

Network geometry shapes the transformation of task representations across the human cortex

Received: 14 July 2025

Accepted: 10 September 2026

Published online: 22 October 2026

 Check for updates

Lakshman N. C. Chakravarthula¹✉, Takuya Ito², Alexandros Tzavallas¹ & Michael W. Cole¹

Flexible cognition requires the brain to generalize across tasks while preserving task-specific information. Previous work showed that task representations are compressed in the association cortex and expanded in sensory and motor systems, but the mechanisms underlying these transformations are unclear. Here we test whether the geometry of intrinsic brain connectivity constrains how representations change between cortical regions. We analyzed functional MRI acquired during 16 diverse cognitive tasks together with activity flow modeling, finding that connectivity dimensionality predicts representational changes across the cortex. Low-dimensional connectivity compressed task representations and increased similarities across tasks, whereas high-dimensional connectivity expanded representations and supported conjunctive coding of task features. Activity flow modeling reproduced the observed compression-to-expansion pattern along the sensory–association–motor hierarchy and generated representational geometries that more closely matched targets than sources, consistent with transformation rather than transfer of information. These findings identify network geometry as a systems-level principle that shapes cortical representations and supports flexible cognition.

The human brain's ability to perform diverse tasks efficiently is a hallmark of cognitive flexibility. A key question in the neuroscience of cognitive flexibility is understanding how distributed neural circuits are recruited and reconfigured to navigate diverse task states¹. Recent studies have shed light on how different brain regions flexibly represent information during the performance of diverse tasks, revealing structured patterns of neural activity that adaptively support diverse cognitive demands^{2–6}. Along the sensory-to-motor cortical axis, task representations are first compressed and then expanded⁶. Concretely, different task conditions may produce distinct high-dimensional patterns in the visual cortex reflecting their unique sensory features (colors, shapes and words). In the association cortex, representations

compress—conditions that share abstract task-relevant features (for example, 'respond left' or 'inhibit response') become more similar despite differing sensory details. In the motor cortex, dimensionality increases again, as abstract features must be translated into specific motor implementations: a 'respond left' command could mean left index finger, left middle finger or other motor specifics, each requiring distinct neural patterns. This suggests that sensory information is initially compressed across task conditions to facilitate transferable representations, before expanding into representations that specify task-specific actions. This compression-then-expansion pattern is thought to provide the representational prerequisites for cognitive flexibility—the capacity to perform diverse tasks while maintaining

¹Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ, USA. ²T.J. Watson Research Center, IBM Research, Yorktown, NY, USA.
✉e-mail: lakastro@gmail.com

task-specific performance quality⁷. Specifically, flexible performance requires representations that balance two competing demands: generalizability (sharing learned components across tasks to enable transfer) and separability (maintaining distinct implementations to prevent interference). This balance—compression enabling generalization and expansion enabling separability—is what we refer to as the representational prerequisites for cognitive flexibility. Compression in some regions may enable generalization by emphasizing shared abstract features across tasks, whereas expansion in other regions supports separability by maintaining task-specific sensory details and motor implementations. However, the systems-level mechanisms that generate this functional architecture remain unclear.

Previous studies have shown that different regions prioritize encoding different aspects of task conditions: some studies have observed regions representing both task-relevant and irrelevant variables⁸, whereas others observed coding schemes supporting selective collapse of task-irrelevant information^{4,9,10} or maintaining abstract representations^{11–13}. Other studies have pointed toward representations encoding conjunctive combinations of variables in a nonlinear mixed-selective manner^{14–16}. Consequently, some brain regions exhibit low-dimensional representations, whereas others represent task conditions in high-dimensional spaces¹⁷. These variations in representational geometry suggest differing computational functions across brain regions^{7,17}.

A central challenge in systems neuroscience is understanding how the brain's network architecture enables flexible cognition. Although local computations within brain regions certainly contribute to representational geometry, the role of inter-regional connectivity architecture remains underexplored. Recent theoretical work predicts that connectivity structure should systematically constrain neural representations^{18–20}, yet empirical validation in biological circuits, particularly during diverse cognitive tasks, has been limited. Previous studies have demonstrated that task-evoked activity can be predicted using network connectivity^{21,22}, and that distributed connectivity patterns can even predict localized functional specialization like category selectivity²³. Building on these findings, we test whether connectivity architecture provides a systematic organizing principle for representational transformations across the cortical hierarchy. Specifically, does the geometric structure of inter-regional connectivity—whether connections are low or high dimensional—systematically shape how task representations compress or expand between brain regions?

To test this hypothesis, we build upon the activity flow modeling framework, which uses empirical functional connections to construct generative models to predict task-evoked activity patterns^{24–26}. In particular, this modeling framework proposes that task-evoked activity in target brain regions can be explained by activity ‘flowing’ (propagating) from source regions to the target region via source–target connectivity. Here we extend this framework to predict the representational geometry of brain regions using diverse task sets.

Specifically, we proposed that the dimensionality of intrinsic (resting-state) connectivity patterns systematically shapes the dimensionality of task representations in target regions (Fig. 1a). These are fundamentally different measures: connectivity dimensionality reflects the geometric structure of how regions connect during a single brain state (rest), whereas representational dimensionality reflects how a region's activity varies across diverse task conditions. Testing whether these independently measured properties are systematically related would reveal whether network architecture constrains representational geometry. Critically, this prediction is non-trivial: low-dimensional connectivity could theoretically lead to representational expansion in target regions if target regions are dominated by local, nonlinear processes, and high-dimensional connectivity could reflect random connection patterns bearing no systematic relationship to representations. Moreover, it is often assumed that inter-region connectivity plays a passive role (like inter-computer connectivity in the internet),

wherein connections simply carry information, and any information transformation occurs within (rather than between) regions. Finding systematic relationships between connectivity geometry and representational transformations would demonstrate that network architecture provides structured constraints on neural computation.

Here we examine how connectivity dimensionality shapes inter-areal representational transformations during diverse task performance. Using functional MRI (fMRI) data from participants performing 16 different tasks (Fig. 1b), we first demonstrate a robust compression-then-expansion pattern in representational dimensionality across the cortical hierarchy and validate its generalizability in an independent Human Connectome Project (HCP) dataset ($n = 352$). We find that the dimensionality of intrinsic connectivity systematically predicts how task representations are transformed between brain regions: low-dimensional connectivity leads to compression of representations, whereas high-dimensional connectivity leads to expansion. Through activity flow modeling, we show that these connectivity patterns actively generate observed representational geometries across the cortex. Moreover, we demonstrate through simulation that strong nonlinear local computations can fully decouple connectivity dimensionality from representational transformations. Rather than implementing transformations through complex local computations in each region, the brain leverages network architecture itself as a computational substrate. Regions with lower-dimensional connectivity exhibit stronger cross-task similarities, suggesting shared latent features that enable generalization across tasks, whereas expansions from high-dimensional connectivity support conjunctive coding of task variables. Higher representational dimensionality, particularly in sensory and association regions, predicts individual differences in general fluid intelligence. Together, these results reveal how the network geometry of the brain's intrinsic connectivity architecture shapes the transformation of task representations to support flexible cognition.

Results

Representational compression (expansion) is associated with low- (high-) dimensional FC

Our main hypothesis was that intrinsic connectivity shapes cross-region representational transformations across diverse tasks. This hypothesis is based on the activity flow modeling framework, which proposes that a region's task-evoked activity is shaped by activity flowing from connected source regions, weighted by their connectivity strengths²⁵. Consistent with this, our hypothesis predicted that there should be a relationship between cross-region representational dimensionality (compression or expansion, where compression indicates lower-dimensional representations in the target relative to source regions, and expansion indicates higher-dimensional representations) and connectivity dimensionality. We expected to observe that low connectivity dimensionality between regions compressed task representations, and that high connectivity dimensionality expanded task representations.

We examined representational transformations between connected regions using a two-step approach to functional connectivity (FC) estimation (Methods). In brief, we estimated a subject-specific FC graph from source vertices (all vertices of non-target regions connected to the target region) using graphical lasso (to identify connected region pairs) and principal component regression (to estimate vertex-level connectivity weights within those pairs). This whole-cortex vertex-level FC formed the basis for our subsequent analyses of connectivity dimensionality and its relationship to representational compression.

We measured representational dimensionality in each region by computing the participation ratio of the eigenspectrum of the cross-validated task representational similarity matrix (RSM). The participation ratio quantifies how many dimensions are needed to capture the variance in neural representations. Higher values indicate that task conditions produce distinct activity patterns across




Fig. 1 | Testing how connectivity geometry shapes representational transformations across the cortex. **a**, Hypothesis: connectivity dimensionality determines compression versus expansion of task representations. Task conditions (colored spheres labeled 1–4) occupy positions in neural activity space. Representational dimensionality determines which pairwise task groupings can be separated by linear decision boundaries (transparent planes). Higher dimensionality enables decoding of more diverse task groupings: in 3D space, all three pairwise splits are linearly separable (1–2 versus 3–4, 2–3 versus 1–4 and 2–4 versus 1–3); in 2D space, two pairwise splits remain possible (1–2 versus 3–4 and 2–3 versus 1–4); in 1D space, only one pairwise split is possible (1–2 versus 3–4). Top, high-dimensional connectivity (dissimilar connectivity patterns among target vertices, shown in the connectivity matrix) preserves representational dimensionality when source activity flows to target (3D → 2D example), maintaining decodability of multiple task groupings. Bottom, low-dimensional connectivity (similar patterns among target vertices) compresses representations (3D → 1D example), eliminating the capacity to decode most task groupings. **b**, Multi-task battery provides diverse representational structure. Left, list of 16 active visual tasks spanning cognitive control, memory, language and motor domains. Right, example task structure using the Stroop task, which includes eight task conditions crossing word–color congruency (congruent: word matches color; incongruent: word differs from color) with four ink colors. Bottom, each task consisted of a 5 s instruction period followed by 30 s of task performance. Across all 16 tasks, we estimated activity patterns for 96 distinct task conditions using a regularized General Linear Model (GLM),

providing rich representational diversity for measuring dimensionality. IAPS, International Affective Picture System. **c**, Activity flow modeling framework. Task activity in source regions (all functionally connected regions identified via sparse graphical lasso; Methods) is projected to each target region via intrinsic (resting-state) functional connectivity weights estimated at the vertex level using principal component regression. For each task condition, predicted target activity is computed as the weighted sum of source activities: predicted target activity = $\sum_{\text{sources}} (\text{source activity} \times \text{FC})$. From predicted target activity patterns across all 96 task conditions, we derive the predicted representational similarity structure (RSM) and predicted representational dimensionality using participation ratio of the eigenspectrum (Methods). **d**, Testing approach. We test two key predictions. Left, does connectivity dimensionality (measured from resting-state FC using participation ratio of the singular value spectrum; Methods) predict the observed change in representational dimensionality from sources to target (compression: target dimensionality < average source dimensionality; expansion: target dimensionality > average source dimensionality)? Each point represents one cortical region. Right, does the activity flow model reproduce the compression-then-expansion pattern of representational dimensionality observed along the cortical hierarchy from sensory through association to motor regions? Solid line shows observed pattern (averaged across regions within each system for each subject); dotted line shows predicted pattern from activity flow model. C-PRO, Concrete Permuted Rule Operations.

many dimensions, whereas lower values indicate that patterns can be described using fewer dimensions. Intuitively, if the 96 task conditions were represented as completely distinct patterns (orthogonal to each other), they would span a near-96-dimensional space (in a region with number of vertices $\gg 96$). However, if some conditions share similar patterns (for example, tasks requiring similar responses), they can be represented in a lower-dimensional subspace. As in previous work⁶, we observed that representational dimensionality strongly correlates with decodability of task conditions from one another (Pearson's $r = 0.755$, spatial autocorrelation (SA)-controlled permutation test;

$P_{\text{perm}} < 0.001$), thus reflecting the granularity of task information in the representations. We defined compression as the difference between the representational dimensionality of the target region and the average representational dimensionality of its source regions (positive values indicate expansion). This allowed us to quantify whether the target region's representation was more compressed (lower dimensionality) or expanded (higher dimensionality) relative to its sources. For connectivity dimensionality, we sought a measure that would parallel our representational dimensionality metric while accounting for the unique structure of connectivity data between source and




Fig. 2 | Connectivity dimensionality predicts representational compression and expansion across the cortex.

