Title: Network geometry shapes multi-task representational transformations across human cortex

Authors:
Lakshman N.C. Chakravarthula[1], Takuya Ito[2], Alexandros Tzalavras[1], Michael W. Cole[1]
[1]Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ
[2]T.J. Watson Research Center, IBM Research, Yorktown, NY

Classification:
Major: Biological Sciences
Minor: Neuroscience

Keywords:
Representational geometry; Functional connectivity; Activity flow modeling; Neural dimensionality; Cognitive flexibility; Task representations; Connectivity dimensionality; Conjunctive coding; Cross-task similarity; Cortical hierarchy

Correspondence:
Correspondence should be addressed to Lakshman N.C. Chakravarthula at lakastro@gmail.com

Data sharing plans:
Primary data is already shared on OpenNeuro database, to which any additional secondary data can be added. Analysis code will be shared on Github.

Competing interest information:
The authors declare no competing interests.

Funding information:
This work was supported by the US National Science Foundation (NSF) under award 2219323.
No open access license has been selected yet.

Abstract

How the human brain performs such a wide variety of possible tasks remains an enigma. Recent work revealed that reduced dimensionality of task representations enables generalization across contexts, yet the mechanisms implementing such compression remain unclear. We hypothesized that network architecture constrains representational compression via low-dimensional region-to-region connectivity mappings. Notably, this contrasts with the common assumption that between-region connections passively carry information. We tested this using fMRI data from participants performing 16 diverse tasks and activity flow modeling, a data-driven neural network modeling approach. As hypothesized, connectivity dimensionality predicted cross-region representational changes: low-dimensional connectivity compressed task representations, while high-dimensional connectivity expanded them. Activity flow modeling demonstrated these connectivity patterns transform rather than passively transfer representations, producing compression-then-expansion from sensory through association to motor regions, enabling cross-task generalization and task-specific implementation. These findings reveal how intrinsic network geometry specifies representational transformations, providing systems-level insights into how network geometry enables flexible cognition.

Main Text

Introduction

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 states [1]. 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, then expanded [6]. Concretely, different task conditions may produce distinct high-dimensional patterns in visual cortex reflecting their unique sensory features (colors, shapes, words). In association cortex, representations compress—conditions that share abstract task-relevant features (e.g., 'respond left' or 'inhibit response') become more similar despite differing sensory details. In 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 transferrable 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 task-specific performance quality[7]. 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, 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, while

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.

Prior studies have shown that different regions prioritize encoding different aspects of task conditions: some studies have observed regions representing both task-relevant and irrelevant variables [8], while others observed coding schemes supporting selective collapse of task-irrelevant information [4,9,10] or maintaining abstract representations [11–13]. Other studies point toward representations encoding conjunctive combinations of variables in a non-linear mixed-selective manner [14–16]. Consequently, some brain regions exhibit low-dimensional representations, while others represent task conditions in high-dimensional spaces [17]. 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. While 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. Prior 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 selectivity[23]. 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 employs 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 hypothesized that the dimensionality of intrinsic (resting-state) connectivity patterns systematically shapes the dimensionality of task representations in target regions. These are fundamentally different measures: connectivity dimensionality reflects the geometric structure of how regions connect during a single brain state (rest), while 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.

