, Greene DJ, Ortega M, Nguyen AL, Schlaggar BL, Petersen SE, Dosenbach NUF, Nelson SM, 2018. Three distinct sets of connector hubs integrate human brain function. Cell Rep. 24, 1687–1695. doi:10.1016/j.celrep.2018.07.050, e4. [PubMed: 30110625]

Grandjean J, Corcoba A, Kahn MC, Upton AL, Deneris ES, Seifritz E, Helmchen F, Mann EO, Rudin M, Saab BJ, 2019. A brain-wide functional map of the serotonergic responses to acute stress and fluoxetine. Nat. Commun 10, 350. doi:10.1038/s41467-018-08256-w. [PubMed: 30664643]

Gratton C, Laumann TO, Nielsen AN, Greene DJ, Gordon EM, Gilmore AW, Nelson SM, Coalson RS, Snyder AZ, Schlaggar BL, Dosenbach NUF, Petersen SE, 2018. Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. Neuron 98, 439–452. doi:10.1016/j.neuron.2018.03.035, e5. [PubMed: 29673485]

Guell X, Schmahmann JD, Gabrieli JD, Ghosh SS, 2018. Functional gradients of the cerebellum. Elife 7, e36652. doi:10.7554/eLife.36652. [PubMed: 30106371]

Guimerà R, Nunes Amaral LA, 2005. Functional cartography of complex metabolic networks. Nature 433, 895–900. doi:10.1038/nature03288. [PubMed: 15729348]

Haak KV, Marquand AF, Beckmann CF, 2018. Connectopic mapping with resting-state fMRI. Neuroimage Segment. Brain 170, 83–94. doi:10.1016/j.neuroimage.2017.06.075.

Hawco C, Dickie EW, Jacobs G, Daskalakis ZJ, Voineskos AN, 2021. Moving beyond the mean: Subgroups and dimensions of brain activity and cognitive performance across domains. NeuroImage 231, 117823. doi:10.1016/j.neuroimage.2021.117823. [PubMed: 33549760]

Hawrylycz MJ, Lein ES, Guillozet-Bongaarts AL, Shen EH, Ng L, Miller JA, van de Lagemaat LN, Smith KA, Ebbert A, Riley ZL, Abajian C, Beckmann CF, Bernard A, Bertagnolli D, Boe AF, Cartagena PM, Chakravarty MM, Chapin M, Chong J, Dalley RA, David Daly B, Dang C, Datta S, Dee N, Dolbeare TA, Faber V, Feng D, Fowler DR, Goldy J, Gregor BW, Haradon Z, Haynor DR, Hohmann JG, Horvath S, Howard RE, Jeromin A, Jochim JM, Kinnunen M, Lau C, Lazarz ET, Lee C, Lemon TA, Li L, Li Y, Morris JA, Overly CC, Parker PD, Parry SE, Reding M, Royall JJ, Schulkin J, Sequeira PA, Slaughterbeck CR, Smith SC, Sodt AJ, Sunkin SM, Swanson BE, Vawter MP, Williams D, Wohnoutka P, Zielke HR, Geschwind DH, Hof PR, Smith SM, Koch C, Grant SGN, Jones AR, 2012. An anatomically comprehensive atlas of the adult human brain transcriptome. Nature 489, 391–399. doi:10.1038/nature11405. [PubMed: 22996553]

Hong S-J, Xu T, Nikolaidis A, Smallwood J, Margulies DS, Bernhardt B, Vogelstein J, Milham MP, 2020. Toward a connectivity gradient-based framework for reproducible biomarker discovery. NeuroImage 223, 117322. doi:10.1016/j.neuroimage.2020.117322. [PubMed: 32882388]

Huntenburg JM, Yeow LY, Mandino F, Grandjean J, 2021. Gradients of functional connectivity in the mouse cortex reflect neocortical evolution. NeuroImage 225, 117528. doi:10.1016/j.neuroimage.2020.117528. [PubMed: 33157264]

Hwang K, Bertolero MA, Liu WB, D'Esposito M, 2017. The human thalamus is an integrative hub for functional brain networks. J. Neurosci 37, 5594–5607. doi:10.1523/JNEUROSCI.0067-17.2017. [PubMed: 28450543]

Karapanagiotidis T, Vidaurre D, Quinn AJ, Vatansever D, Poerio GL, Turnbull A, Ho NSP, Leech R, Bernhardt BC, Jefferies E, Margulies DS, Nichols TE, Woolrich MW, Smallwood J, 2020. The psychological correlates of distinct neural states occurring during wakeful rest. Sci. Rep 10, 21121. doi:10.1038/s41598-020-77336-z. [PubMed: 33273566]

Kong X, Kong R, Orban C, Wang P, Zhang S, Anderson K, Holmes A, Murray JD, Deco G, van den Heuvel M, Yeo BTT, 2021. Sensory-motor cortices shape functional connectivity dynamics in the human brain. Nat. Commun 12, 6373. doi:10.1038/s41467-021-26704-y. [PubMed: 34737302]

Li J, Bolt T, Bzdok D, Nomi JS, Yeo BTT, Spreng RN, Uddin LQ, 2019. Topography and behavioral relevance of the global signal in the human brain. Sci. Rep 9, 1–10. doi:10.1038/s41598-019-50750-8. [PubMed: 30626917]

Liu TT, Nalci A, Falahpour M, 2017. The global signal in fMRI: Nuisance or Information? NeuroImage 150, 213–229. doi:10.1016/j.neuroimage.2017.02.036. [PubMed: 28213118]

Liu X, Zhang N, Chang C, Duyn JH, 2018a. Co-activation patterns in resting-state fMRI signals. NeuroImage Brain Connec. Dyn 180, 485–494. doi:10.1016/j.neuroimage.2018.01.041.

Liu X, Zwart J.A.de, Schölvinck ML, Chang C, Ye FQ, Leopold DA, Duyn JH, 2018b. Subcortical evidence for a contribution of arousal to fMRI studies of brain activity. Nature Communications 9, 395. doi:10.1038/s41467-017-02815-3.

Lynn CW, Cornblath EJ, Papadopoulos L, Bertolero MA, Bassett DS, 2021. Broken detailed balance and entropy production in the human brain. Proc. Natl. Acad. Sci. USA 118, e2109889118. doi:10.1073/pnas.2109889118. [PubMed: 34789565]

Margulies DS, Ghosh SS, Goulas A, Falkiewicz M, Huntenburg JM, Langs G, Bezgin G, Eickhoff SB, Castellanos FX, Petrides M, Jefferies E, Smallwood J, 2016. Situating the default-mode network along a principal gradient of macroscale cortical organization. PNAS 113, 12574–12579. doi:10.1073/pnas.1608282113. [PubMed: 27791099]

Mckeown B, Strawson WH, Wang H-T, Karapanagiotidis T, Vos de Wael R, Benkarim O, Turnbull A, Margulies D, Jefferies E, McCall C, Bernhardt B, Smallwood J, 2020. The relationship between individual variation in macroscale functional gradients and distinct aspects of ongoing thought. NeuroImage 220, 117072. doi:10.1016/j.neuroimage.2020.117072. [PubMed: 32585346]