a, Cortical organization characterized by discrete networks and continuous gradients. Top, cortical regions colored by functional network identity according to the Cole-Anticevic Brain Network Partition²⁸, dividing the cortex into 12 networks including visual (VIS1 and secondary visual (VIS2)), somatomotor (frontal/motor (SMN-Fr) and parietal/somatosensory (SMN-Par)), cingulo-opercular (CON), dorsal attention (DAN), language (LAN), FPN, AUD, DMN, posterior-multimodal (PMULTI), ventral-multimodal (VMM) and orbito-affective (ORA) networks. Bottom, continuous FC gradients derived from principal component (PC) analysis of the group-averaged, thresholded, parcel-level connectivity graph. Each dot represents one of 360 cortical regions, colored by network identity. The first PC (y axis) captures a uni-transmodal gradient separating sensory/motor regions from association regions. The second PC (x axis) captures a sensory–motor gradient running from sensory through association to motor regions, providing a continuous measure of cortical hierarchy position. We use this sensory–motor gradient (second PC) for visualization throughout. **b**, Representational dimensionality shows compression-then-expansion pattern along the cortical hierarchy. Top, brain maps showing representational dimensionality (D_{Rep} , participation ratio of task RSM eigenspectrum) for each region, averaged across subjects. Bottom, representational dimensionality plotted along the sensory–motor gradient, with regions colored by network. Sensory regions (left) and motor regions (right) show high dimensionality, whereas association regions (center) show compressed representations, forming a U-shaped pattern (black curve: quadratic regression fit with 95% bootstrap CI as gray error band, 1,000 iterations; quadratic model favored over linear, $\Delta\text{AIC} = 3.94$, quadratic coefficient $P_{\text{perm}} = 0.001$, SA-preserving permutation test). **c**, Connectivity dimensionality mirrors the U-shaped pattern. Top, brain

maps showing connectivity dimensionality (D_{FC} , participation ratio of source-to-target connectivity singular value spectrum) for each region as target. Color scale ranges from 8.91 to 38.16. Bottom, connectivity dimensionality plotted along the sensory–motor gradient, quadratic regression fit in black curve and bootstrapped 95% confidence intervals (gray error band) widen owing to fewer regions. Whereas model comparison favors a linear fit ($\Delta\text{AIC} = -9.76$), visual inspection shows curvature driven by regions at gradient extremes (black curve: quadratic fit with 95% bootstrap CI). **d**, Compression and expansion relative to source regions follow the gradient pattern. Top, brain maps showing change in representational dimensionality ($\Delta D_{\text{Rep}} = \text{target dimensionality minus average source dimensionality}$). Negative values (blue, compression) indicate that the target has lower dimensionality than its sources; positive values (red, expansion) indicate higher dimensionality than sources. Bottom, compression/expansion plotted along the sensory–motor gradient, recapitulating the U-shaped pattern (quadratic regression fit in black curve and bootstrapped 95% confidence intervals as gray error band). **e**, Connectivity dimensionality predicts compression and expansion across all regions. Direct correlation between connectivity dimensionality and representational compression/expansion across all 360 cortical regions (Pearson's correlation $r = 0.44$, $P < 0.001$ using SA-preserving permutation tests (using BrainSMASH, two-sided); relationship remains significant after controlling for region size, signal-to-noise ratio, inter-vertex distance, temporal variance and spatial homogeneity). Each dot represents one region, colored by network identity; linear regression fit in black curve and bootstrapped 95% confidence intervals as gray error band. Higher connectivity dimensionality (diverse connectivity patterns among target vertices) predicts representational expansion; lower connectivity dimensionality (similar patterns) predicts compression.

target regions. We therefore performed singular value decomposition on the connectivity matrix—an established approach for analyzing rectangular matrices—and computed the participation ratio of its singular value spectrum. As with representational dimensionality, this yields a continuous measure where higher values indicate more distinct components in the source-to-target connectivity structure.

Connectivity dimensionality reflects the similarity structure of source-to-target connectivity patterns rather than connection sparsity: it was strongly related to the mean pairwise correlation among target vertices' connectivity patterns ($r = -0.79$, $P_{\text{perm}} < 0.001$) but unrelated to connection sparsity ($P > 0.05$; Extended Data Fig. 1).

Using these metrics, we tested whether changes in representational dimensionality from sources to each target correlated with the dimensionality of the intermediate FC. Building on previous work⁶, we plotted transformations along an FC gradient—a continuous axis derived from the principal structure of the whole-brain connectivity matrix, which runs from the sensory through association to motor cortex and captures the cortical processing hierarchy (Fig. 2a), given the intuition that activity flow-based transformations likely occur along

a sensory–association–motor axis during sensory–motor tasks. As expected, we observed a compression-then-expansion trend in task representations along the FC gradient (Fig. 2b,d). As targets, sensory regions typically showed expansion, association regions showed compression, and motor regions showed expansion, forming a U-shaped trajectory of representational compression/expansion. This pattern can be understood through a concrete example: in the visual cortex, different task conditions produce distinct high-dimensional patterns reflecting their unique sensory features (colors, shapes, words and spatial locations). In association regions, representations compress—conditions that share abstract task-relevant features (for example, ‘respond left’ or ‘inhibit response’ or ‘remember this stimulus’) become more similar despite differing sensory details. In the motor cortex, dimensionality increases again because the abstract association-level representations must be translated into specific motor implementations: a ‘respond left’ command could mean left index finger, left middle finger, left hand grip or other motor specifics; a visual attention task and a working memory task might both require ‘maintain fixation’ but with different timing parameters. Model comparisons confirmed

the U-shape for representational dimensionality change (quadratic favored over linear, $\Delta\text{AIC} = 3.94$, $P_{\text{perm}} = 0.001$).

Consistent with our hypothesis, connectivity dimensionality mirrored this trend with a similar U-shaped trajectory (connectivity dimensionality favored a linear fit overall over quadratic ($\Delta\text{AIC} = -9.76$), although with visible curvature at gradient extremes). These observations revealed a significant positive relationship between compression and connectivity dimensionality ($r = 0.44$; group t -test against zero, $t(17) = 16.94$, $P = 4.44 \times 10^{-12}$; SA-preserving permutation test, $P_{\text{perm}} < 0.001$; Fig. 2e). This positive relationship was consistent across sensory, association and motor systems (all $P < 0.001$; Extended Data Fig. 2), demonstrating that this relationship holds across all major cortical systems.

This relationship between connectivity and representational dimensionality is non-trivial for several reasons. First, the two dimensionality measures capture fundamentally different aspects of brain organization. Connectivity dimensionality reflects the geometric structure of how regions connect during rest—essentially, whether vertices of two regions share similar or diverse connectivity patterns. By contrast, representational dimensionality reflects how a single region's activity varies across 96 different task conditions during active task performance. These are distinct sources of variance measured in different brain states. Second, if representational dimensionality were driven by local (within-region) computations alone, it would be independent of task-independent connectivity properties like cross-region connectivity dimensionality. Third, low-dimensional connectivity could theoretically support representational expansion through nonlinear transformations in the target region (for example, if target neurons applied complex, task-dependent computations to their inputs beyond simple weighted summation). Finally, high-dimensional connectivity might have simply reflected noisy or random connectivity patterns, bearing no relationship to representational dimensionality. To rule out confounds, we regressed anatomical (region size and inter-vertex distance), rest-related (temporal signal-to-noise ratio and signal homogeneity) and task-related (condition reliability and betaSNR) properties from connectivity and representational dimensionality (Extended Data Fig. 3). The compression–expansion relationship remained essentially unchanged ($r = 0.4357$, $P_{\text{perm}} < 0.001$). To demonstrate that this result is not a given, we constructed a simulation showing that strong nonlinear local computations in target regions can decouple the relationship between connectivity dimensionality and representational change (Extended Data Fig. 4), demonstrating that the observed relationship is not mathematically inevitable. Together, these results support our hypothesis that lower-dimensional FC is associated with greater representational compression.

A potential concern is that the blood-oxygenation-level-dependent (BOLD) signal in the association cortex, where mixed-selective neurons are spatially intermingled, may poorly reflect underlying representational diversity²⁷. However, we found robust task decodability throughout the cortex (Extended Data Fig. 5), with frontoparietal regions comparable with sensory and motor regions; compression was instead concentrated in default-mode (DMN) and auditory (AUD) networks, indicating specific functional network organization rather than a general BOLD integration artifact.

Cross-region activity flows transform representational geometries

Although these analyses revealed a robust correlation between connectivity dimensionality and representational compression, they leave the generative question open: does low-dimensional connectivity produce representational compression or merely correlate with it? We addressed this using activity flow modeling^{21,22}, which tests whether simulated propagation of task activity over empirical connectivity patterns can reproduce observed representational transformations.

This involved comparing model-generated RSMs (by first generating task activation patterns for all task conditions) with the true RSMs in target regions. The model-generated estimates were based solely on empirical source activity and task-free source-to-target FC. Strong correspondence between model-generated and empirical representational properties would suggest that connectivity patterns play an important role in shaping representational transformations. Importantly, our goal is not to claim that activity flow provides a complete model of representational transformation. Rather, we test whether connectivity geometry plays a systematic role in shaping these transformations. The linear activity flow model establishes a baseline: if connectivity patterns alone (without additional factors like local circuit dynamics, attention or neuromodulation) can predict substantial variance in target representations, this demonstrates that connectivity architecture provides fundamental constraints on how representations transform between regions.

Activity flow modeling simulates how task-evoked activity propagates through the network: for each target vertex, we multiply source-vertex activity patterns by their connectivity weights and sum the weighted contributions (Fig. 1c). This generates predicted activity patterns in each target region based solely on (1) observed task activity in source regions and (2) resting-state connectivity structure. We then compared these model-generated predictions with observed task activity. The model-generated activity patterns were significantly predictive in all regions. Pearson's correlations between observed and predicted activity patterns (averaged across subjects, task conditions and sessions) ranged from $r = 0.16$ to 0.83 across regions (false discovery rate (FDR)-corrected P values, 6.17×10^{-14} to 0.041 ; Fig. 3a, top). To robustly test the importance of structured FC patterns in generating the observed activity patterns, we randomly permuted the FC weights between source and target vertices and generated 1,000 group-level null means. The observed group-level means were found to be significantly outside the null distribution ($P_{\text{connperm}} < 0.001$ for all target regions; Fig. 3a, bottom).

A more conservative test, shuffling connectivity only on the source side while preserving each target's intrinsic connectivity structure, confirmed this specificity: 337 of 360 regions still showed significantly better predictions from intact than source-permuted connectivity ($P_{\text{connperm}} < 0.05$), indicating genuine source-to-target structure rather than target-geometry artifacts (Fig. 3a).

We next tested whether these connectivity-generated activity patterns reproduced the observed representational geometry—specifically, the similarity structure across task conditions captured by RSMs. This similarity structure can reveal important patterns, such as how representations generalize across tasks. We tested the ability of observed connectivity patterns to transform representational geometry between regions, as quantified by RSMs and dimensionality. As expected, model-generated and observed RSMs showed strong correspondence across brain regions. Cosine similarity (averaged across subjects) ranged across regions from 0.22 to 0.84 , whereas P_{FDR} values for t -tests against zero ranged from 3.30×10^{-13} to 2.51×10^{-5} (Fig. 3b, top). Comparison against null values of RSM predictions generated from FC shuffling resulted in group means of all regions' predictions lying significantly outside of the null distribution (all $P_{\text{connperm}} < 0.001$; source-side shuffling only: $P_{\text{connperm}} < 0.05$ in 358 out of 360 regions; Fig. 3b, bottom). The model-predicted representational dimensionality values (averaged across subjects) also showed significant correspondence with observed dimensionality values in target regions ($r = 0.55$, t -test against zero, $t(17) = 10.94$, $P = 4.07 \times 10^{-9}$; Fig. 3c, bottom; comparison to FC-shuffled null distribution, $P_{\text{connperm}} < 0.001$; source-side shuffling only, $P_{\text{connperm}} < 0.001$). Predictions were also consistent within each cortical system (sensory, $r = 0.76$; association, $r = 0.53$; motor, $r = 0.66$; all $P < 0.0001$).

Prediction residuals (observed minus predicted dimensionality) revealed systematic under- and overprediction in specific regions




Fig. 3 | Activity flow over connectivity generates observed representational geometries. **a**, Activity flow accurately predicts task-evoked activity patterns. Top, brain maps showing Pearson's correlation between observed and activity flow-predicted activity patterns for each region, averaged across subjects, sessions and 96 task conditions. Bottom, prediction accuracy (red) for each region plotted against null distribution (gray) generated by randomly shuffling source-to-target connectivity weights (1,000 iterations). All 360 regions show significantly better predictions from intact connectivity than shuffled connectivity (all $P < 0.001$, two-sided permutation test). A more conservative permutation shuffling only source-side connections, which preserves the target region's own connectivity structure, confirmed this specificity in 337 of 360 regions ($P_{\text{comp perm}} < 0.05$, two-sided permutation test). **b**, Activity flow captures representational similarity structure. Top, brain maps showing cosine similarity between observed and predicted RSMs for each region, averaged across subjects. This tests whether connectivity-based predictions capture not just mean activity levels but the full geometry of how task conditions relate to one another. Bottom, RSM prediction accuracy (red) significantly exceeds shuffled-FC null (gray) in all regions (all $P < 0.001$, two-sided permutation test). **c**, Activity flow predicts representational dimensionality and reveals where additional factors operate. Top, brain maps showing prediction residuals (observed minus predicted representational dimensionality). Blue regions indicate systematic underprediction in early visual cortex and frontoparietal regions, suggesting that local dynamics or task-dependent reconfigurations increase dimensionality beyond connectivity-based flow. Red regions indicate overprediction in late

visual areas, medial temporal lobe and salience network, suggesting that specialized local computations compress representations beyond connectivity constraints. Bottom, Pearson's correlation between observed and predicted dimensionality across all 360 regions ($r = 0.55$, $P < 0.001$ using SA-preserving permutation tests, two-sided). Each dot represents one region, colored by network identity; linear regression fit in black curve and bootstrapped 95% confidence intervals as gray error band. The dimensionality prediction was consistent within each cortical system (sensory, $r = 0.76$; association, $r = 0.53$; motor, $r = 0.66$; all $P < 0.0001$). **d**, Activity flow reproduces compression-then-expansion pattern at the subject level. Representational dimensionality for observed (left) and activity flow-predicted (right) representations, plotted for major cortical networks arranged along the sensory–motor gradient. Each dot represents one subject ($n = 18$); box plots show the median and quartiles as box center and bounds, respectively, with whiskers indicating the range of data points within 1.5 times the interquartile range. For both observed and predicted data, dimensionality is averaged across all regions within each network for each subject individually and then plotted here. Statistical comparisons (two-sided paired t -tests, Bonferroni corrected) between networks (VIS1 versus DMN and DMN versus SMN-Fr) test whether the compression-then-expansion pattern holds at the individual subject level. The model successfully reproduces the U-shaped pattern: high dimensionality in visual networks, compression in default-mode network and re-expansion in motor regions. For definitions of abbreviations, refer to Fig. 2.