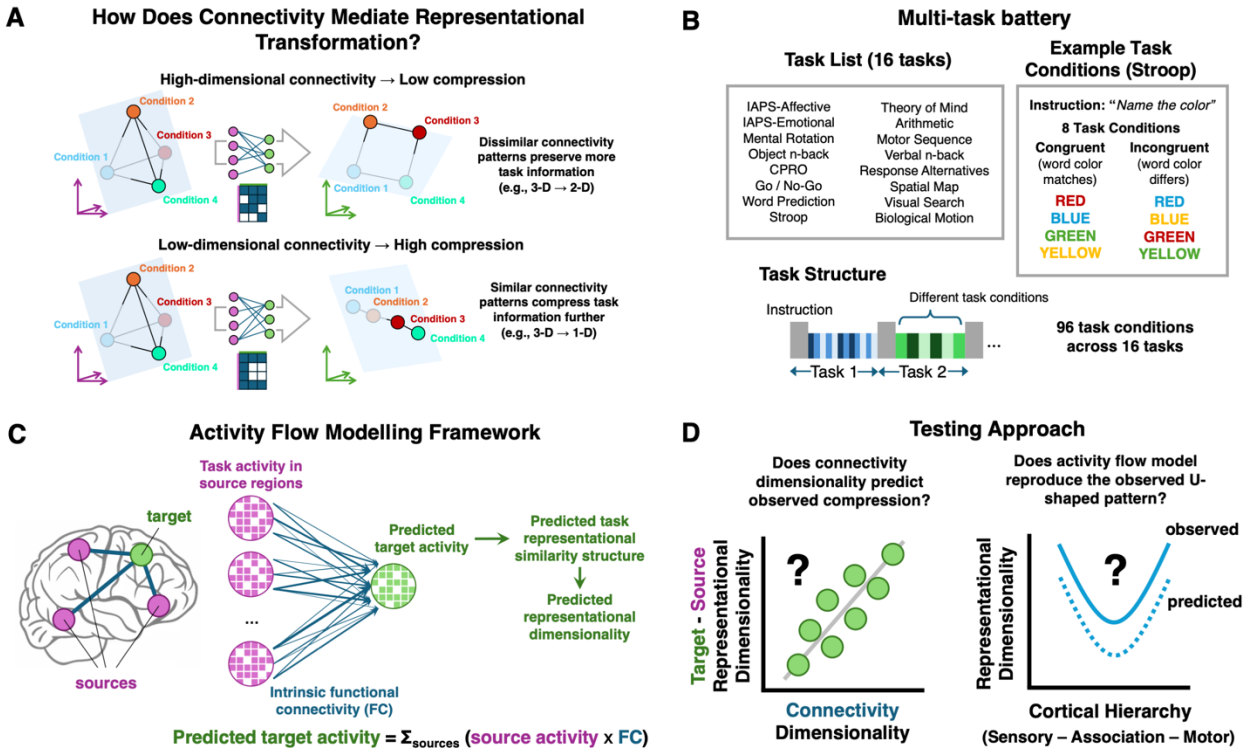

Figure 1: Testing how connectivity geometry shapes representational transformations across cortex. A) Hypothesis: Connectivity dimensionality determines compression versus expansion of task representations. Task conditions (colored spheres labeled 1-4) occupy positions in neural activity space. Representational 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 vs 3-4; 2-3 vs 1-4; 2-4 vs 1-3); in 2D space, two pairwise splits remain possible (1-2 vs 3-4; 2-3 vs 1-4); in 1D space, only one pairwise split is possible (1-2 vs 3-4). *Top:* High-dimensional connectivity (dissimilar connectivity patterns among target vertices, shown in 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 8 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-second instruction period followed by 30 seconds of task performance. Across all 16 tasks, we estimated activity patterns for 96 distinct task conditions using regularized GLM, providing rich representational diversity for measuring dimensionality. **C) Activity flow modeling framework.** Task activity in source regions (all functionally connected regions identified via sparse graphical lasso; see 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 = $\Sigma_{sources}$ (source activity × 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 (see 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; see 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. (End of Figure 1 caption)

Here we examine how connectivity dimensionality shapes inter-areal representational transformations during diverse task performance. Using fMRI data from participants performing 16 different tasks, 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 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, while 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 non-linear local computations can fully decouple connectivity dimensionality from representational transformations (Supplementary Fig. S4). 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, while 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 strengths[25]. 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 FC estimation (see 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 components 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 many dimensions, while 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 >> 96). However, if some conditions share similar patterns (e.g., tasks requiring similar responses), they can be represented in a lower-dimensional subspace. As in prior work[6], 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: $p_{perm} < 0.001$), 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 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.