Munn BR, Müller EJ, Wainstein G, Shine JM, 2021. The ascending arousal system shapes neural dynamics to mediate awareness of cognitive states. Nat. Commun 12, 6016. doi:10.1038/s41467-021-26268-x. [PubMed: 34650039]

Orban C, Kong R, Li J, Chee MWL, Yeo BTT, 2020. Time of day is associated with paradoxical reductions in global signal fluctuation and functional connectivity. PLOS Biol. 18, e3000602. doi:10.1371/journal.pbio.3000602. [PubMed: 32069275]

Paquola C, Wael RVD, Wagstyl K, Bethlehem RAI, Hong S-J, Seidlitz J, Bullmore ET, Evans AC, Misic B, Margulies DS, Smallwood J, Bernhardt BC, 2019. Microstructural and functional gradients are increasingly dissociated in transmodal cortices. PLOS Biol. 17, e3000284. doi:10.1371/journal.pbio.3000284. [PubMed: 31107870]

Pasquini L, Toller G, Staffaroni A, Brown JA, Deng J, Lee A, Kurcyus K, Shdo SM, Allen I, Sturm VE, Cobigo Y, Borghesani V, Battistella G, Gorno-Tempini ML, Rankin KP, Kramer J, Rosen HH, Miller BL, Seeley WW, 2020. State and trait characteristics of anterior insula time-varying functional connectivity. NeuroImage 208, 116425. doi:10.1016/j.neuroimage.2019.116425. [PubMed: 31805382]

Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL, Petersen SE, 2011. Functional network organization of the human brain. Neuron 72, 665–678. doi:10.1016/j.neuron.2011.09.006. [PubMed: 22099467]

Power JD, Laumann TO, Plitt M, Martin A, Petersen SE, 2017. On Global fMRI signals and simulations. Trends Cogn. Sci 21, 911–913. doi:10.1016/j.tics.2017.09.002. [PubMed: 28939332]

Power JD, Mitra A, Laumann TO, Snyder AZ, Schlaggar BL, Petersen SE, 2014. Methods to detect, characterize, and remove motion artifact in resting state fMRI. NeuroImage 84, 320–341. doi:10.1016/j.neuroimage.2013.08.048. [PubMed: 23994314]

Power JD, Schlaggar BL, Lessov-Schlaggar CN, Petersen SE, 2013. Evidence for Hubs in Human Functional Brain Networks. Neuron 79, 798–813. doi:10.1016/j.neuron.2013.07.035. [PubMed: 23972601]

Rand WM, 1971. Objective criteria for the evaluation of clustering methods. J. Am. Stat. Assoc 66, 846–850. doi:10.2307/2284239.

Raut RV, Snyder AZ, Mitra A, Yellin D, Fujii N, Malach R, Raichle ME, 2021. Global waves synchronize the brain's functional systems with fluctuating arousal. Sci. Adv 7. doi:10.1126/sciadv.abf2709.

Richiardi J, Altmann A, Milazzo A-C, Chang C, Chakravarty MM, Banaschewski T, Barker GJ, Bokde ALW, Bromberg U, Büchel C, Conrod P, Fauth-Bühler M, Flor H, Frouin V, Gallinat J, Garavan H, Gowland P, Heinz A, Lemaître H, Mann KF, Martinot J-L, Nees F, Paus T, Pausova Z, Rietschel M, Robbins TW, Smolka MN, Spanagel R, Ströhle A, Schumann G, Hawrylycz M, Poline J-B, Greicius MD, 2015. Correlated gene expression supports synchronous activity in brain networks. Science 348, 1241–1244. doi:10.1126/science.1255905. [PubMed: 26068849]

Rubinov M, Sporns O, 2010. Complex network measures of brain connectivity: uses and interpretations. NeuroImage 52, 1059–1069. doi:10.1016/j.neuroimage.2009.10.003. [PubMed: 19819337]

Saggar M, Sporns O, Gonzalez-Castillo J, Bandettini PA, Carlsson G, Glover G, Reiss AL, 2018. Towards a new approach to reveal dynamical organization of the brain using topological data analysis. Nature Commun. 9, 1399. doi:10.1038/s41467-018-03664-4. [PubMed: 29643350]

Satterthwaite TD, Elliott MA, Gerraty RT, Ruparel K, Loughead J, Calkins ME, Eickhoff SB, Hakonarson H, Gur RC, Gur RE, Wolf DH, 2013. An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. NeuroImage 64, 240–256. doi:10.1016/j.neuroimage.2012.08.052. [PubMed: 22926292]

Schaefer A, Kong R, Gordon EM, Laumann TO, Zuo X-N, Holmes AJ, Eickhoff SB, Yeo BTT, 2018. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. Cereb. Cortex 28, 3095–3114. doi:10.1093/cercor/bhx179. [PubMed: 28981612]

Seeley WW, Menon V, Schatzberg AF, Keller J, Glover GH, Kenna H, Reiss AL, Greicius MD, 2007. Dissociable intrinsic connectivity networks for salience processing and executive control. J. Neurosci 27, 2349–2356. doi:10.1523/JNEUROSCI.5587-06.2007. [PubMed: 17329432]

Seth AK, Barrett AB, Barnett L, 2015. Granger causality analysis in neuroscience and neuroimaging. J. Neurosci 35, 3293–3297. doi:10.1523/JNEUROSCI.4399-14.2015. [PubMed: 25716830]

Shafiei G, Markello RD, Vos de Wael R, Bernhardt BC, Fulcher BD, Misic B, 2020. Topographic gradients of intrinsic dynamics across neocortex. Elife 9. doi:10.7554/eLife.62116.

Shafiei G, Zeighami Y, Clark CA, Coull JT, Nagano-Saito A, Leyton M, Dagher A, Mišic B, 2019. Dopamine signaling modulates the stability and integration of intrinsic brain networks. Cereb. Cortex 29, 397–409. doi:10.1093/cercor/bhy264. [PubMed: 30357316]

Shine JM, Breakspear M, Bell PT, Martens KE, Shine R, Koyejo O, Sporns O, Poldrack RA, 2019. Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. Nature Neurosci. 1. doi:10.1038/s41593-018-0312-0.

Shine JM, van den Brink RL, Hernaus D, Nieuwenhuis S, Poldrack RA, 2018. Catecholaminergic manipulation alters dynamic network topology across cognitive states. Netw.. Neurosci 2, 381–396. doi:10.1162/netn_a_00042. [PubMed: 30294705]

Smith SM, Fox PT, Miller KL, Glahn DC, Fox PM, Mackay CE, Filippini N, Watkins KE, Toro R, Laird AR, Beckmann CF, 2009. Correspondence of the brain's functional architecture during activation and rest. Proc. Natl. Acad. Sci 106, 13040–13045. doi:10.1073/pnas.0905267106. [PubMed: 19620724]

Smith SM, Miller KL, Moeller S, Xu J, Auerbach EJ, Woolrich MW, Beckmann CF, Jenkinson M, Andersson J, Glasser MF, Van Essen DC, Feinberg DA, Yacoub ES, Ugurbil K, 2012. Temporally-Independent functional modes of spontaneous brain activity. PNAS 109, 3131–3136. doi:10.1073/pnas.1121329109. [PubMed: 22323591]