(Fig. 3c, top), indicating where local circuit dynamics operate beyond the connectivity-based baseline.

We proposed that activity flows over observed connectivity patterns also generate the previously observed compression-then-expansion pattern in dimensionality⁶ (Fig. 2b). This pattern of representational transformations is likely essential for flexible performance of diverse tasks, as sensory information is compressed (generalized) across task conditions for learning transfer, before task-specific representational expansions specify task-appropriate motor responses. Although we used the same dataset as the original study reporting the compression-then-expansion pattern⁶, we expanded the number of modeled task conditions from 45 to 96 to better characterize representational dimensionality (Methods). With this expanded task condition set, we first replicated this compression-then-expansion pattern in our data's task condition split. Examining network-level dimensionality patterns, we found that visual (primary visual (VIS1)) and motor (frontal/motor (SMN-Fr)) networks showed significantly higher dimensionality than DMN and AUD, which exhibited the strongest compression (VIS1 versus DMN, $t(17) = 9.81$, $P = 2.05 \times 10^{-8}$; SMN-Fr versus DMN, $t(17) = 6.46$, $P = 5.86 \times 10^{-6}$; VIS1 versus AUD, $t(17) = 13.61$, $P = 1.43 \times 10^{-10}$; SMN-Fr versus AUD, $t(17) = 10.02$, $P = 1.49 \times 10^{-8}$; Fig. 3d, left). To test whether the activity flow model could reproduce this

pattern, we used a double cross-validated approach (Methods): we constructed hybrid RSMs by measuring the cosine similarity between observed activity patterns of tasks from one run and predicted activity patterns of tasks from a different run. We then analyzed the dimensionality of these double cross-validated RSMs using the participation ratio measure. Our analysis revealed that the model successfully reconstructed the compression-then-expansion pattern of dimensionality along the sensory–association–motor axis, with visual and motor networks again showing significantly higher dimensionality than DMN and AUD (VIS1 versus DMN, $t(17) = 5.49$, $P = 4.02 \times 10^{-5}$; SMN-Fr versus DMN, $t(17) = 5.11$, $P = 8.65 \times 10^{-5}$; VIS1 versus AUD, $t(17) = 9.57$, $P = 2.93 \times 10^{-8}$; SMN-Fr versus AUD, $t(17) = 13.93$, $P = 1.00 \times 10^{-10}$; Fig. 3d, right). Together, these results demonstrate that activity flows shaped by specific geometric properties of connectivity patterns generate the observed representational geometries.

Activity flows over connectivity patterns drive transformations of task representations

Previous analyses established that empirical connectivity patterns reliably generate observed representational geometries in target regions. However, it remains possible that the representations generated by the activity flow models transformed representations only minimally,

with similarities between source and target representations explaining the accurate predictions of target representations. For instance, a source and a target region may both represent the word ‘green’ during the Stroop paradigm, with the activity flow model simply transferring (rather than truly transforming) a representation between these two regions. This would result in flow-predicted RSMs that are no more similar to the target than the source. Furthermore, it remains possible that predicted RSMs of a target region were more similar to the source RSM than the true target RSM. Instead, we proposed that activity flows implement true representational transformations most of the time, predicting that generated RSMs would be more similar to the target RSM than the source RSM for most regions. We therefore next examined whether the geometry of connectivity-generated predictions more closely resembled their source or target geometries.

We quantified this using two cosine distances: d_{pred} (predicted-to-target RSM) and $d_{\text{transform}}$ (predicted-to-source RSM). $d_{\text{pred}} < d_{\text{transform}}$ indicates genuine transformation rather than transfer (Extended Data Fig. 6a). Connectivity-based transformations were non-random, with 357 of 360 regions differing significantly from shuffled-connectivity nulls ($P_{\text{connperm}} < 0.001$; Extended Data Fig. 6c), and predicted transformation magnitude closely tracked observed source-to-target distances ($r = 0.62$, $P_{\text{perm}} < 0.001$; Extended Data Fig. 6b).

Predicted geometries more closely matched target than source geometries in 339 of 360 regions ($P_{\text{FDR}} < 0.05$; 292 at $P_{\text{FDR}} < 0.0001$; Extended Data Fig. 6d), supporting a fundamental role for connectivity patterns in transforming representational geometries across regions. The regions where predictions were closest to targets (smallest d_{pred}) versus sources included both compression zones (association cortex, where low-dimensional connectivity compressed representations) and expansion zones (motor cortex, where high-dimensional connectivity expanded representations), demonstrating that connectivity geometry shapes transformations bidirectionally.

Low-dimensional connectivity shapes cross-task representational similarity

The transition from sensory to the association cortex was accompanied by substantial reductions in FC dimensionality (Fig. 2c). To explore the cognitive benefits of such low-dimensional connectivity, we tested whether target regions with lower connectivity dimensionality show greater similarity among task representations, reflecting latent factors shared across tasks that support the brain’s capacity for diverse task performance.

To isolate cross-task similarity while excluding within-task effects, we computed mean cosine similarities between task condition combinations for each task pair (16 tasks), retaining only positive values across subjects (Methods; Fig. 4a).

Our analysis yielded two key findings. First, we proposed that low-dimensional connectivity actively *generates* cross-task similarities—not merely that low-dimensional regions happen to show similar representations, but that connectivity geometry creates these similarities through activity flow. Consistent with this, regions with lower-dimensional connectivity showed stronger cross-task representational similarity across all task pairs ($r = -0.26$, $t(17) = -24.92$, $P = 8.01 \times 10^{-15}$; SA-preserving permutation test, $P_{\text{perm}} = 0.01$; Fig. 4b).

Activity flow over these same connectivity patterns reproduced the observed cross-task similarity structure across most regions (cosine similarity 0.13–0.99; $P_{\text{connperm}} < 0.05$ in 355/360 regions; Extended Data Fig. 7a).

A simulation using random source activity (ensuring near-zero cross-task similarity at the source) confirmed this as a general property of connectivity-based transformations: lower-dimensional connectivity consistently produced greater cross-task similarity ($\rho = -0.988$, $P = 1.08 \times 10^{-72}$; Methods).

To begin exploring whether these cross-task similarities might reflect interpretable task factors, we conducted an exploratory analysis

at the network level by averaging the cross-task similarity patterns of regions among 12 cortex-wide FC networks according to an established network parcellation²⁸. Although our task set was not designed to systematically investigate shared cognitive components, selective examination of two networks revealed potentially meaningful contrasts. For instance, in VIS1, we observed high cross-task similarity among the ‘go/no-go’, ‘word prediction’ and ‘arithmetic’ tasks, which share visual word processing demands. By contrast, the frontoparietal network (FPN) showed high similarity between the ‘object *n*-back’, ‘verbal *n*-back’ and ‘spatial map’ tasks despite their different stimulus types, likely reflecting shared cognitive control demands (Fig. 4c, top).

This pattern of network-relevant latent task factors was replicated when connectivity-predicted estimates were analyzed, demonstrating the role of connectivity patterns in the generation of meaningful latent task factors (Fig. 4c, bottom). These network-specific patterns are non-trivial: VIS1’s high similarity for word-processing tasks and FPN’s high similarity for cognitive control tasks demonstrate that connectivity geometry generates functionally interpretable latent task factors aligned with each network’s computational role (see observed and predicted network comparison results in Extended Data Table 1). Cross-task similarities at the network level significantly differed from null models based on shuffled connectivity for all networks (Extended Data Fig. 7b). Future work with a more systematic task battery, perhaps involving the reuse of instruction elements across tasks (for example, to allow for cross-condition generalization analysis), is needed to thoroughly investigate these network- and region-specific patterns.

Together, these results show that low-dimensional connectivity generates shared representational structure across tasks, further supporting a fundamental role for intrinsic connectivity in specifying task representations throughout the cortex.

Activity flow-driven expansions generate task-variable conjunctions

We proposed that connectivity patterns expanding representational dimensionality specifically support task-variable conjunctions: activity patterns present only when two or more task features co-occur (for example, ink color blue *and* Stroop-incongruent), rather than whenever any single feature is present. Such conjunctive coding allows cognitive variables to interact in situation-specific ways—for instance, processing color blue differently depending on congruency—unlike low-dimensional compositional representations that are reused identically across contexts. Although our task set was not systematically designed to investigate conjunctive coding, we identified a subset of five tasks containing separable sensorimotor and cognitive variables that allowed for an exploratory test of this hypothesis.

To test which task features drive representational similarity, we identified five tasks with separable sensorimotor and cognitive variables (Fig. 5a) and constructed four template RSMs: (1) sensory similarity (based on stimulus features like color, shape and words), (2) cognitive similarity (based on task demands like congruency and difficulty), (3) motor similarity (based on required button press responses) and (4) conjunction similarity (multiplicative combination of sensory $\times$ cognitive $\times$ motor features, capturing non-additive interactions). Using ridge regression, we fit these four vectorized matrices to the vectorized observed neural RSM (Fig. 5a).

As expected, sensory, cognitive and motor similarities were encoded in their respective brain systems (Fig. 5b). Critically, conjunction-term coefficients correlated positively with representational expansion ($r_{\text{conj}} = 0.46$, $P = 2.92 \times 10^{-7}$; Fig. 5c, top), a relationship stronger than for cognitive or motor terms individually and comparable with the sensory term (Extended Data Table 2). Activity flow-predicted conjunction strength similarly correlated with predicted expansion ($r = 0.34$, $P = 0.006$; Fig. 5c, bottom), exceeding predicted correlations for the other three terms (Extended Data Table 2),




Fig. 4 | Low-dimensional connectivity generates functionally interpretable cross-task similarities. **a**, Deriving cross-task similarity from RSMs. Cross-task similarity quantifies how similar a region's representations are across different tasks, revealing shared latent factors. Left, example 96 × 96 task condition RSM from a single region showing cosine similarity between all pairs of task conditions (averaged across subjects). Blocked structure reflects tasks with multiple conditions (for example, conditions 1–8 are Stroop task conditions, showing high within-task similarity). Right, to isolate cross-task similarity, we first retained only positive similarity values from the 96 × 96 condition RSM (negative similarities are difficult to interpret) and then averaged these positive values within each task pair to create a 16 × 16 task-level matrix. Off-diagonal elements represent cross-task similarity (for example, how similar are representations between Stroop and word prediction tasks?). For each region, we computed mean cross-task similarity by averaging all off-diagonal values in this 16 × 16 matrix, yielding a single measure of how much the region shares representational structure across tasks. **b**, Regions with low-dimensional connectivity show higher cross-task similarity. Pearson's correlation between connectivity dimensionality (D_{FC}) and mean cross-task similarity across all 360 regions ($r = -0.26$, $P < 0.001$ using SA-preserving permutation tests, two-sided). Each dot represents one region, colored by network identity; linear fit in black curve and bootstrapped 95% confidence intervals as gray error band. Lower connectivity dimensionality (similar connectivity patterns among target vertices) predicts greater similarity of task representations, suggesting that these regions encode shared latent factors that generalize across tasks. Higher connectivity dimensionality predicts

task-specific representations. **c**, Activity flow-predicted cross-task similarities reflect network-specific functional specialization. Comparison of cross-task similarity patterns between VIS1 (purple) and FPN (red) for task pairs with shared features. For each subject, cross-task similarity values were averaged across all regions within each network. Each box plot shows the distribution across subjects ($n = 18$); each dot represents one subject's network-averaged cross-task similarity for that task pair; box plots show the median and quartiles as box center and bounds, respectively, with whiskers indicating the range of data points within 1.5 times the interquartile range. Top, observed cross-task similarities. Bottom, model-predicted cross-task similarities from activity flow. Left, task pairs sharing sensory features (word stimuli: go/no-go, word prediction and arithmetic tasks): VIS1 shows significantly higher cross-task similarity than FPN across all three task pairs shown (all $P < 0.00001$, two-sided paired t -tests, Bonferroni corrected; exact values of all comparisons in Extended Data), reflecting VIS1's sensitivity to shared visual word processing demands. Right, task pairs sharing cognitive features (working memory: object n -back, verbal n -back and spatial map tasks): FPN shows higher cross-task similarity than VIS1, reflecting FPN's sensitivity to shared cognitive control demands. Critically, activity flow predictions (bottom) reproduce these network-specific patterns, demonstrating that connectivity geometry alone generates functionally interpretable latent task factors aligned with each network's computational specialization. C-PRO, Concrete Permuted Rule Operations. For color keys and definitions of other abbreviations, refer to Fig. 2.

confirming that connectivity-driven expansion specifically supports conjunctive coding.