To ground this abstract measure in interpretable connectivity properties, we examined what drives variation in connectivity dimensionality across regions – whether it is similar connectitivty patterns or sparsity. Connectivity dimensionality reflects the similarity structure of how source regions connect to target vertices: when target vertices receive similar input patterns from source regions, dimensionality is low; when they receive diverse patterns, dimensionality is high. We confirmed this by computing the mean pairwise correlation among connectivity patterns of target vertices (where each pattern is a vector of connection weights from all source vertices to that target vertex; Supplementary Fig. S1A). This correlation was strongly negatively related to connectivity dimensionality ($r = -0.79$, SA-preserving permutation test: $p_{perm} < 0.001$; Supplementary Fig. S1B), while connection sparsity (proportion of near-zero weights) showed no relationship ($p > 0.05$; Supplementary Fig. S1C-E), confirming that connectivity dimensionality captures geometric structure rather than the density of connections.

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 prior work[6], we plotted transformations along an FC gradient – a continuous axis derived from the principal structure of the whole-brain connectivity matrix, which runs from 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 trend of compression-then-expansion in task representations along the FC gradient (Fig. 2B,D). As targets, sensory regions typically showed expansion, association regions exhibited compression, and motor regions expansion, forming a U-shaped trajectory of representational compression/expansion. This pattern can be understood through a concrete example: In visual cortex, different task conditions produce distinct high-dimensional patterns reflecting their unique sensory features (colors, shapes, words, spatial locations). In association regions, representations compress—conditions that share abstract task-relevant features (e.g., 'respond left' or 'inhibit response' or 'remember this stimulus') become more similar despite differing sensory details. In 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. To assess the robustness of the U-shaped patterns in Fig. 2B-D, we performed model comparisons between quadratic and linear fits. For representational dimensionality change (Fig. 2D), the quadratic model was favored (ΔAIC = 3.94), with the quadratic coefficient reaching significance (p_{perm} = 0.001 using spatial-autocorrelation preserving permutation tests). For connectivity dimensionality (Fig. 2C), model comparison actually favored a linear fit (ΔAIC = -9.76), though visual inspection shows curvature driven by regions at gradient extremes. We used bootstrap resampling (1000 iterations) to generate 95% confidence intervals for both fits, showing increased uncertainty at gradient extremes where fewer regions are located.

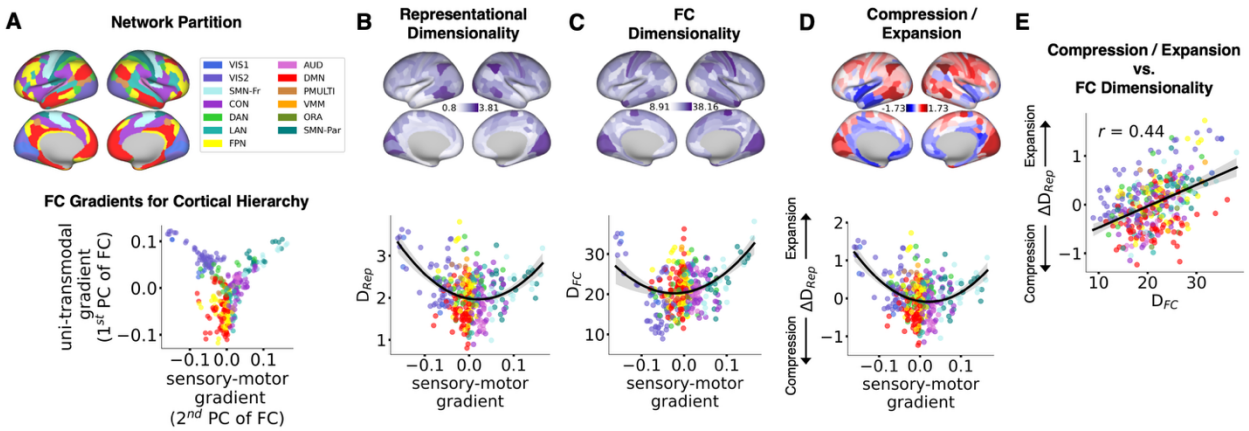

Figure 2: **Connectivity dimensionality predicts representational compression and expansion across cortex. A) Cortical organization characterized by discrete networks and continuous gradients.** *Top:* Cortical regions colored by functional network identity according to the Cole-Anticevic Brain Network Partition[27], dividing cortex into 12 networks including visual (VIS1, VIS2), somatomotor