Sporns O, Betzel RF, 2016. Modular brain networks. Annu. Rev. Psychol 67, 613–640. doi:10.1146/annurev-psych-122414-033634. [PubMed: 26393868]

Sydnor VJ, Larsen B, Bassett DS, Alexander-Bloch A, Fair DA, Liston C, Mackey AP, Milham MP, Pines A, Roalf DR, Seidlitz J, Xu T, Raznahan A, Satterthwaite TD, 2021. Neurodevelopment of the association cortices: Patterns, mechanisms, and implications for psychopathology. Neuron 109, 2820–2846. doi:10.1016/j.neuron.2021.06.016. [PubMed: 34270921]

Tian Y, Margulies DS, Breakspear M, Zalesky A, 2020. Topographic organization of the human subcortex unveiled with functional connectivity gradients. Nat. Neurosci 23, 1421–1432. doi:10.1038/s41593-020-00711-6. [PubMed: 32989295]

Tsunemoto R, Lee S, Sz cs A, Chubukov P, Sokolova I, Blanchard JW, Eade KT, Bruggemann J, Wu C, Torkamani A, Sanna PP, Baldwin KK, 2018. Diverse reprogramming codes for neuronal identity. Nature 557, 375–380. doi:10.1038/s41586-018-0103-5. [PubMed: 29743677]

Turchi J, Chang C, Ye FQ, Russ BE, Yu DK, Cortes CR, Monosov IE, Duyn JH, Leopold DA, 2018. The basal forebrain regulates global resting-state fMRI fluctuations. Neuron 97, 940–952. doi:10.1016/j.neuron.2018.01.032, e4. [PubMed: 29398365]

Uddin LQ, 2020. Bring the noise: reconceptualizing spontaneous neural activity. Trends Cogn. Sci 24, 734–746. doi:10.1016/j.tics.2020.06.003. [PubMed: 32600967]

van den Brink RL, Pfeffer T, Warren CM, Murphy PR, Tona K-D, van der Wee NJA, Giltay E, van Noorden MS, Rombouts SARB, Donner TH, Nieuwenhuis S, 2016. Catecholaminergic neuromodulation shapes intrinsic mri functional connectivity in the human brain. J. Neurosci 36, 7865–7876. doi:10.1523/JNEUROSCI.0744-16.2016. [PubMed: 27466332]

Vidaurre D, Smith SM, Woolrich MW, 2017. Brain network dynamics are hierarchically organized in time. PNAS 114, 12827–12832. doi:10.1073/pnas.1705120114. [PubMed: 29087305]

Vogel JW, La Joie R, Grothe MJ, Diaz-Papkovich A, Doyle A, Vachon-Presseau E, Lepage C, Vos de Wael R, Thomas RA, Iturria-Medina Y, Bernhardt B, Rabinovici GD, Evans AC, 2020. A molecular gradient along the longitudinal axis of the human hippocampus informs large-scale behavioral systems. Nat. Commun 11, 960. doi:10.1038/s41467-020-14518-3. [PubMed: 32075960]

Vos de Wael R, Benkarim O, Paquola C, Lariviere S, Royer J, Tavakol S, Xu T, Hong S-J, Langs G, Valk S, Misic B, Milham M, Margulies D, Smallwood J, Bernhardt BC, 2020. BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. Commun. Biol 3, 103. doi:10.1038/s42003-020-0794-7. [PubMed: 32139786]

Wang X-J, 2020. Macroscopic gradients of synaptic excitation and inhibition in the neocortex. Nat. Rev. Neurosci 21, 169–178. doi:10.1038/s41583-020-0262-x. [PubMed: 32029928]

Yang GJ, Murray JD, Glasser M, Pearlson GD, Krystal JH, Schleifer C, Repovs G, Anticevic A, 2017. Altered global signal topography in schizophrenia. Cerebral Cortex 27, 5156–5169. doi:10.1093/cercor/bhw297. [PubMed: 27702810]

Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR, Fischl B, Liu H, Buckner RL, 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J. Neurophysiol 106, 1125–1165. doi:10.1152/jn.00338.2011. [PubMed: 21653723]

Yousefi B, Keilholz S, 2021. Propagating patterns of intrinsic activity along macroscale gradients coordinate functional connections across the whole brain. NeuroImage 231, 117827. doi:10.1016/j.neuroimage.2021.117827. [PubMed: 33549755]

Zhang J, Abiose O, Katsumi Y, Touroutoglou A, Dickerson BC, Barrett LF, 2019. Intrinsic functional connectivity is organized as three interdependent gradients. Sci. Rep 9, 1–14. doi:10.1038/s41598-019-51793-7. [PubMed: 30626917]
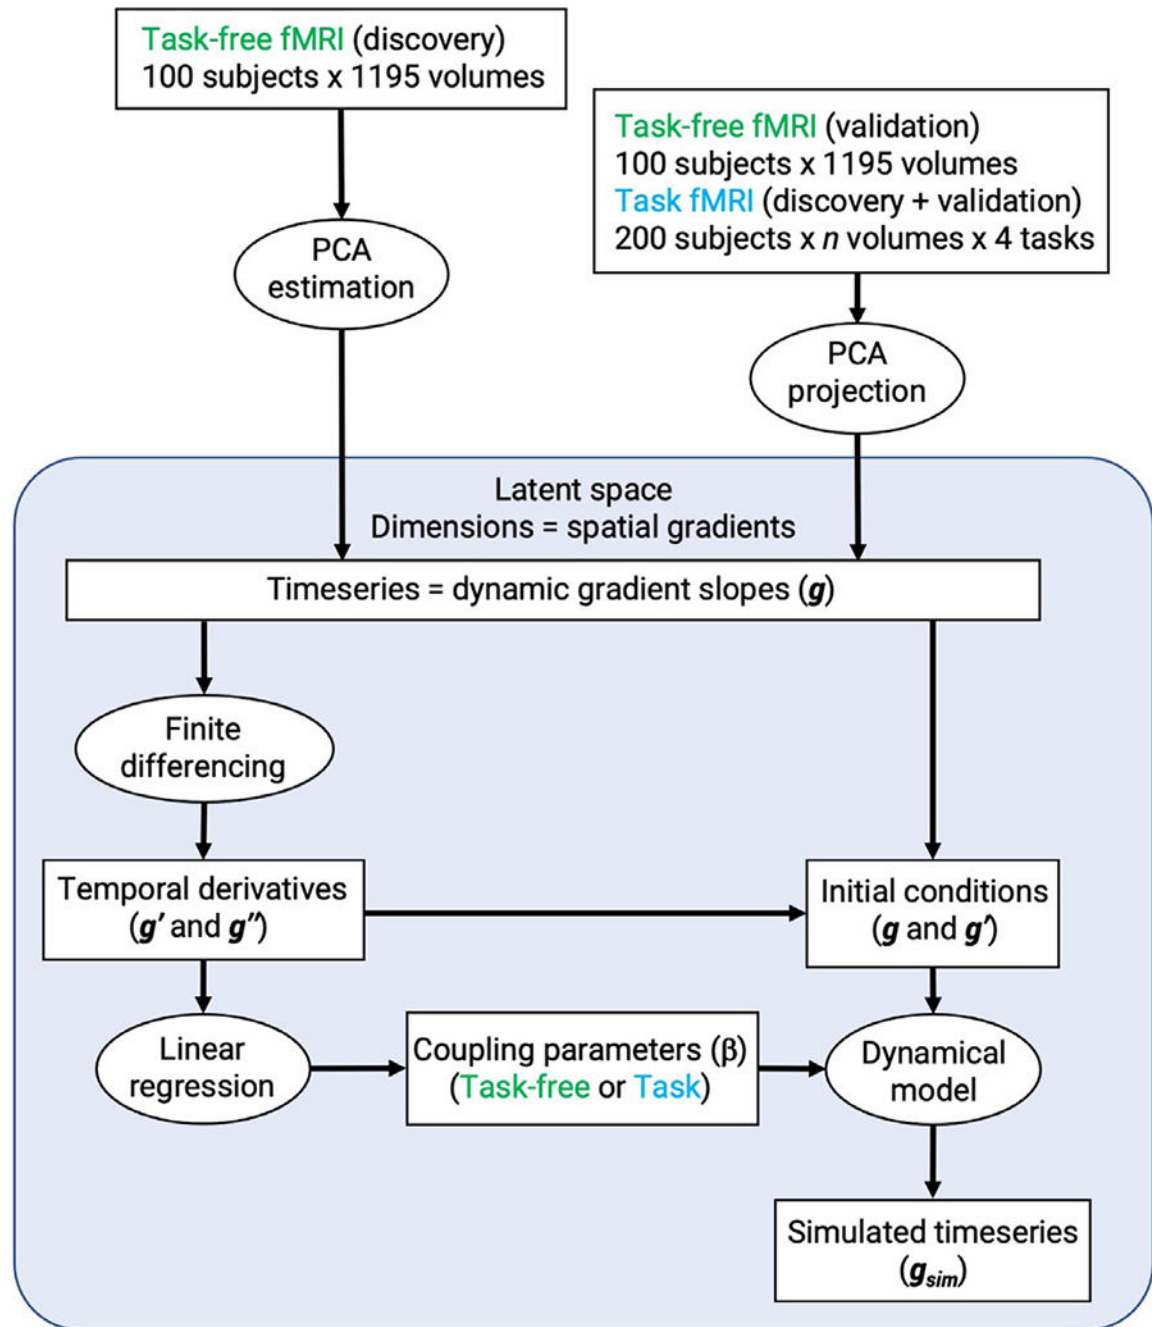
Fig. 1. Study workflow.