Beyond conjunctive coding, we also tested whether representational dimensionality relates to cognitive flexibility more broadly, using fluid intelligence data from the HCP^{29,30} ($n = 352$). Region-averaged dimensionality positively predicted general fluid intelligence ($r = 0.29$, $P = 1.5 \times 10^{-8}$), particularly in sensory and association regions (Extended Data Fig. 8a–d). These findings generate specific hypotheses that can be tested in future studies using carefully designed task batteries that systematically vary sensory, cognitive and motor demands to more definitively determine how connectivity-based transformations contribute to high-dimensional task representations and support flexible cognitive control through

conjunctive coding. As a replication check, we checked for generalizability of the compression-then-expansion pattern in representational dimensionality in the HCP dataset as well. Using cross-validated RSMs constructed from 24 task conditions across seven tasks, we found that regional representational dimensionality exhibited the same U-shaped pattern along the sensorimotor gradient at the group level (Extended Data Fig. 8e). To quantify the consistency of compression in association regions relative to sensory and motor regions across individuals, we calculated a compression ratio for each subject (mean dimensionality of sensory and motor regions divided by association region dimensionality). A compression ratio greater than 1 indicated representational compression in the association cortices. This pattern proved highly reliable, with 98.6% (347/352) of subjects showing




Fig. 5 | Connectivity-driven dimensionality expansion supports conjunctive encoding of task features. a, Representational similarity analysis to estimate conjunction strength. We identified a subset of 5 tasks (44 task conditions total) with separable sensorimotor and cognitive features: Stroop, go/no-go, IAPS-affective, IAPS-emotional and mental rotation. For each region, we fit the observed 44×44 task condition RSM (leftmost plot, example shown from right area PH in the visual cortex) using ridge regression against four template RSMs encoding different feature types. Sensory template: similarity based on stimulus type (word, image and spatial array). Cognitive template: similarity based on task demand level (low versus high cognitive effort). Motor template: similarity based on required button press responses (left/right index/middle finger). Conjunction template: multiplicative combination of all three features (sensory × cognitive × motor), capturing non-additive feature interactions where specific combinations (for example, ‘word stimulus + high demand + left index response’) produce distinct representations beyond the sum of individual features. Ridge regression yields four coefficients per region, with the conjunction coefficient quantifying how strongly that region encodes feature combinations. **b**, Observed feature encoding shows expected anatomical specificity. Brain maps showing fitted coefficients (averaged across subjects) for each feature template, with only significant values shown (one-sample *t*-test against zero, FDR corrected). Sensory: highest in the visual cortex (9/10 highest coefficients). Cognitive: highest in the association cortex (7/10 highest in frontoparietal and default-mode regions). Motor: highest in the motor

and premotor cortices (8/10 highest). Conjunction: more distributed, with highest values in regions showing representational expansion. This anatomical specificity validates the representational similarity analysis approach and demonstrates that sensory, cognitive and motor features are encoded in their respective brain systems. **c**, Dimensionality expansion predicts conjunction strength in both observed and predicted data. Top, correlation between observed conjunction coefficient and observed representational compression/expansion (ΔD_{Rep}) across all 360 regions ($r = 0.46$, $P < 0.001$ using SA-preserving permutation tests, two-sided; linear fit in black curve and bootstrapped 95% confidence intervals as gray error band). Each dot represents one region, colored by network. Regions showing expansion (positive ΔD_{Rep}) encode stronger conjunctive representations, whereas regions showing compression encode weaker conjunctions. This relationship is stronger than correlations with individual feature terms (sensory, $r = 0.15$; cognitive, $r = 0.28$; motor, $r = 0.04$), demonstrating that expansion specifically supports conjunctive coding. Bottom, correlation between activity flow-predicted conjunction coefficient and predicted ΔD_{Rep} ($r = 0.34$, $P < 0.001$ using SA-preserving permutation tests, two-sided; linear fit in black curve and bootstrapped 95% confidence intervals as gray error band). The model reproduces the relationship between expansion and conjunction strength using only connectivity geometry and source activity, demonstrating that high-dimensional connectivity patterns systematically support conjunctive representations. For color keys and definitions of abbreviations, refer to Fig. 2.

compression in association regions, providing strong validation of this organizational principle (Extended Data Fig. 8f).

Discussion

We identified a systems-level organizing principle by which intrinsic connectivity dimensionality directs the cross-region compression and expansion of task representations: low-dimensional connectivity produces compressed representations with shared similarities across tasks, whereas high-dimensional connectivity expands representations to integrate task variables for specific cognitive demands. Activity flow modeling demonstrated that this connectivity-based mechanism actively generates observed representational geometries across sensory, association and motor systems, rather than merely correlating with them. Regions with lower-dimensional connectivity showed stronger cross-task similarities, suggesting shared latent features supporting generalization, whereas high-dimensional regions supported conjunctive representations situating that information for specific contexts. This mechanism was replicated in an independent HCP dataset, where representational dimensionality also predicted general fluid intelligence.

Our simulations show that strong nonlinear local computations can fully decouple connectivity dimensionality from representational transformations (Extended Data Fig. 4), yet we empirically observe robust coupling ($r = 0.44$), positioning the brain in an intermediate regime where network architecture provides systematic constraints while local circuit dynamics contribute region-specific modulation, as evidenced by our prediction residuals (Fig. 3c). This suggests a design principle of architectural parsimony: connectivity patterns, once established, provide stable constraints shaping representations across tasks without requiring task-dependent reconfiguration of local circuits, with residual patterns indicating a principled partitioning of computational labor between distributed architecture and specialized local dynamics. This framework generates testable predictions: diseases affecting white matter connectivity should produce predictable distortions in representational geometry, and connectivity-targeted interventions should have systematic downstream effects on representational transformations.

Our findings provide empirical validation of theoretical predictions from computational neuroscience regarding how connectivity structure constrains representational geometry^{18–20,31–33}, which have

been extensively characterized in artificial neural networks, but not previously demonstrated in biological circuits³⁴. Our resting-state FC-based approach offers a functionally defined measure of connectivity constraints complementary to the structural connectivity emphasized in theoretical work.

The present study expands previous connectivity-based models of task processing from structured rule combinations to diverse cognitive demands. Earlier work using the Concrete Permuted Rule Operations paradigm demonstrated that task rule information transfers between brain regions via resting-state connectivity²²; we show that this principle generalizes to diverse cognitive tasks lacking such structured relationships, extending the activity flow modeling framework to show that connectivity patterns actively transform, not merely transfer, task representations between regions. Although previous work emphasized information transfer between regions encoding similar information, we demonstrate that intrinsic connectivity patterns support representational transformations between functionally diverse connected regions. Our results also extend previous findings showing that pairwise similarity between regional RSMs correlates with resting-state connectivity at the parcel level⁶. We build upon this finding by developing a mechanistic model based on finer-grained, vertex-level connectivity that characterizes how connectivity geometry constrains these representational transformations.

Our study instantiates a key theoretical framework for understanding cognitive flexibility: examining how multiple tasks mutually constrain neural activity patterns to reveal mechanisms of generalization³⁵. Activity flows over intrinsic connectivity generated widespread shared representations across 355 of 360 regions, comparable with cognitive encoding models explaining variance in 86% of cortical voxels², but through direct connectivity-mediated mappings rather than predefined cognitive factors⁵. This connectivity-based framework leverages a brain-based ontology³⁶ that remains open to discovering latent cross-task factors not corresponding to existing cognitive constructs.

Our findings complement a growing body of work in animal models examining how inter-regional connectivity shapes neural computation^{37–39}. Studies in non-human primates have demonstrated that functional connectivity between sensory and association areas constrains how information is communicated and transformed across regions⁴⁰, with communication subspaces between areas shaping what information can be transmitted^{41–43}. Recent work in mice similarly shows that inter-regional coordination structures cortical dynamics and supports flexible behavior⁴⁴. Our findings extend these principles to human whole-brain organization during diverse cognitive tasks, demonstrating that connectivity geometry operates as a systematic organizing principle across species and scales, from sensory processing in animal models to cognitive flexibility in humans. A remaining limitation is that BOLD-based measures cannot resolve contributions of mixed selectivity at the sub-voxel level, which may manifest as low dimensionality and low decodability and could also affect connectivity dimensionality estimation. Resolving this fully requires cross-modal validation—for example, estimating connectivity dimensionality from noise correlations between neuronal populations in electrophysiology recordings and testing whether this predicts representational dimensionality change between those populations across task conditions.

A key conceptual clarification is needed: activity flow modeling here involves literal convergence (many source vertices projecting to fewer target vertices); however, we define compression/expansion relative to the average dimensionality of individual source regions rather than pooled source-vertex dimensionality, allowing us to characterize region-to-region transformations while acknowledging the underlying integration across many sources. This framework reveals that convergent architectures can produce dimensionality expansion through a plausible linear mechanism: when source regions occupy distinct low-dimensional subspaces and target vertices receive diverse

weighted mixtures (high connectivity dimensionality), the target can occupy a higher-dimensional space spanning combinations of those subspaces, potentially explaining our conjunctive coding results (Fig. 5), as linear mixing of diverse sources naturally creates feature combinations approximating conjunctive coding. This generates a testable prediction: expansion-showing regions should receive input from sources with lower cross-source representational similarity.

The present results reveal how the brain's intrinsic connectivity architecture enables both generalization across tasks and specific task implementation through systematic control of representational dimensionality. The geometric properties of neural representations have proven highly informative for understanding how brain regions support multiple task contexts^{8,11,14,17}. Building on this framework, we show that these geometries are fundamentally shaped by connectivity patterns: low-dimensional connectivity supports abstraction of shared features across tasks, as evidenced by meaningful task groupings in our network-level analyses, such as visual word processing demands in visual networks and cognitive control demands in FPNs. These findings align with the neural reuse framework^{45,46}, which proposes that brain regions and circuits support multiple cognitive functions by sharing computational resources across tasks; here we characterize this reuse in terms of the geometric properties of connectivity and representations. Complementarily, high-dimensional connectivity enables the integration of multiple task variables into conjunctive representations, which we measured using established representational similarity analysis methods^{15,16} to show how abstract features are situated for specific cognitive demands. For example, a task-implementing conjunctive representation may only be active when the rule representation 'press your left index finger when you see green' is active and a representation of green is active, with activity flowing from that representation to a left index finger representation in the motor system implementing task behavior. This connectivity-based mechanism—abstraction through compression and specificity through expansion—provides a principled basis for cognitive flexibility, enabling both generalization across tasks and task-specific processing.

Higher representational dimensionality, particularly in sensory and association regions, predicted general fluid intelligence⁴⁷. Although this finding seems to conflict with the compression-then-expansion pattern, it likely reflects complementary levels of organization: the compression-then-expansion pattern is a canonical within-subject architecture supporting efficient task processing, whereas absolute dimensionality within that architecture reflects an individual's capacity for flexible cognition (Extended Data Fig. 8).

In summary, our study reveals a fundamental organizing principle by which intrinsic connectivity architecture enables cognitive flexibility through systematic transformation of task representations, providing a framework for understanding both healthy cognitive function and its disruption in neurological and psychiatric conditions.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41593-026-02486-8>.

References

1. Miller, E. K. & Cohen, J. D. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* **24**, 167–202 (2001).
2. Nakai, T. & Nishimoto, S. Quantitative models reveal the organization of diverse cognitive functions in the brain. *Nat. Commun.* **11**, 1142 (2020).
3. Nakai, T. & Nishimoto, S. Representations and decodability of diverse cognitive functions are preserved across the human cortex, cerebellum, and subcortex. *Commun. Biol.* **5**, 1–13 (2022).

4. Flesch, T., Juechems, K., Dumbalska, T., Saxe, A. & Summerfield, C. Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* **110**, 1258–1270.e11 (2022).
5. Walters, J. et al. Predicting brain activation maps for arbitrary tasks with cognitive encoding models. *NeuroImage* **263**, 119610 (2022).
6. Ito, T. & Murray, J. D. Multitask representations in the human cortex transform along a sensory-to-motor hierarchy. *Nat. Neurosci.* **26**, 306–315 (2023).
7. Badre, D., Bhandari, A., Keglovits, H. & Kikumoto, A. The dimensionality of neural representations for control. *Curr. Opin. Behav. Sci.* **38**, 20–28 (2021).
8. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
9. Takagi, Y., Hunt, L. T., Woolrich, M. W., Behrens, T. E. & Klein-Flügge, M. C. Adapting non-invasive human recordings along multiple task-axes shows unfolding of spontaneous and over-trained choice. *eLife* **10**, e60988 (2021).
10. Brincat, S. L., Siegel, M., von Nicolai, C. & Miller, E. K. Gradual progression from sensory to task-related processing in cerebral cortex. *Proc. Natl Acad. Sci. USA* **115**, E7202–E7211 (2018).
11. Bernardi, S. et al. The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell* **183**, 954–967.e21 (2020).
12. Ito, T. et al. Compositional generalization through abstract representations in human and artificial neural networks. In *Advances in Neural Information Processing Systems* (eds Koyejo, S. et al.) Vol. 35 32225–32239 (Curran Associates, 2022).
13. Vaidya, A. R., Jones, H. M., Castillo, J. & Badre, D. Neural representation of abstract task structure during generalization. *eLife* **10**, e63226 (2021).
14. Rigotti, M. et al. The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).
15. Kikumoto, A. & Mayr, U. Conjunctive representations that integrate stimuli, responses, and rules are critical for action selection. *Proc. Natl Acad. Sci. USA* **117**, 10603–10608 (2020).
16. Kikumoto, A., Bhandari, A., Shibata, K. & Badre, D. A transient high-dimensional geometry affords stable conjunctive subspaces for efficient action selection. *Nat. Commun.* **15**, 8513 (2024).
17. Fusi, S., Miller, E. K. & Rigotti, M. Why neurons mix: high dimensionality for higher cognition. *Curr. Opin. Neurobiol.* **37**, 66–74 (2016).
18. Mastrogiuseppe, F. & Ostojic, S. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron* **99**, 609–623.e29 (2018).
19. Dubreuil, A., Valente, A., Beiran, M., Mastrogiuseppe, F. & Ostojic, S. The role of population structure in computations through neural dynamics. *Nat. Neurosci.* **25**, 783–794 (2022).
20. Ostojic, S. & Fusi, S. Computational role of structure in neural activity and connectivity. *Trends Cogn. Sci.* **28**, 677–690 (2024).
21. Cole, M. W., Ito, T., Bassett, D. S. & Schultz, D. H. Activity flow over resting-state networks shapes cognitive task activations. *Nat. Neurosci.* **19**, 1718–1726 (2016).
22. Ito, T. et al. Cognitive task information is transferred between brain regions via resting-state network topology. *Nat. Commun.* **8**, 1027 (2017).
23. Cocuzza, C. V. et al. Distributed network flows generate localized category selectivity in human visual cortex. *PLoS Comput. Biol.* **20**, e1012507 (2024).
24. Ito, T., Hearne, L., Mill, R., Cocuzza, C. & Cole, M. W. Discovering the computational relevance of brain network organization. *Trends Cogn. Sci.* **24**, 25–38 (2020).
25. Cole, M. W. Chapter 4 - Network modeling: The explanatory power of activity flow models of brain function. In *Computational and Network Modeling of Neuroimaging Data* (ed. Kay, K.) 81–117 (Academic Press, 2024).
26. Cole, M. W. Cognitive flexibility as the shifting of brain network flows by flexible neural representations. *Curr. Opin. Behav. Sci.* **57**, 101384 (2024).
27. Bhandari, A., Gagne, C. & Badre, D. Just above chance: is it harder to decode information from prefrontal cortex hemodynamic activity patterns? *J. Cogn. Neurosci.* **30**, 1473–1498 (2018).
28. Ji, J. L. et al. Mapping the human brain's cortical-subcortical functional network organization. *NeuroImage* **185**, 35–57 (2019).
29. Barch, D. M. et al. Function in the human connectome: task-fMRI and individual differences in behavior. *NeuroImage* **80**, 169–189 (2013).
30. Van Essen, D. C. et al. The WU-Minn Human Connectome Project: an overview. *NeuroImage* **80**, 62–79 (2013).
31. Langdon, C. & Engel, T. A. Latent circuit inference from heterogeneous neural responses during cognitive tasks. *Nat. Neurosci.* **28**, 665–675 (2025).
32. Hu, Y. & Sompolinsky, H. The spectrum of covariance matrices of randomly connected recurrent neuronal networks with linear dynamics. *PLoS Comput. Biol.* **18**, e1010327 (2022).
33. Recanatesi, S., Ocker, G. K., Buice, M. A. & Shea-Brown, E. Dimensionality in recurrent spiking networks: global trends in activity and local origins in connectivity. *PLoS Comput. Biol.* **15**, e1006446 (2019).
34. Langdon, C., Genkin, M. & Engel, T. A. A unifying perspective on neural manifolds and circuits for cognition. *Nat. Rev. Neurosci.* **24**, 363–377 (2023).
35. Nau, M., Schmid, A. C., Kaplan, S. M., Baker, C. I. & Kravitz, D. J. Centering cognitive neuroscience on task demands and generalization. *Nat. Neurosci.* **27**, 1656–1667 (2022).
36. Buzsáki, G. *The Brain from Inside Out* (Oxford Univ. Press, 2019).
37. Brynildsen, J. K., Rajan, K., Henderson, M. X. & Bassett, D. S. Network models to enhance the translational impact of cross-species studies. *Nat. Rev. Neurosci.* **24**, 575–588 (2023).
38. Higley, M. J. & Cardin, J. A. Spatiotemporal dynamics in large-scale cortical networks. *Curr. Opin. Neurobiol.* **77**, 102627 (2022).
39. Wang, X.-J. Theory of the multiregional neocortex: large-scale neural dynamics and distributed cognition. *Annu. Rev. Neurosci.* **45**, 533–560 (2022).
40. Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M. & Kohn, A. Cortical areas interact through a communication subspace. *Neuron* **102**, 249–259.e4 (2019).
41. Semedo, J. D. et al. Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nat. Commun.* **13**, 1099 (2022).
42. Srinath, R., Ruff, D. A. & Cohen, M. R. Attention improves information flow between neuronal populations without changing the communication subspace. *Curr. Biol.* **31**, 5299–5313.e4 (2021).
43. Stoll, F. M. & Rudebeck, P. H. Decision-making shapes dynamic inter-areal communication within macaque ventral frontal cortex. *Curr. Biol.* **34**, 4526–4538.e5 (2024).
44. MacDowell, C. J. et al. Multiplexed subspaces route neural activity across brain-wide networks. *Nat. Commun.* **16**, 3359 (2025).
45. Anderson, M. L. Neural reuse: a fundamental organizational principle of the brain. *Behav. Brain Sci.* **33**, 245–266 (2010).
46. Anderson, M. L. Evolution of cognitive function via redeployment of brain areas. *Neuroscientist* **13**, 13–21 (2007).
47. Gray, J. R., Chabris, C. F. & Braver, T. S. Neural mechanisms of general fluid intelligence. *Nat. Neurosci.* **6**, 316–322 (2003).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with

the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2026

Methods

Multi-domain task battery dataset

Portions of this section are paraphrased from the original data study⁴⁸.

We used the multi-domain task battery dataset⁴⁸, originally designed to examine functional task boundaries in the human cerebellum. The dataset comprises whole-brain resting-state and task-state fMRI data from 24 right-handed participants (16 women and 8 men; mean age 23.8 ± 2.6 years) at Western University. Participants provided informed consent under an approved experimental protocol. A total of 26 participants were originally scanned over two sessions, but 2 of them were excluded for not completing all the scanning runs. Participants were compensated based on an experimental protocol approved by the institutional review board at Western University, where the data were originally collected.

The multi-domain task battery includes 26 cognitive tasks with up to 45 distinct task conditions per participant, divided into sets A and B. Participants completed set A and B tasks in separate identical sessions, with inter-session intervals of 2–3 weeks for half the participants and approximately 1 year for the others. Each set involved two 10 min imaging runs. For 18 out of 24 participants, an additional resting-state scan (two 10 min runs) was conducted separately from the ‘rest’ block in task sessions. Our analyses utilize task data from these 18 participants to relate task and resting-state data within the same individuals.

The dataset encompasses a broad range of cognitive processes. Sets A and B shared eight common tasks and nine unique tasks each, totaling 17 tasks per set. Tasks were performed once per imaging run using an interleaved block design. Each task block consisted of a 5 s instruction screen followed by 30 s of continuous task performance, totaling 35 s blocks. Most tasks comprised 10–15 trials per block, with the number of trials per task ranging from 1 to 30. For 17 tasks requiring motor input, participants used a four-button box with the index or middle finger of either or both hands. All tasks within each set were completed in a single imaging run, ensuring a consistent baseline across tasks for all participants.

We selected 16 active, visual tasks (excluding the active auditory interval timing task), decomposing them into 96 task conditions based on stimulus type, motor response and context variables (Supplementary Table 1; decomposition unique to this study). Focusing on visual tasks allowed us to treat visual regions as sensory in the sensory–association–motor hierarchy. The fMRI acquisition parameters were as follows: repetition time = 1 s, field of view = 20.8 cm, phase-encoding direction P to A, 48 slices, 3 mm thickness and in-plane resolution = 2.5×2.5 mm².

fMRI preprocessing

This section is paraphrased from the preprocessing details of Ito and Murray⁶, which used the same preprocessing pipeline.

The preprocessing of the fMRI data in this study utilized the HCP preprocessing pipeline⁴⁹, which was executed through the Quantitative Neuroimaging Environment & Toolbox (QuNex, version 0.61.17)⁵⁰. The HCP pipeline encompassed several key steps in data preparation. These included the reconstruction and segmentation of anatomical images, similar procedures for EPI (echo-planar imaging) data, alignment of brain images to the standardized MNI152 template for spatial normalization and corrections for head movement during scanning. In addition to the standard HCP pipeline, we applied further steps of nuisance regression, as detailed in the next section.

fMRI task activation estimation

This study used a modified version of the single-subject beta series regression technique to analyze the fMRI task data⁵¹. The analysis focused on estimating vertex-wise activations throughout the cortical sheet. Each task condition was represented by a distinct task regressor, modeled as boxcar functions (with 1 for on during each trial period of the corresponding task condition and 0 for off, with a separate

regressor for each task’s instruction period). Given the high practice accuracy of tasks (>85%) prior to scanning, we did not exclude error trials during this estimation. These regressors were then convolved with the Statistical Parametric Mapping (SPM) canonical hemodynamic response function to account for temporal lags in hemodynamic responses⁵². The ‘rest’ block in each run was left unmodeled and thus formed part of the baseline. The coefficients of these regressors were used to represent task condition activations.

Task General Linear Models (GLMs) were fit per run using ridge regression (Python 3.8.5, ‘ridge_corr’ function, <https://github.com/alexhuth/ridge/blob/master/ridge.py>) to address multicollinearity from closely spaced task conditions. In each run, the optimal regularization parameter (alpha) was selected using 4-fold cross-validation, testing a range of alpha values from 1 to 1,000 in 20 logarithmic steps. The optimal alpha was identified by maximizing the average correlation with the testing fold across cross-validation iterations. Once the optimal alpha was determined, task estimates were obtained for the entire run using this value.

To minimize non-neural noise and avoid potential artifacts, task GLMs were performed concurrently with nuisance regression, as suggested by previous research⁵³. The nuisance parameters included 24 motion-related regressors (6 motion parameters, their time derivatives and quadratic terms) and 8 physiological nuisance signals (mean time series from white matter and ventricle voxels, their derivatives and quadratic terms). In total, 32 nuisance regressors were incorporated alongside the task regressors to extract refined task activation estimates.

Prior to running the regression, both the time series data and the regressor time series were z scored along the time axis. Multiple values of activity estimates across different runs were averaged to obtain one estimate per task condition per session. To assess the reliability of activity estimates, we computed Pearson’s correlation between the estimates from the two identical sessions after averaging across task conditions in each parcel (region) for each subject (range across parcels, averaged across subjects, 0.14–0.83). Although these correlations were not used to filter out any task conditions, regions or subjects, they provided a measure of estimated stability across sessions.

FC estimation

We used the two 10 min resting-state runs per subject to measure FC. Following preprocessing and nuisance regression (as above), residual time series from the two runs were concatenated per subject before estimating FC in two steps.

First, we estimated sparse parcel-level FC between mean Glasser parcellation⁵⁴ time series using graphical lasso (regularized partial correlation, L1 penalty; ‘GGLasso’ package⁵⁵), which improves validity against structural connectivity and individual-level reliability⁵⁶. This two-step approach—sparse parcel-level graph followed by dense vertex-level estimation—balances statistical power with spatial specificity while avoiding the intractability of estimating full vertex-to-vertex connectivity across the cortex (millions of connections). The regularization parameter (alpha) was optimized over 0.001–1 (10 log-spaced steps), selecting the value minimizing negative log-likelihood averaged across folds.

The parcel-level graph identified each region’s connected neighbors (sources). Dense vertex-level FC was then estimated between each target region’s vertices and all vertices of its source regions, producing a rectangular connectivity matrix (source vertices $\times$ target vertices) per target region. For each target vertex, source vertices served as regressors in a cross-validated principal component regression, chosen to handle thousands of source vertices with only ~1,200 timepoints while reducing overfitting and preserving major covariance patterns^{21,22,57}. The number of components was selected via 10-fold cross-validation (minimizing mean squared error), yielding a mean of 57.1 components across subjects (range, 51–71). Source vertices within 10 mm of the

target region's boundary were excluded to prevent circularity from SA. Goodness of fit (r^2) of the regression to training data ranged from 0.24 to 0.66 across regions, averaged across subjects.

RSM estimation

For each of the 360 Glasser parcels, we computed cosine similarity between task condition activity vectors, cross-validated between the two sessions (for example, condition A from session 1 versus condition B from session 2). Cosine similarity was chosen over Pearson's correlation because it captures both activation pattern and magnitude across vertices. Cross-validating between sessions yielded a non-trivial diagonal, reflecting test–retest reliability. This produced a 96×96 RSM per parcel per subject, corresponding to the 96 task conditions (16 tasks).

Representational and connectivity dimensionality

We measured representational dimensionality using the participation ratio of the RSM's eigenvalue spectrum (eigendecomposition), which quantifies the rate of spectral decay:

$$D_{\text{Rep}} = \frac{\left(\sum_{i=1}^n \lambda_i\right)^2}{\sum_{i=1}^n \lambda_i^2}$$

where λ_i are the eigenvalues of the RSM. This measure corresponds to the number of principal components explaining ~80% of variance⁵⁸, but was preferred for its continuous scale, enabling finer-grained comparisons across regions and subjects.