The primary workflow of the study. Rectangular boxes show data and ellipses show functional operations applied to data.
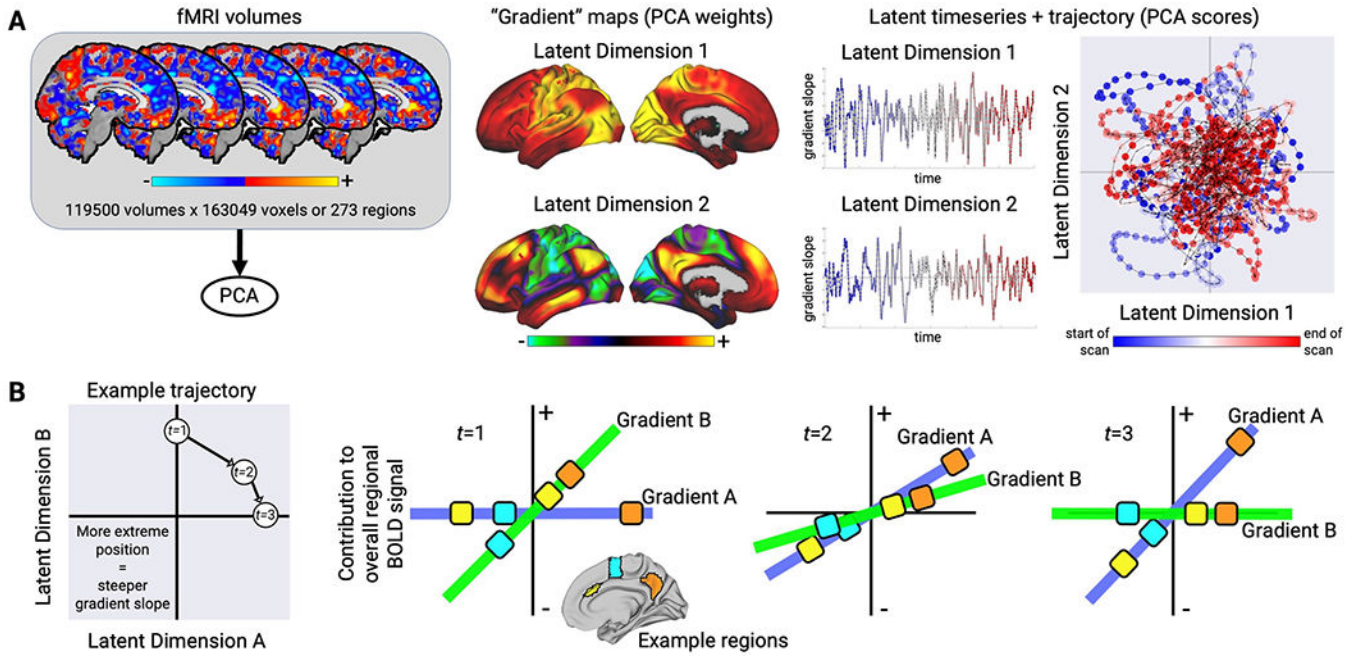
Fig. 2.