We related representational dimensionality to the decoding of 96 task conditions using a minimum distance classifier based on cosine distance (cross-validated across 2 sessions), intuitable directly from the RSM: a condition is decodable if its diagonal value exceeds the other values in its row/column.

For connectivity dimensionality, we sought a parallel measure accounting for the rectangular structure of connectivity data (connecting distinct source and target vertex sets). We performed singular value decomposition on the dense FC matrix between each target region's vertices and its source regions:

$$W = U \Sigma V^T.$$

We then calculated the participation ratio of the singular value spectrum as a continuous measure of connectivity dimensionality for each target region:

$$D_{\text{FC}} = \frac{\left(\sum_{i=1}^t s_i\right)^2}{\sum_{i=1}^t s_i^2}$$

where s_i are singular values following singular value decomposition.

These measures were calculated separately for each subject. Higher representational dimensionality indicates more distinct condition-specific patterns; lower connectivity dimensionality reflects more convergent, homogeneous source input (see Results for interpretation).

We tested parcel-specific properties that could confound both dimensionality measures: target region size (vertex count), signal-to-noise ratio (mean RSM diagonal), mean inter-vertex distance (mean geodesic distance to neighboring mesh-connected vertices⁵⁹), temporal variance of resting-state activity (s.d. of each vertex's time series across time, averaged across vertices) and spatial homogeneity (s.d. across vertices within each parcel at each timepoint, averaged across time). These confounds were regressed out of connectivity dimensionality prior to correlating the residuals with representational dimensionality. Across-parcel correlations used SA-preserving permutation tests (using the BrainSMASH toolbox⁶⁰, 1,000 null surrogate maps were generated, making the minimum possible P value 0.001, with any value under that reported as $P < 0.001$).

Discrete brain systems and continuous FC gradient estimation

We delineated sensory, association, and motor regions using the Glasser et al. (2016) multimodal parcellation (HCP-MMP1.0; 180 areas/hemisphere), selected for its integration of cortical architecture (myelin and thickness), function (task-fMRI), connectivity and topography, providing improved spatial localization over unimodal approaches^{54,61}, combined with the Cole-Anticevic Brain Network Partition²⁸, which divides the parcellation into 12 functional networks. We categorized primary and secondary visual network regions as sensory (restricted to visual networks given our focus on visual tasks); somatomotor regions were split into somatosensory (SMN-parietal, posterior to the central sulcus) and motor (SMN-frontal, anterior); all remaining regions were categorized as association.

To derive continuous connectivity gradients, we thresholded the parcel-level sparse FC graph (averaged across subjects) to retain the top 20% of values⁶² and then applied principal component analysis. The first component ran from unimodal (sensory/motor) to transmodal (association) regions. The second exhibited continuous variation from the sensory through association to motor regions, representing the broad cortical processing hierarchy—we refer to this as the sensorimotor FC gradient, which we visually confirmed replicates previous reports and complements our discrete network categorization.

Connectivity pattern similarity and sparsity

To interpret connectivity dimensionality in concrete terms, we analyzed target-vertex connectivity patterns (vectors of connectivity weights from source vertices to each target vertex), motivated by the activity flow framework's focus on how activity is generated in target populations; connectivity dimensionality is also capped at target region size, the smaller dimension of the dense connectivity matrix. For each target region, we computed Pearson's correlation between all pairs of vertices' connectivity patterns ($t \times t$ matrix, where t is the number of vertices) and derived the mean correlation as a measure of overall pattern similarity.

Sparsity analysis. To test whether connectivity dimensionality reflects sparsity rather than geometric structure, we defined approximate sparsity (as principal component regression produces dense, all-nonzero matrices) as the proportion of weights falling within thresholds near zero (± 0.0001 , ± 0.00005 and ± 0.00001), and then we examined its correlation with connectivity dimensionality across regions.

Activity flow modeling

We implemented activity flow modeling^{21,22} to generate task-evoked activity via network flow. For each subject and target region, source regions were identified from the parcel-level FC graph (as in 'FC estimation'). Source region activity estimates (task conditions $\times$ source vertices) were projected onto each target region via matrix multiplication with the dense FC graph (source vertices $\times$ target vertices) in each session: each target vertex's predicted activity is the sum, over all n connected source vertices, of source activity weighted by source–target connectivity strength (that is, connectivity-specified activity flow from each source vertex). Model evaluation used Pearson's correlation between the flow-generated and empirically observed task condition activity matrices per target region, averaged over task conditions and sessions.

From model-generated activity estimates, we derived predicted RSMs and representational dimensionality, evaluated against observed values using cosine similarity (RSMs) and Pearson's correlation (dimensionality) per region. Prediction residuals (observed minus predicted) were examined to identify regions where additional factors beyond connectivity-based flow—for example, local circuit dynamics, neuromodulation and attention—may contribute.

Significance was estimated by projecting source activity over 100 iterations of randomly shuffled connectivity per subject, generating 1,000 group-level bootstrapped means⁶³. A more stringent

test constrained shuffling to the source side only—randomly permuting connectivity values among vertices within each source region—preserving each source’s spatial distribution properties and leaving target-vertex connectivity properties unaltered. Significance against this null demonstrates that specific source-to-target connectivity patterns, rather than general properties like target vertex degree distribution, drive predictive accuracy.

Double cross-validated RSMs for cortical dimensionality pattern comparison

To test whether model-predicted dimensionality patterns replicate the compression-then-expansion pattern along the sensorimotor FC gradient, we developed double cross-validated RSMs, cross-validating between observed and model-predicted estimates in addition to between sessions. Specifically, we computed cosine similarity between observed activity for condition A (session 1) and model-predicted activity for condition B (session 2), yielding a task conditions $\times$ task conditions RSM, and we repeated this for the reverse combination (predicted session 1 versus observed session 2). Representational dimensionality was computed from each and averaged within each region, providing a stringent test that reduces noise and session-specific artifacts.

RSM-based transformation index

We developed an RSM-based transformation index to quantify how much model-generated target representations differ from observed source representations. For each target region, we computed three cosine distances ($1 - \text{cosine similarity}$), averaged across source–target pairs: (a) $d_{\text{transform}}$, between observed source and observed target RSMs; (b) $\hat{d}_{\text{transform}}$, between observed source and predicted target RSMs; and (c) d_{pred} , between predicted and observed target RSMs (a reference measure of prediction error).

Cross-task representational similarity

To isolate cross-task from within-task similarity, we performed a one-sample t -test on each RSM cell across subjects (FDR corrected), yielding significance-survived RSMs of reliable task condition combinations. From these, we computed average positive similarity values among survived combinations (restricted to positive values, as negative cosine similarities are difficult to interpret), providing a measure of cross-task representational relationships.

To assess this relationship more generally, we simulated activity propagation across connectivity matrices with varying dimensionality. Source and target region sizes were sampled from uniform distributions of (100,150) and (50,100) vertices, respectively. For each source region, 100 random task activity patterns (standard normal distribution) were propagated through low-rank connectivity matrices (also standard normal) at rank factors of (0.01, 0.025, 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1.0) $\times$ target region size (10 source–target pairs per rank factor). Cross-task similarity was computed as the cosine similarity between resulting target task activity patterns, averaging positive values in the upper-triangular similarity matrix.

Representational similarity analysis in multiple tasks with shared features

We identified five tasks (44 task conditions) with especially interpretable shared features to test whether conjunctive feature coding relates to representational expansion. We constructed ground-truth similarity matrices for sensory, cognitive and motor features, and their multiplicative conjunction.

Sensory and motor features were coded from stimuli and expected responses in each task condition (Supplementary Table 1). Stimuli were categorized hierarchically (higher-order features: text, images, videos and spatial arrays; specific elements: scenes, faces and objects), yielding a graded stimulus RSM (1 = same stimulus type 1, for example, images; 2 = same stimulus type 2, for example, objects; 3 = identical

stimuli reused across conditions). Given practiced task accuracy (>85%) prior to scanning, we constructed a binary motor similarity matrix from expected button press responses (1 = same response, 0 = different).

Cognitive (task demand) similarity was computed by assigning demand levels to each task’s cognitive manipulation—most tasks had three levels (easy, medium and hard, coded 1–3), with Stroop coded 1 (congruent) or 3 (incongruent)—and computing pairwise similarity as

$$\text{sim}_{\text{cognitive}}(\text{cond}_A, \text{cond}_B) = 3 - |\text{td}(\text{cond}_A) - \text{td}(\text{cond}_B)|,$$

where td represents task demand, yielding values between one and three.

All three matrices were normalized to [0, 1] and multiplied to produce the conjunction term. In each region and subject, the observed 44×44 RSM (averaged across subjects for visualization) was fit to the four ground-truth matrices using ridge regression with 5-fold cross-validation.

HCP: neural and behavioral data analysis

We analyzed 352 unrelated, low-motion subjects from the HCP ‘1200 Subjects’ release, computing representational dimensionality from HCP task data^{29,30} using procedures analogous to our primary dataset. Following the same preprocessing pipeline, we obtained activity estimates for 24 task conditions across seven tasks via GLM-based multiple regression, with separate regressors per condition across two identical runs differing only in phase-encoding direction. Cross-validated RSMs (cosine similarity between runs) yielded dimensionality metrics, and a compression ratio quantified relative compression in association versus sensory/motor regions:

$$\text{CompressionRatio} = \frac{(D_{\text{Rep,Sensory}} + D_{\text{Rep,Motor}})}{2 \times D_{\text{Rep,Association}}},$$

where D_{Rep} represents the average representational dimensionality for each region type.

For behavioral analysis, we used a general fluid intelligence composite score⁶⁴, combining the z-scored fluid cognition composite (Dimensional Change Card Sort, Flanker, Picture Sequence Memory, List Sorting and Pattern Comparison^{65,66}) with z-scored Penn Matrix Reasoning Test performance, reflecting broad cognitive ability across working memory, cognitive control and reasoning⁴⁷.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The activation maps and functional parcellations from the original study are available at <http://www.diedrichsenlab.org/imaging/mdtb.htm>. The raw behavioral and imaging data for the cerebellum are also available on the data sharing repository at <https://openneuro.org/>. Example derivatives (representational and connectivity dimensionality values of all subjects) can be accessed at https://github.com/LakshmanChakravarthy/multitask_generalization/tree/master/example_derivatives. Source data are provided with this paper.

Code availability

The codebase can be accessed at https://github.com/LakshmanChakravarthy/multitask_generalization.

References

- King, M., Hernandez-Castillo, C. R., Poldrack, R. A., Ivry, R. B. & Diedrichsen, J. Functional boundaries in the human cerebellum revealed by a multi-domain task battery. *Nat. Neurosci.* **22**, 1371–1378 (2019).

49. Glasser, M. F. et al. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* **80**, 105–124 (2013).
50. Ji, J. L. et al. QuNex—An integrative platform for reproducible neuroimaging analytics. *Front. Neuroinform.* **17**, 1104508 (2023).
51. Rissman, J., Gazzaley, A. & D'Esposito, M. Measuring functional connectivity during distinct stages of a cognitive task. *NeuroImage* **23**, 752–763 (2004).
52. Friston, K. J. et al. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* **2**, 189–210 (1994).
53. Lindquist, M. A., Geuter, S., Wager, T. D. & Caffo, B. S. Modular preprocessing pipelines can reintroduce artifacts into fMRI data. *Hum. Brain Mapp.* **40**, 2358–2376 (2019).
54. Glasser, M. F. et al. A multi-modal parcellation of human cerebral cortex. *Nature* **536**, 171–178 (2016).
55. Schaipp, F., Müller, C. L. & Vlasovets, O. GGLasso—a Python package for General Graphical Lasso computation. Preprint at <https://doi.org/10.48550/arXiv.2110.10521> (2021).
56. Peterson, K. L., Sanchez-Romero, R., Mill, R. D. & Cole, M. W. Regularized partial correlation provides reliable functional connectivity estimates while correcting for widespread confounding. *Imaging Neurosci.* **3**, IMAG.a.162 (2025).
57. Mill, R. D. et al. Network modeling of dynamic brain interactions predicts emergence of neural information that supports human cognitive behavior. *PLOS Biol.* **20**, e3001686 (2022).
58. Gao, P. et al. A theory of multineuronal dimensionality, dynamics and measurement. Preprint at <https://doi.org/10.1101/214262> (2017).
59. Jeganathan, J. et al. Spurious correlations in surface-based functional brain imaging. *Imaging Neurosci.* **3**, imag_a_00478 (2025).
60. Burt, J. B., Helmer, M., Shinn, M., Anticevic, A. & Murray, J. D. Generative modeling of brain maps with spatial autocorrelation. *NeuroImage* **220**, 117038 (2020).
61. Coalson, T. S., Van Essen, D. C. & Glasser, M. F. The impact of traditional neuroimaging methods on the spatial localization of cortical areas. *Proc. Natl Acad. Sci. USA* **115**, E6356–E6365 (2018).
62. Margulies, D. S. et al. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc. Natl Acad. Sci. USA* **113**, 12574–12579 (2016).
63. Stelzer, J., Chen, Y. & Turner, R. Statistical inference and multiple testing correction in classification-based multi-voxel pattern analysis (MVPA): random permutations and cluster size control. *NeuroImage* **65**, 69–82 (2013).
64. Podvalny, E., Sanchez-Romero, R. & Cole, M. W. Functionality of arousal-regulating brain circuitry at rest predicts human cognitive abilities. *Cereb. Cortex* **34**, bhae192 (2024).
65. Scott, E. P., Sorrell, A. & Benitez, A. Psychometric properties of the NIH Toolbox Cognition Battery in healthy older adults: reliability, validity, and agreement with standard neuropsychological tests. *J. Int. Neuropsychol. Soc.* **25**, 857–867 (2019).
66. Heaton, R. K. et al. Reliability and validity of composite scores from the NIH Toolbox Cognition Battery in adults. *J. Int. Neuropsychol. Soc.* **20**, 588–598 (2014).

Acknowledgements

Multi-domain Task Battery data were provided by M. King, C. Hernandez-Castillo, R. A. Poldrack, J. Walters, R. B. Ivry and J. Diedrichsen via OpenNeuro (<https://doi.org/10.18112/openneuro.ds002105.v1.1.0>). HCP data were provided by the Human Connectome Project, WU-Minn Consortium (principal investigators: D. Van Essen and K. Ugurbil; 1U54MH091657), funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. We thank the Office of Advanced Research Computing at Rutgers, The State University of New Jersey, for providing access to the Amarel cluster and associated research computing resources that have contributed to the results reported here. We also thank the members of the Cole Neurocognition Lab for helpful conversations during the course of the research.

Author contributions

L.N.C.C., T.I. and M.W.C. conceptualized the project and wrote the paper. L.N.C.C. performed the formal analysis and visualization, developed software and wrote the original draft. A.T. collaborated in discussions of functional connectivity computation. M.W.C. acquired funding. M.W.C. (senior author) and T.I. supervised the project.

Funding

This work was supported by the US National Science Foundation (NSF) under award 2219323. The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

Competing interests

The authors declare no competing interests.

Additional information


Extended data is available for this paper at <https://doi.org/10.1038/s41593-026-02486-8>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41593-026-02486-8>.

Correspondence and requests for materials should be addressed to Lakshman N. C. Chakravarthula.


Peer review information *Nature Neuroscience* thanks Casey Paquola, Lucina Uddin and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.




Extended Data Fig. 1 | Connectivity dimensionality reflects similarity of connectivity patterns, not sparsity. **a)** Schematic of connectivity pattern similarity. Each target vertex (yellow and purple boxes on right) receives weighted inputs from all source vertices (red boxes on left), forming a connectivity pattern (vector of connection weights). Connectivity pattern similarity is measured by computing the pairwise Pearson correlation between all target vertices' connectivity patterns within each region, then averaging these correlations. **b)** Connectivity dimensionality is strongly determined by connectivity pattern similarity. Correlation between connectivity pattern similarity (mean pairwise correlation among target vertices' connectivity patterns) and connectivity dimensionality (D_{FC}) across all 360 regions ($r = -0.79$, $p < 0.001$, spatial-autocorrelation preserving permutation test). Each dot represents one region, colored by network. When target vertices have similar connectivity patterns (high correlation), dimensionality is low.

When target vertices have diverse connectivity patterns (low correlation), dimensionality is high. This demonstrates that connectivity dimensionality captures the geometric structure of how regions connect, not merely the number or strength of connections. **c-e)** Connection sparsity does not correlate with connectivity dimensionality. Principal component regression produces dense connectivity matrices (all weights non-zero). To test whether approximately sparse connectivity (many near-zero weights) explains dimensionality, we computed sparsity as the proportion of connection weights falling within small thresholds around zero: ± 0.0001 (panel **c**), ± 0.00005 (panel **d**), and ± 0.00001 (panel **e**). Across all three thresholds, sparsity showed no significant relationship with connectivity dimensionality (all $p > 0.05$). This confirms that connectivity dimensionality reflects the similarity structure among connectivity patterns rather than connection density.




Extended Data Fig. 2 | Connectivity dimensionality predicts compression/expansion within sensory, association, and motor systems. The relationship between connectivity dimensionality (D_{FC}) and representational compression/expansion (ΔD_{Rep}) holds consistently across all three major cortical divisions. **a)** Sensory regions (visual networks VIS1 and VIS2; $n = 59$ regions): $r = 0.76$, $p < 0.001$. **b)** Association regions (default-mode, frontoparietal, dorsal attention, language, cingulo-opercular, auditory, and somatomotor-parietal networks;

$n = 250$ regions): $r = 0.44$, $p < 0.001$. **c)** Motor regions (somatomotor-frontal network SMN-Fr; $n = 51$ regions): $r = 0.69$, $p = 0.0012$. Each dot represents one region, colored by network identity within each system. Gray shading shows 95% confidence interval of the linear fit. These within-system correlations demonstrate that connectivity geometry predicts representational transformations as a general principle operating across functionally distinct cortical systems, not merely as a whole-brain statistical artifact.



Extended Data Fig. 3 | Relationship between connectivity dimensionality and representational compression/expansion remains significant after controlling for potential confounds. a-f) Spatial distribution of potential confounding variables across cortex. Brain maps showing six region-specific properties varying across cortex that could potentially confound connectivity dimensionality, representational dimensionality change, or both. **a) Region size** (vertex count). **b) Inter-vertex distance** (mean geodesic distance on cortical surface). **c) Betas reliability** (mean task condition reliability, estimated by taking the mean of the RSM diagonal). **d) Task Condition Signal-to-Noise Ratio, tcSNR** (mean/SD of task condition betas at each vertex, averaged across vertices in each region). **e) Temporal Signal-to-Noise Ratio, tSNR** (mean/SD of rest time series at each vertex, averaged across vertices in each region). **f) Spatial variance** (standard deviation of rest time series across vertices at each timepoint, averaged across time). **g-j) Individual confound relationships with compression/expansion.** Each panel shows correlation between one confound and ΔD_{Rep} across all 360 regions. Region size ($r = 0.42$), betas reliability


($r = 0.51$), and task condition SNR ($r = 0.41$) show significant relationships; inter-vertex distance shows no relationship ($r = -0.01$). **k-n) Individual confound relationships with connectivity dimensionality.** Each panel shows correlation between one confound and D_{FC} across all 360 regions. Region size shows the strongest relationship with D_{FC} ($r = 0.75$), while inter-vertex distance ($r = 0.15$), temporal SNR ($r = -0.17$), and spatial SD ($r = 0.18$) show smaller but significant relationships. These correlations demonstrate that several region-specific properties do vary with compression/expansion and connectivity dimensionality, necessitating control analyses. **o) Central result remains significant after controlling for confounds.** To test whether connectivity dimensionality independently predicts compression/expansion, we regressed the respective confounds from ΔD_{Rep} and D_{FC} and correlated the residuals, and the resultant correlation is still significant ($r = 0.4357$, $p < 0.001$). These results demonstrate that connectivity dimensionality predicts representational transformations independent of region size, inter-vertex distance, beta reliability, task condition SNR, temporal SNR, and spatial variance.



Extended Data Fig. 4 | Non-linear local computations can decouple connectivity dimensionality from representational dimensionality change.


To test whether this relationship is mathematically inevitable, we simulated source and target regions (100 vertices each) with varying connectivity dimensionality. Source activity patterns for 100 tasks were generated from a standard normal distribution and projected through connectivity matrices with dimensionality controlled by rank (participation ratio ranging from 10–100). **a)** Linear projection preserves the relationship. When source activity is linearly projected through the connectivity matrix and representational dimensionality is computed directly from target activity, the relationship between connectivity dimensionality and change in representational dimensionality (target minus source) is extremely strong ($r = 0.981, p < 0.001$). Each dot represents one simulated source-target pair (10 pairs per dimensionality level, across 9 levels). Lower connectivity dimensionality produces compression (negative Δ representational dimensionality); higher dimensionality produces less compression or expansion, demonstrating that connectivity dimensionality strongly constrains representational dimensionality under purely linear transformations. **b)** Non-linear local computations eliminate the relationship. To simulate nonlinear transformation at the target region, connectivity-projected activity additionally underwent a leaky ReLU non-linearity ($\alpha = 0.01$) with region-specific and task-dependent bias terms before representational

dimensionality was computed: target vertex activity $y = f(T \cdot W + b_{\text{region}} + b_{\text{task}})$, where source task activity $T \sim N(0, 1)$, f is leaky ReLU with $\alpha = 0.01$, W is the trainable parameter, but in this case are connectivity weights with set dimensionality, bias pattern that simulates idiosyncratic local computations in the target region, where each vertex has a baseline bias drawn from a Gaussian distribution with a large mean for demonstration, $b_{\text{region}} \sim N(400, 50)$, and task-specific biases with smaller variance on top, $b_{\text{task}} \sim N(0, 25)$. When simulated this way, the relationship between compression and connectivity dimensionality disappears ($r = 0.046, p = 0.650$). This exaggerated scenario demonstrates that strong non-linear local operations can override connectivity-based constraints on representational dimensionality, and that such relationship is not mathematically inevitable. The empirical observation that connectivity dimensionality robustly predicts representational transformations across cortex (Fig. 2f, $r = 0.44$) therefore suggests that while local computations certainly contribute, connectivity geometry provides systematic constraints that shape representational transformations throughout the brain. Note that this linear projection assumption applies to the activity flow model, not to the primary empirical result; it does not imply that neuronal transformations are themselves linear, but rather that linear projection approximates their net effect at the fMRI spatial scale.



Extended Data Fig. 5 | Representational compression corresponds to reduced task decodability across cortical networks. *Motivation:* BOLD signal in association cortex, where mixed-selective neurons are spatially intermingled, may poorly reflect underlying representational diversity, consistent with prior reports of difficulty decoding task information from prefrontal BOLD⁴⁷. **a)** Task decodability varies across cortical networks. Decodability of 16 individual tasks from representational similarity matrices (2-fold minimum distance classification based on cosine distance from task-averaged RSMs; chance = $1/16 = 0.0625$, dashed line). Each dot represents one region; box plots show distribution across regions within each network arranged along the sensory-motor gradient. Visual networks (VIS1, VIS2) and motor regions (SMN-Fr) show highest task decodability. Default-mode (DMN) and auditory (AUD) networks show lowest decodability, falling near or below the median of other networks. **b)** Compression/expansion pattern mirrors decodability. Representational compression/expansion (ΔD_{Rep}) plotted by network along the sensory-motor


gradient. Default-mode and auditory networks show strongest compression (most negative ΔD_{Rep}), while visual and motor networks show expansion or minimal compression. This pattern closely parallels the decodability results in panel A, demonstrating that compression reflects genuine loss of task discriminability. **c)** Task decodability is robust throughout cortex. Brain maps showing decodability of 16 tasks (left; range 0.25-1.0) and 96 task conditions (right; range 0.03-0.3). Despite the compression observed in association cortex, task decodability remains well above chance (0.0625) even in default-mode regions, indicating that BOLD signals in association cortex meaningfully reflect task information rather than averaging over mixed selectivity to the point of indiscriminability. **d)** Mean task decodability strongly correlates with compression/expansion. Correlation between mean task decodability and ΔD_{Rep} across all 360 regions ($r = 0.85$, $p < 0.001$). Each dot represents one region, colored by network. Regions showing expansion (positive ΔD_{Rep}) have higher task decodability; regions showing compression have lower decodability.



Extended Data Fig. 6 | Activity flow over connectivity generates genuine representational transformations, not simple transfer. a) Measuring transformation distance. To test whether activity flow implements true representational transformations rather than simply transferring source representations to targets, we computed three distances using cosine distance ($1 - \text{cosine similarity}$) between RSMs. $\hat{d}_{\text{transform}}$: distance between observed source RSM and predicted target RSM, quantifying how much the model transforms representations. $d_{\text{transform}}$: distance between observed source RSM and observed target RSM, quantifying the actual transformation. d_{pred} : distance between predicted target RSM and observed target RSM, quantifying prediction error. **b)** Predicted transformation distance correlates with observed transformation distance. Correlation between $\hat{d}_{\text{transform}}$ (model-predicted source-to-target transformation) and $d_{\text{transform}}$ (observed source-to-target transformation) across all 360 regions ($r = 0.62$, $p < 0.001$, spatial-autocorrelation preserving permutation test). Each dot represents one region, colored by network. The model accurately captures the magnitude of representational transformations


across cortex. **c)** Connectivity-based transformations differ significantly from null models. Predicted transformation distances ($\hat{d}_{\text{transform}}$, red) plotted for each region sorted by magnitude, compared against null distributions generated by shuffled connectivity (gray; 1000 iterations). All 360 regions show significantly lower transformation distances with intact connectivity than with shuffled connectivity (all $p < 0.001$), demonstrating that the specific structure of empirical connectivity patterns drives transformations. **d)** Predictions resemble targets more than sources in most regions. Each dot represents one region, plotting prediction distance (d_{pred} , similarity to target) against predicted transformation distance ($\hat{d}_{\text{transform}}$, similarity to source). Points below the diagonal (339/360 regions, $\text{pFDR} < 0.05$) indicate the predicted RSM is closer to the observed target than to the observed source, demonstrating genuine transformation rather than simple transfer. This holds for both compression zones (association cortex) and expansion zones (motor cortex), showing that connectivity geometry shapes transformations bidirectionally.