Latent space derivation and spatiotemporal interpretation. **A**. The procedure for deriving the brain activity latent space with PCA on task-free fMRI timeseries data, yielding spatial BOLD activity gradient maps, gradient activity timeseries, and resultant individual latent trajectories. **B.** Illustration of the relationship between latent space trajectories and regional BOLD activity. Across three successive timepoints, latent space positions on each dimension reflect the current slopes of the corresponding gradients. The resultant BOLD signal in each region depends on the region's weight on each gradient, shown for the anterior cingulate cortex (yellow), premotor cortex (cyan), and precuneus (orange).
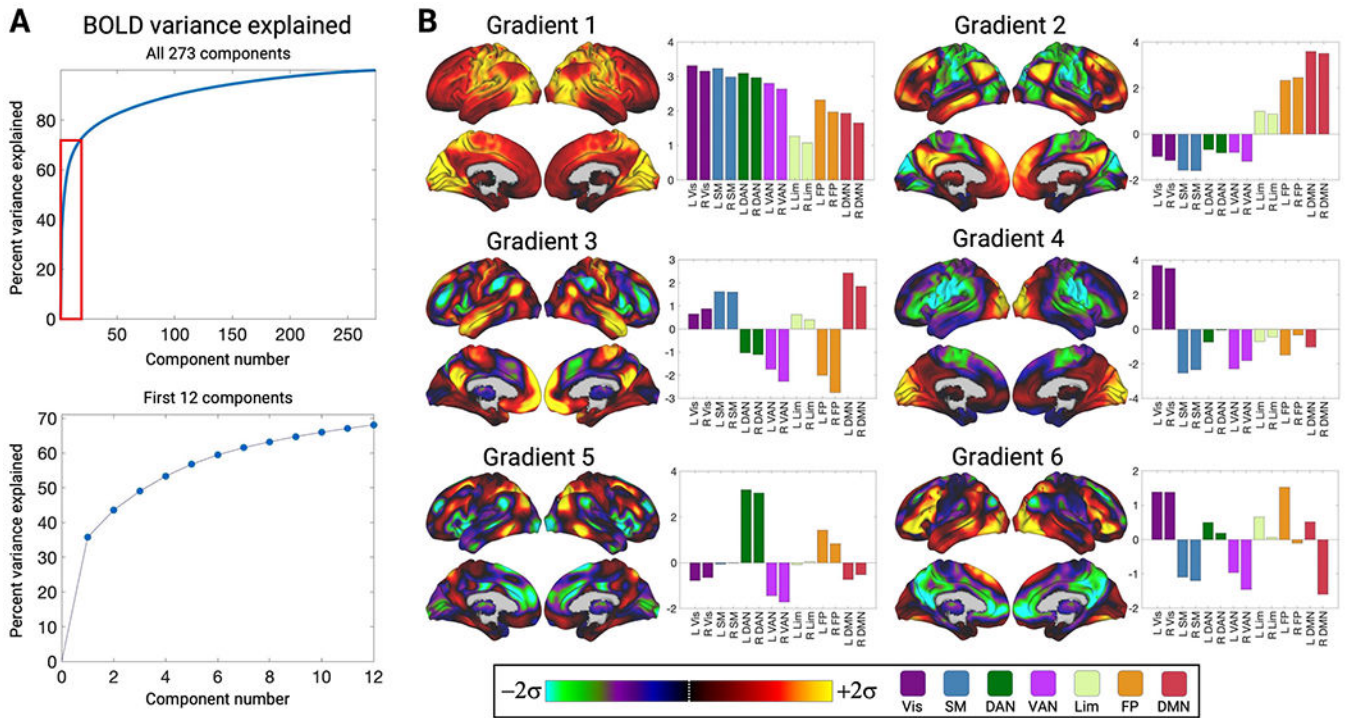
Fig. 3. Latent space variance explained and activity gradient maps.
**A.** The amount of cumulative variance explained by the PCA components derived from the 119500×273 regional timeseries. The components within the red box are expanded in the lower plot. **B.** Gradient maps for the first six latent space dimensions. Colormaps show gradient weights from −2 to + 2 standard deviations (σ), scaled separately for each gradient. Bar plots show the mean gradient weight for voxels overlapping with each of seven canonical cortical networks in the left or right hemisphere (Yeo et al., 2011). Units are PCA loadings multiplied by 1000. Vis: Visual; SM: Somatomotor; DAN: Dorsal Attention Network; VAN: Ventral Attention Network; Lim: Limbic; FP: Frontoparietal; DMN: Default Mode Network.
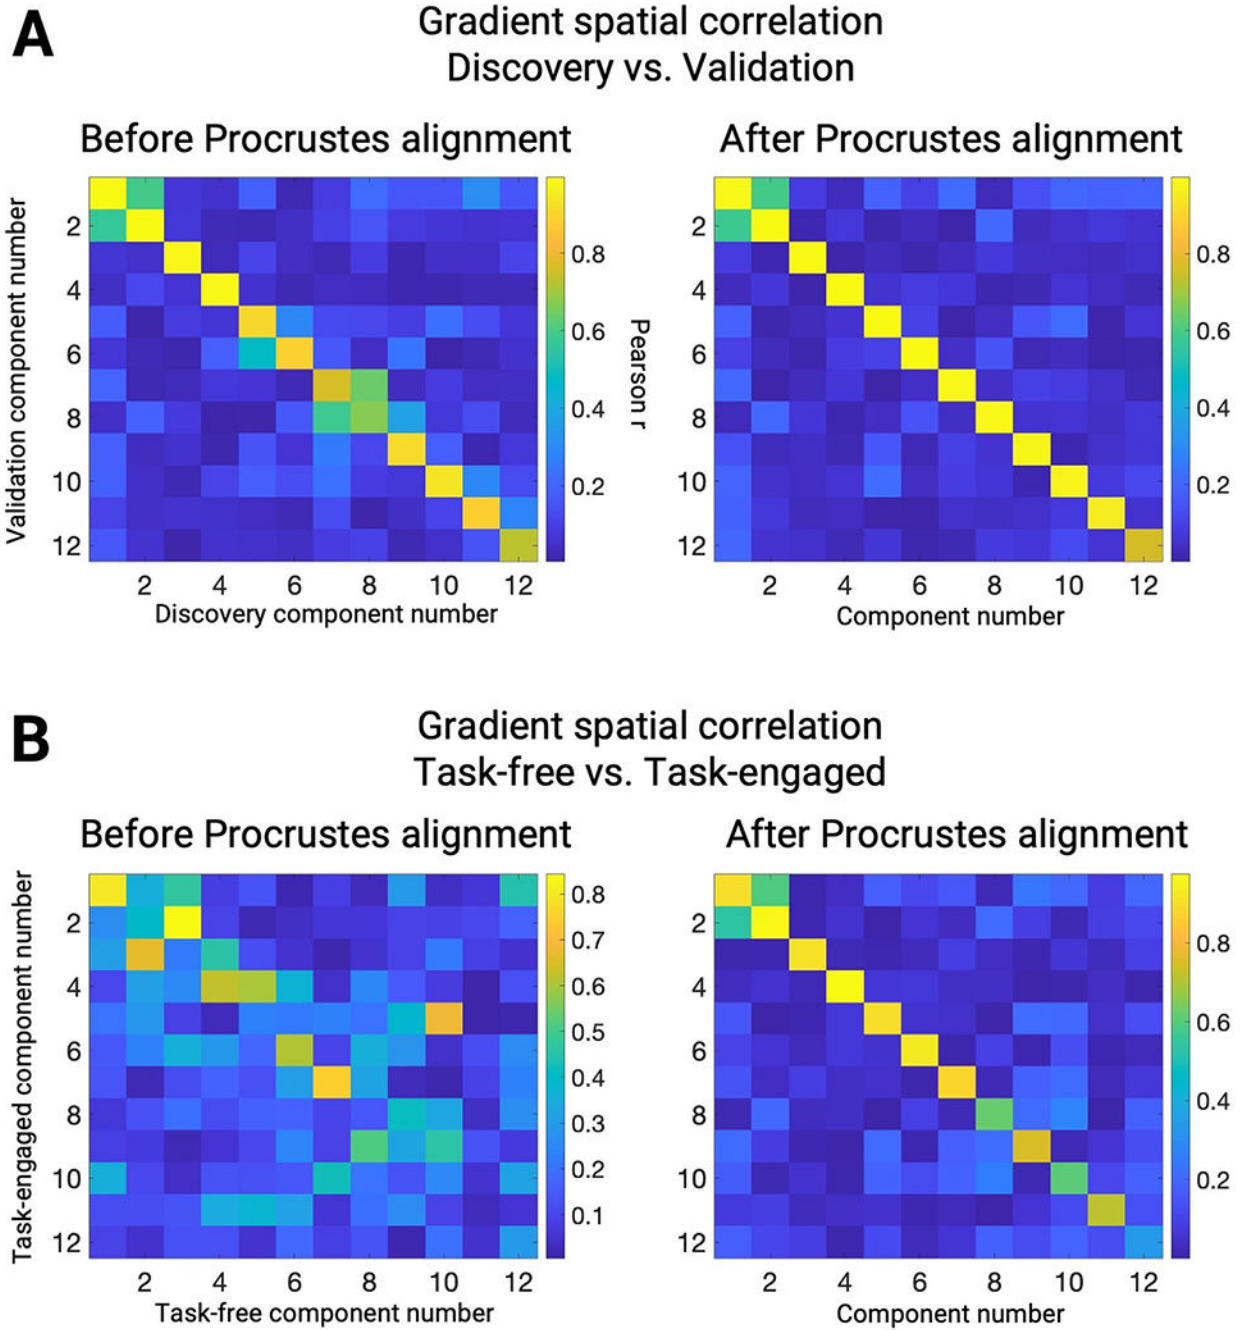
**Fig. 4. Gradient map reproducibility across individuals and cognitive states.**
A. The spatial correlation between gradient maps derived from the discovery vs. validation task-free datasets, before and after Procrustes alignment. Absolute values are shown. **B.** The spatial correlation between gradient maps derived from the task-free vs. task datasets, before and after Procrustes alignment.
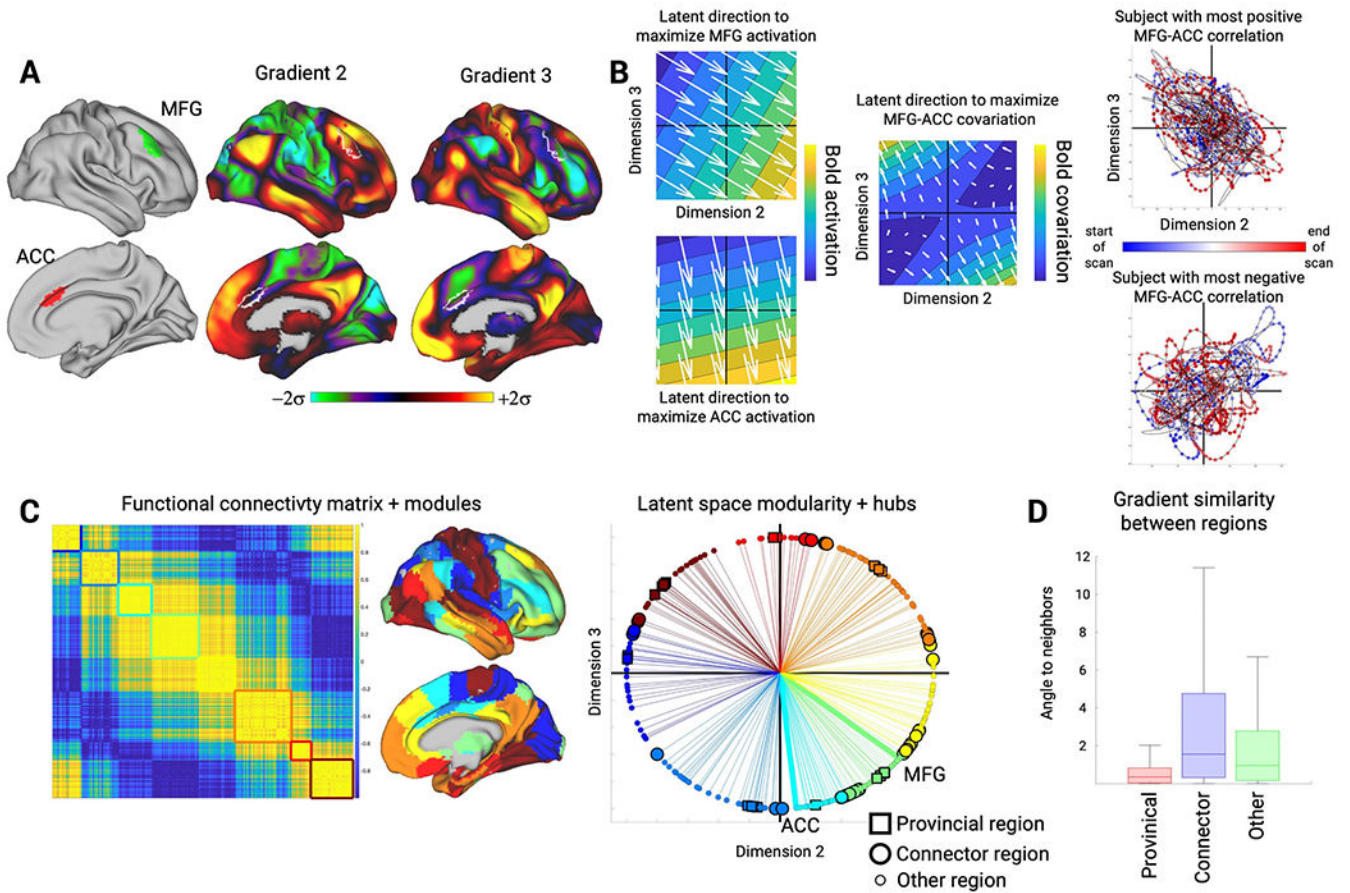
Fig. 5. The relationship between latent space trajectories, regional activity, and functional modularity.
A. Gradient weights of the middle frontal gyrus (MFG) and anterior cingulate cortex (ACC) on Gradients 2 and 3, the first two bipolar gradients. **B.** Latent trajectory directions that maximize BOLD activity for the MFG, ACC, or the functional connectivity (BOLD covariation) between them. Also shown are latent trajectories during task-free scans for the subjects with the most positive (top) or negative (bottom) ACC-MFG functional connectivity. **C. Left:** Group-average functional connectivity matrix based on Gradients 2 and 3. The eight modules