Observed vs. predicted cross-task similarity pattern comparison (only FC shuffling permutation survived)



Extended Data Fig. 7 | Cross-task similarity predictions at the level of individual regions and large-scale brain networks. (a) Cosine similarity between vectorized cross-task similarity values from observed (empirically-estimated) and from connectivity-predicted estimates (range 0.13–0.99; $p_{\text{connperm}} < 0.05$ in 355/360 regions). Regions with no values did not survive FC shuffling permutation test. (b) Similar comparison of the network-level cross-task similarity patterns. Grey dots indicate null distributions derived from FC

shuffling. Network abbreviations from Cole-Anticevic Brain Network Partition: VIS1: primary visual network, VIS2: secondary (higher-order) visual network, SMN: somato-motor network, CON: cingulo-opercular network, DAN: dorsal attention network, LAN: language network, FPN: fronto-parietal network, AUD: auditory network, DMN: default-mode network, PMULTI: posterior multimodal network, VMM: ventral multimodal network, ORA: orbito-affective network.



Extended Data Fig. 8 | Representational dimensionality is predictive of behavioral measures of cognitive flexibility. **a)** Average representational dimensionality across regions was positively correlated with cognitive flexibility scores across subjects. **b)** Similar relationship in individual regions (following multiple comparison corrections) was reflected in positive correlations in different regions. **c)** Dimensionality values averaged across sensory, association and motor regions were predictive of flexibility scores in a cross-validated (leave one subject out) manner. **d)** Coefficients from the cross-validation approach show higher values for sensory and association regions' dimensionality values compared to motor regions in explaining the flexibility

scores. **e)** The representational dimensionality values in the new dataset also follow compression-then-expansion pattern (average values across subjects). **f)** Relative compression in association regions compared to sensory and motor regions was reflected in almost all subjects, measured using compression ratio. These findings address an apparent paradox: although association regions show representational compression on average (Fig. 2), individual differences in absolute dimensionality within this organization—particularly in sensory and association regions—predict fluid intelligence, suggesting the compression-then-expansion pattern reflects canonical task organization while absolute dimensionality reflects individual capacity for flexible cognition.

Extended Data Table 1 | Cross-task similarity comparison between tasks and networks

Data Type	Task Comparison	Category	Mean VIS1	Mean FPN	SD VIS1	SD FPN	t Statistic	p-value	p (Bonferroni)	Cohen's d
Observed	Object N-back – Spatial Map	Shared Function (WM)	0.2808	0.3347	0.0488	0.0476	-10.4621	7.95×10^{-9}	4.77×10^{-8}	-2.5374
Observed	Verbal N-back – Spatial Map	Shared Function (WM)	0.2212	0.3177	0.0530	0.0558	-9.5790	2.90×10^{-8}	1.74×10^{-7}	-2.3233
Observed	Object N-back – Verbal N-back	Shared Function (WM)	0.2773	0.3632	0.0618	0.0706	-6.1485	1.07×10^{-5}	6.43×10^{-5}	-1.4912
Observed	Go No-Go – Arithmetic	Shared Stimulus (Word)	0.4427	0.2807	0.0845	0.0498	11.9830	1.03×10^{-9}	6.17×10^{-9}	2.9063
Observed	Go No-Go – Word Prediction	Shared Stimulus (Word)	0.4819	0.2761	0.1036	0.0428	9.5515	3.02×10^{-8}	1.81×10^{-7}	2.3166
Observed	Word Prediction – Arithmetic	Shared Stimulus (Word)	0.4119	0.2482	0.0868	0.0438	9.4209	3.68×10^{-8}	2.21×10^{-7}	2.2849
Model-Predicted	Verbal N-back – Spatial Map	Shared Function (WM)	0.4612	0.5918	0.0780	0.0755	-12.1001	8.86×10^{-10}	5.32×10^{-9}	-2.9347
Model-Predicted	Object N-back – Spatial Map	Shared Function (WM)	0.5649	0.6363	0.0630	0.0580	-5.7340	2.43×10^{-5}	1.46×10^{-4}	-1.3907
Model-Predicted	Object N-back – Verbal N-back	Shared Function (WM)	0.5867	0.6634	0.0767	0.0719	-4.4696	3.37×10^{-4}	2.02×10^{-3}	-1.0840
Model-Predicted	Go No-Go – Word Prediction	Shared Stimulus (Word)	0.7982	0.4514	0.0766	0.0625	17.2661	3.26×10^{-12}	1.95×10^{-11}	4.1876
Model-Predicted	Go No-Go – Arithmetic	Shared Stimulus (Word)	0.7620	0.4592	0.0637	0.0489	16.1177	9.85×10^{-12}	5.91×10^{-11}	3.9091
Model-Predicted	Word Prediction – Arithmetic	Shared Stimulus (Word)	0.7323	0.4165	0.0889	0.0556	14.6785	4.36×10^{-11}	2.62×10^{-10}	3.5600

Cross-task similarity comparison between tasks and networks.

Extended Data Table 2 | Correlations between representational expansion and RSA template term coefficients, observed and predicted

Term	Observed r	Observed p	vs. conjunction (paired t)	Predicted r	Predicted p
Conjunction	0.46	2.92e-07	—	0.34	0.006
Sensory	0.47	8.41e-10	t(17)=-1.94, p=0.07 (n.s.)	0.16	0.002
Cognitive	0.33	0.013	t(17)=2.74, p=0.014	0.13	0.636
Motor	-0.19	0.002	t(17)=6.56, p=4.86e-06	-0.15	0.039

Anatomical specificity of the four RSA template terms: 9/10 highest sensory coefficients in sensory regions, 7/10 highest cognitive coefficients in association regions, 7/10 highest motor coefficients in motor regions (Fig. 5b). Conjunction-term coefficients correlated with representational expansion (ΔD_{rep}) more strongly than cognitive or motor terms alone, and comparably to the sensory term.

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- ☐ ☒ The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- ☐ ☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- ☐ ☒ The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- ☐ ☒ A description of all covariates tested
- ☐ ☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- ☐ ☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- ☐ ☒ For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- ☒ ☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- ☒ ☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- ☐ ☒ Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection	Data was downloaded from public repositories (OpenNeuro, Human Connectome Project). No software was used for data collection in this study.
Data analysis	QuNex (version 0.61.17) was used to preprocess the raw data. Workbench (version 1.5.0) was used to visualize cortical maps. Custom python code (version 3.8.5) was used for analysis, and will be made publicly available prior to publication. We used NumPy (version 1.18.5) and SciPy (version 1.6.0).

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

Multi-domain task battery data was downloaded from <https://openneuro.org/datasets/ds002105/versions/1.1.0/1.1.0>. Human Connectome project was downloaded from <https://www.humanconnectome.org/study/hcp-young-adult/data-releases>

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender	We report using biological sex, as reported in the original data resource article (King et al., 2019, Nature Neuroscience).
Reporting on race, ethnicity, or other socially relevant groupings	No data was reported in the original data resource article
Population characteristics	From the original Data Resource article (King et al., 2019, Nature Neuroscience): "The final sample consisted of 24 healthy, right-handed individuals (16 females, 8 males; mean age=23.8 years old, SD=2.6) with no self-reported history of neurological or psychiatric illness."
Recruitment	From the original Data Resource article (King et al., 2019, Nature Neuroscience): "Undergraduate and graduate students were recruited (via posters) from the larger student body at Western University. Thus, our sample was biased towards relatively high-functioning, healthy and young individuals. While we don't expect cerebellar organization to be dramatically different in this group, caution needs to be exercised when generalizing the results to the general population."
Ethics oversight	From the original Data Resource article (King et al., 2019, Nature Neuroscience): "The Ethics committee at Western University approved all experimental protocols (Protocol number: 107293)"

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	Sample size was determined by the previous Data Resource article (King et al. 2019, Nature Neuroscience): "No formal power analysis was conducted, but the number of participants and the amount of data per participant was determined from prior experience. For the withinsubject sample size, we required 5.5 hours of fMRI data for each participant. This gave us enough power to reliably estimate activity patterns for the large number of tasks, as well as having enough data to perform out-of-sample prediction tests. A sample size of n=24 participants (after exclusions) was deemed sufficient to both estimate a representative mean organization, as well as obtaining reliable estimates of the between-subject variability." The present study analyzed 18 out of 24 subjects, also performed group-level analyses, ensuring similar statistical power to the original study.
Data exclusions	Data from 6 subjects was not included in analysis, as their resting-state data was not provided, which is critical for computing functional connectivity.
Replication	While no direct replication was performed with an external data set, cross-validation (within subjects) was performed across two independent recording sessions per participant to ensure robustness and reliability of results.
Randomization	N/A -- participants were not allocated into separate groups.
Blinding	No blinding was performed, since between-group differences were not assessed.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☐	☒ MRI-based neuroimaging

Plants

Seed stocks	<i>Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.</i>
Novel plant genotypes	<i>Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.</i>
Authentication	<i>Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosaicism, off-target gene editing) were examined.</i>

Magnetic resonance imaging

Experimental design

Design type	Task- and resting-state.
Design specifications	From the original Data Resource article (King et al., 2019, Nature Neuroscience): "Two task sets. 2 fMRI scanning sessions per task set. 8 functional imaging runs per session (10-min each). 17 tasks per imaging run (35 s each)."
Behavioral performance measures	From the original Data Resource article (King et al., 2019, Nature Neuroscience): "Variables recorded: response made, number of correct responses, false alarms, missed responses, response time. Accuracy (% correct) and reaction time (ms) were collected and averaged across tasks per participant"

Acquisition

Imaging type(s)	From the original Data Resource article (King et al., 2019, Nature Neuroscience): "EPI, MPAGE, and GRE field maps"
Field strength	3T
Sequence & imaging parameters	From the original Data Resource article (King et al., 2019, Nature Neuroscience): "EPI: Gradient echo, multi-band (factor 3, interleaved) with an in-plane acceleration (factor 2). Imaging parameters were: TR=1 sec, FOV=20.8cm, phase encoding direction was P to A, acquiring 48 slices with in-plane resolution of 2.5 mm x 2.5 mm and 3 mm thickness. For anatomical localization and normalization, a 5-min high-resolution scan of the whole brain was acquired (MPAGE, FOV=15.6 cm x 24 cm x 24 cm, at 1x1x1 mm voxel size)."
Area of acquisition	Whole-brain
Diffusion MRI	☐ Used ☒ Not used

Preprocessing

Preprocessing software	Resting-state and task-state fMRI data were minimally preprocessed using the Human Connectome Project (HCP) preprocessing pipeline within the Quantitative Neuroimaging Environment & Toolbox (QuNex, version 0.61.17). The HCP
------------------------	---

	preprocessing pipeline consisted of anatomical reconstruction and segmentation, EPI reconstruction and segmentation, spatial normalization to the MNI152 template, and motion correction.
Normalization	Nonlinear spatial normalization was performed to the MNI 152 template using QuNex.
Normalization template	MNI152
Noise and artifact removal	Additional nuisance regression was performed on the minimally preprocessed time series. Consistent with previous reports, this included six motion parameters, their derivatives, and the quadratics of those parameters (24 motion regressors in total). We also removed the mean physiological time series extracted from the white matter and ventricle voxels. We also included the quadratic, derivatives, and the derivatives of the quadratic time series of each of the white matter and ventricle time series (8 physiological nuisance signals). This amounted to 32 nuisance parameters in total, and was a nuisance regression model that was previously benchmarked (Ciric et al. 2017).
Volume censoring	We excluded the first five volumes of each run

Statistical modeling & inference

Model type and settings	We used univariate GLMs to estimate task activations, and multivariate pattern analyses (RSA and decoding) analyses at the vertex-level. Additional simulations include demonstrating effect of cross-task similarity based on connectivity dimensionality, and non-linearity introduction at target region to demonstrate possible activity generation at target region whose representational geometry does not relate to connectivity properties involved in its generation.
Effect(s) tested	Mean differences were tested for: Representational and connectivity dimensionality, associations (Pearson correlation) between them, activity flow model predictions of activation patterns, representational geometry, and transformation metrics; contrasting similarity between tasks sharing sensory and cognitive features in respective regions, association between feature conjunction strength and task dimensionality In simulations, relationship between change of target representational dimensionality and connectivity dimensionality under linear and non-linear activity projection of source activity
Specify type of analysis:	☒ Whole brain ☐ ROI-based ☐ Both
Statistic type for inference (See Eklund et al. 2016)	Whole-brain analyses and inferences were made at the parcel-level using pre-defined cortical parcels (defined in Glasser et al., 2016, Nature).
Correction	Since inferences were made at the whole-cortex level (or between systems across cortex), multiple comparisons correction was not necessary.

Models & analysis

n/a	Involved in the study
☐	☒ Functional and/or effective connectivity
☐	☒ Graph analysis
☐	☒ Multivariate modeling or predictive analysis
Functional and/or effective connectivity	Regularized partial correlation was used for region-level functional connectivity scaffold. Among the areal relationships that survived, fine-grained vertex-level functional connectivity was estimated using cross-validated principal component regression
Graph analysis	Weighted matrices thresholded at 20% were used to estimate principal gradient loadings, consistent with previous reports (Margulies et al., 2016, PNAS). Connectivity dimensionality was measured by performing singular value decomposition on connectivity matrices between each region and its source regions, and the decay in singular value spectrum was computed using participation ratio (Gao et al. 2017)
Multivariate modeling and predictive analysis	Multivariate modeling was performed using multivariate pattern decoding on vertices within each cortical parcel. Additional estimates of multivariate analyses were included, such as representational dimensionality, which measures the participation ratio of the cross-validated representational similarity matrix. Representational similarity matrices were constructed using cross-validated cosine similarity across imaging sessions (within subjects).