detected in this connectivity matrix are highlighted along the diagonal and colored on the brain surface. **Middle:** Latent space plot showing the optimal trajectory direction for maximizing activity in each region, including the MFG and ACC from panel B. Color represents module membership. Connector regions (large circles) and provincial regions (large squares) are shown. **D.** Box plots for provincial regions, connector regions, and non-provincial/connector regions, showing angle to functionally connected neighbors. Plots show the median, interquartile range, and minimum/maximum values.
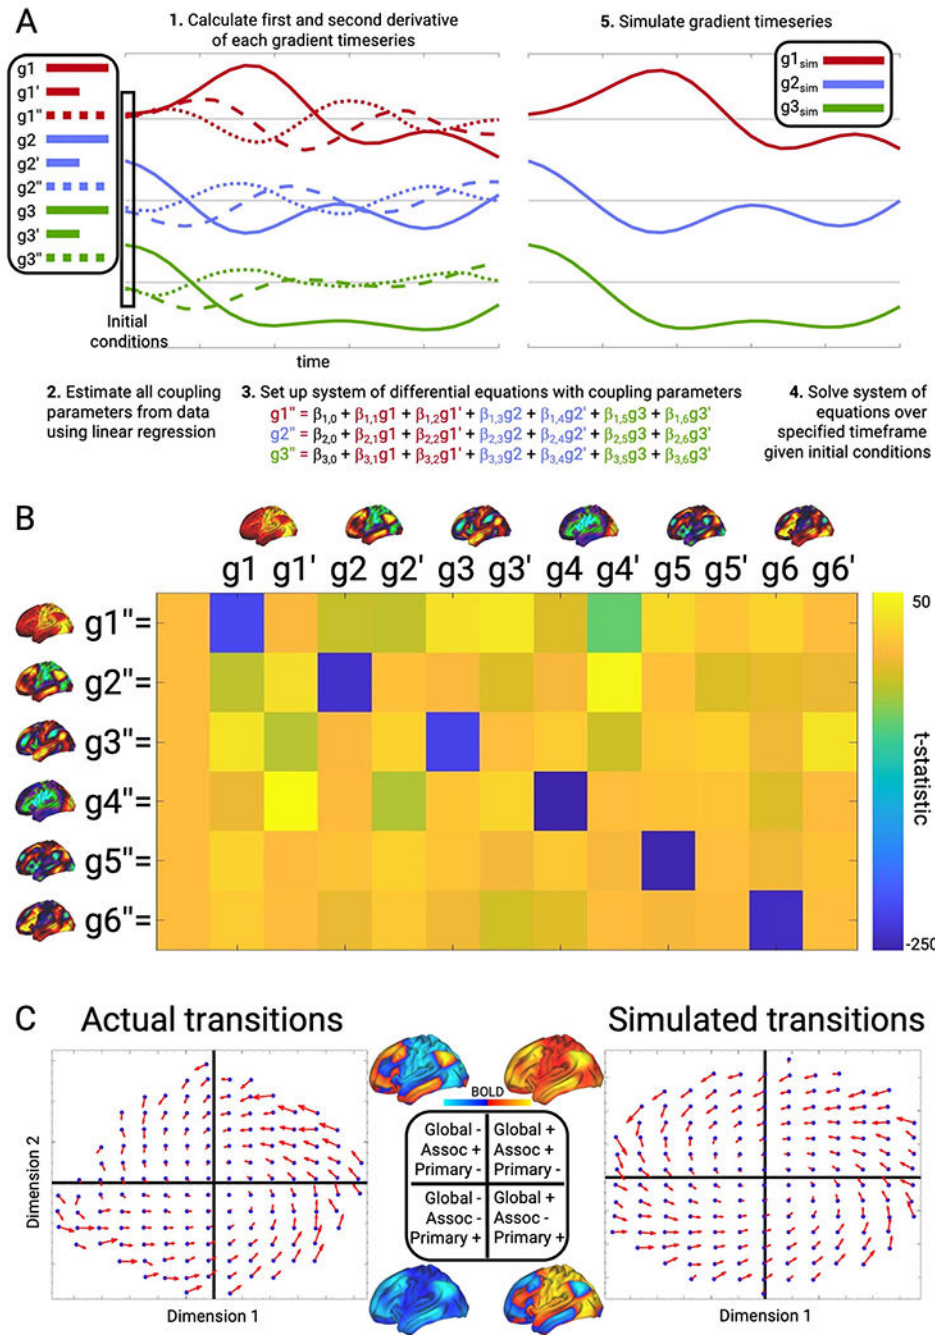
Fig. 6. Dynamical systems model overview, gradient coupling parameters, and simulated trajectories.
A. The dynamical systems modeling procedure. **B.** Task-free coupling parameters for the gradients 1-6. T-statistics for the parameter estimates are shown. **C.** Average latent space trajectory transitions from time *t* to *t*+1 for dimensions 1 and 2. The left plot shows average transitions in the actual task-free data and the right plot shows the average simulated transitions generated by the dynamical systems model. Each latent space quadrant reflects
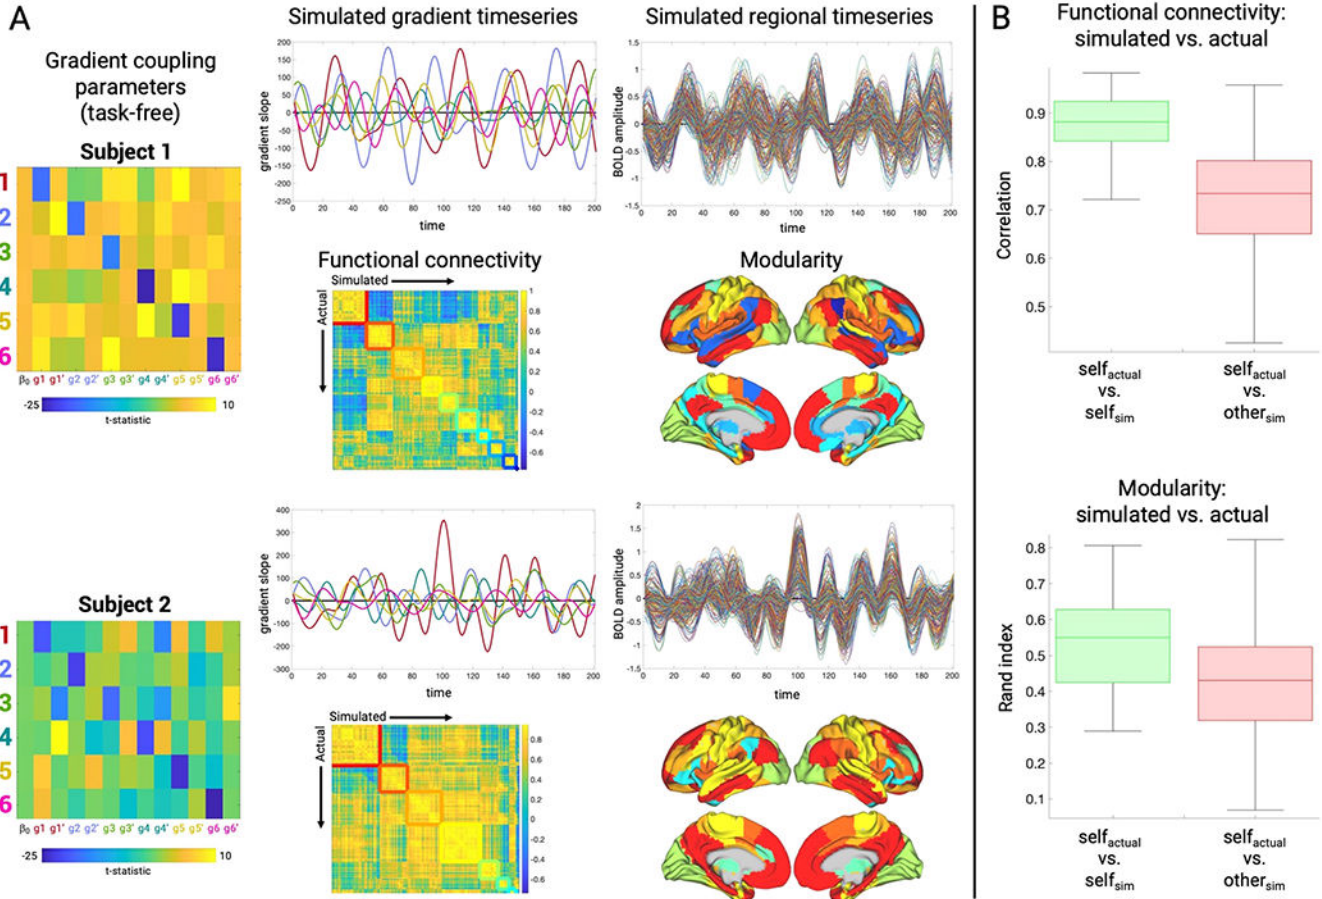
either a positive or negative slope for Gradient 1 ("Global") and Gradient 2 ("Primary" vs. "Assoc").

Fig. 7. Individual-specific simulated task-free activity and connectivity patterns.
A. Simulations for two example subjects showing their task-free gradient coupling parameters, the simulated gradient timeseries, the resulting regional timeseries, and the resulting functional connectivity matrix on the upper triangle. The actual functional connectivity matrix is shown on the lower triangle. The resulting differences in modularity for the two subjects are shown, including the default mode network (red). **B. Top:** Box plots showing simulated and actual functional connectivity matrix similarity based on using either the subject's own task-free coupling parameters (self vs. self, green) or the 99 other subject's coupling parameters (self vs. other, red). Plots show the median, interquartile range, and minimum/maximum values. **Bottom:** the modularity similarity from the simulated and actual functional connectivity matrices.
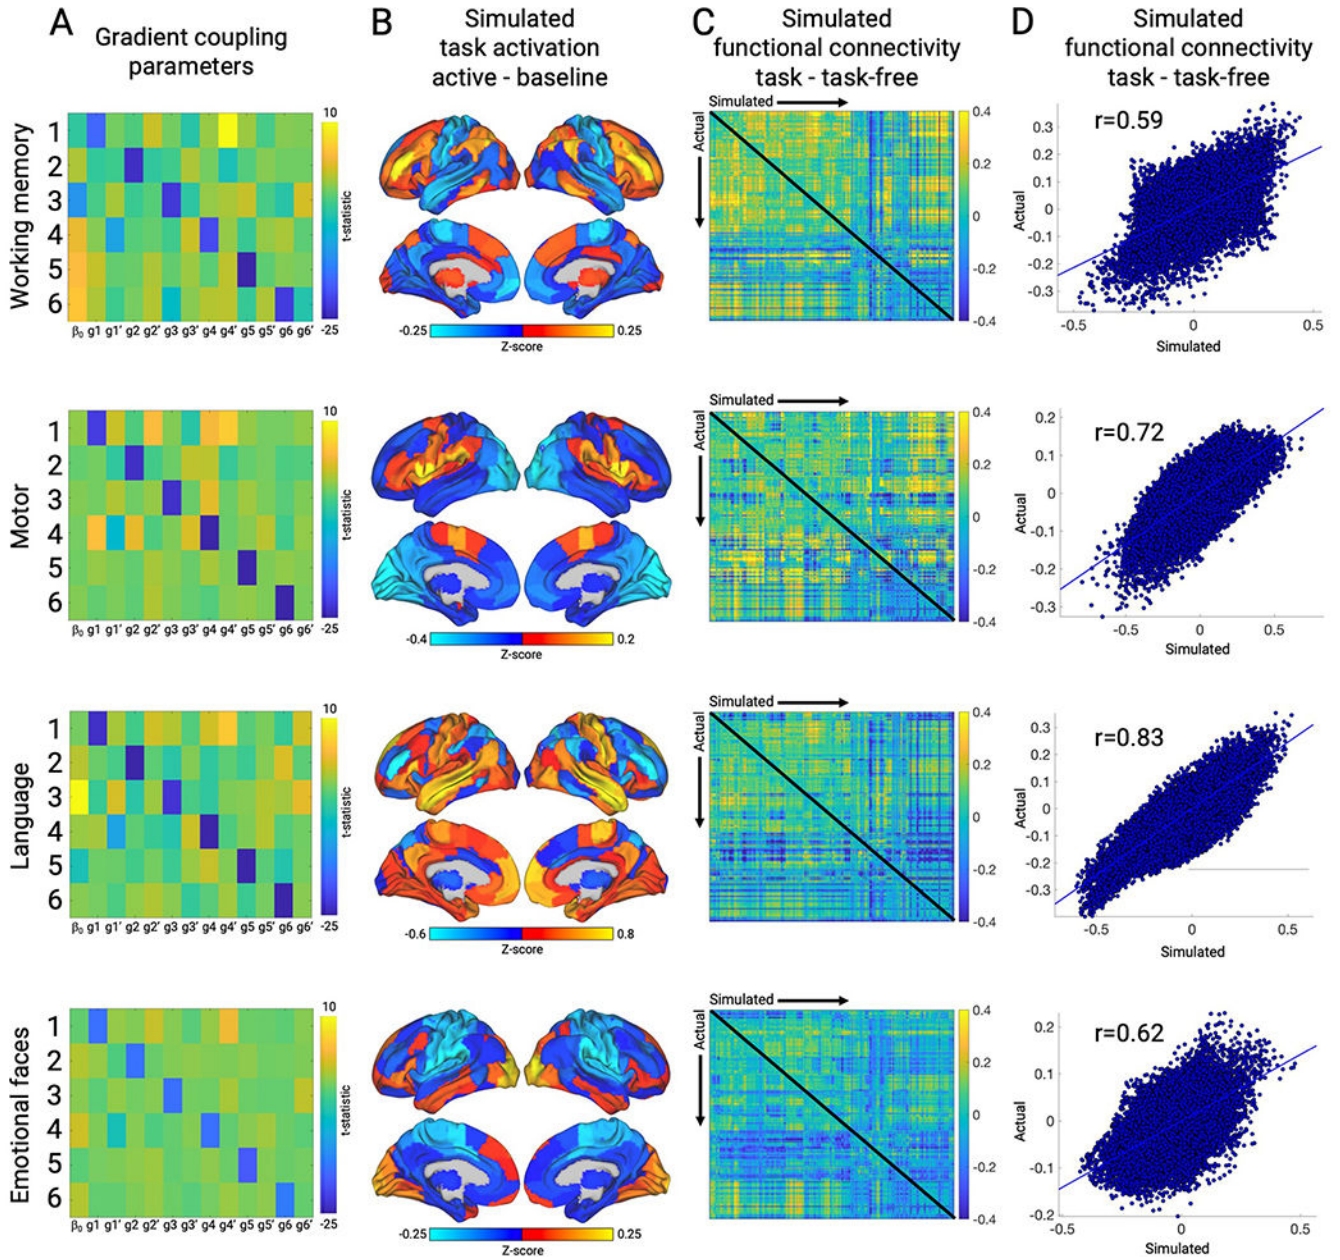
Fig. 8. Task-specific simulated activity and connectivity patterns.
A. Gradient coupling parameters for each task. **B.** Simulated task activation differences for the active task condition versus the baseline task condition. Color represents the difference in units of normalized BOLD standard deviation (Z-score). **C.** Functional connectivity (FC) matrices resulting from task-specific simulations. Simulated FC for the active task condition minus the simulated task-free FC is shown on the upper triangle of the matrix and actual active task condition FC minus the actual task-free FC (from the validation dataset) is shown on the lower triangle. **D.** Partial correlations between simulated and actual FC. FC values represent the region pairwise correlation strengths, adjusted for the actual FC strengths (task - task-free) in each of the other three tasks.

video 1.