(SMN-Fr: frontal/motor, SMN-Par: parietal/somatosensory), cingulo-opercular (CON), dorsal attention (DAN), language (LAN), frontoparietal (FPN), auditory (AUD), default-mode (DMN), posterior-multimodal (PMULTI), ventral-multimodal (VMULTI), and orbito-affective (ORA) networks. *Bottom:* Continuous FC gradients derived from principal components analysis of the group-averaged, thresholded parcel-level connectivity graph. Each dot represents one of 360 cortical regions, colored by network identity. The first principal component (y-axis) captures a uni-transmodal gradient separating sensory/motor regions from association regions. The second principal component (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 (2nd PC) for visualization throughout. **B) Representational dimensionality shows compression-then-expansion along the cortical hierarchy.** *Top:* Brain maps showing representational dimensionality (DRep, 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, while association regions (center) show compressed representations, forming a U-shaped pattern (black curve: quadratic fit with 95% bootstrap CI; ΔAIC = 3.94 favoring quadratic over linear). **C) Connectivity dimensionality mirrors the U-shaped pattern.** *Top:* Brain maps showing connectivity dimensionality (DFC, 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. While model comparison favors a linear fit (ΔAIC = -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 (ΔDRep = target dimensionality minus average source dimensionality). Negative values (blue, compression) indicate 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. **E) Connectivity dimensionality predicts compression and expansion across all regions.** Direct correlation between connectivity dimensionality and representational compression/expansion across all 360 cortical regions ($r = 0.44$, $p < 0.001$ using spatial-autocorrelation preserving permutation tests; relationship remains significant $r = 0.29$ 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. Higher connectivity dimensionality (diverse connectivity patterns among target vertices) predicts representational expansion; lower connectivity dimensionality (similar patterns) predicts compression. (End of Figure 2 caption)

To assess the generalizability of this compression-then-expansion pattern, we analyzed an independent dataset from the Human Connectome Project[28,29] (HCP; n=352). 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 (Supplementary Fig. S8E). 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 association cortices. This pattern proved highly reliable, with 98.6% (347/352) of subjects showing compression in association regions, providing strong validation of this organizational principle (Supplementary Fig. S8F). Consistent with our hypothesis, connectivity dimensionality mirrored this trend with a similar U-shaped trajectory. These observations revealed a significant positive relationship between compression and connectivity dimensionality ($r = 0.44$;

group t-test against zero: $t(17) = 16.94$, $p = 4.44e-12$; SA-preserving permutation test: $p_{perm} < 0.001$; Fig. 2E). This positive relationship was consistent across sensory, association, and motor systems (all $p < 0.001$, Supplementary Fig. S2), 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. In 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 non-linear transformations in the target region (e.g., 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 control for potential confounds, we tested whether regional properties that vary across cortex could explain both connectivity and representational dimensionality. We measured region size (vertex count), signal-to-noise ratio (mean task condition reliability), mean inter-vertex distance (surface mesh geometry), temporal variance of resting-state activity (standard deviation of time series averaged across vertices) (Supplementary Fig. S3 A-J. See supplementary figure for spatial homogeneity of activity – standard deviation across vertices averaged across time – as another confound). After regressing these confounds from connectivity dimensionality, the relationship with representational compression/expansion remained significant ($r = 0.29$, SA-preserving permutation test: $p_{perm} < 0.001$, Supplementary Fig. S3 K-L). To directly test this, we constructed a simulation showing that strong non-linear local computations in target regions can decouple the relationship between connectivity dimensionality and representational change (Supplementary Fig. S4), 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 for the above findings is that BOLD signals in association cortex may poorly reflect the diversity of underlying neural representations. In association regions, neurons with heterogeneous, mixed-selective responses are spatially intermingled, so the BOLD signal—which integrates activity across many neurons within a voxel—can average over distinct responses and reduce condition-discriminability. Prior work has indeed reported difficulty decoding task information from prefrontal cortex BOLD signals[30], raising the question of whether our representational dimensionality estimates in association regions are meaningfully constrained by this integration problem. However, we found robust decodability of both task conditions (96 conditions, minimum distance classification from RSMs: accuracy range 0.03-0.3 across regions) and individual tasks (16 tasks, accuracy range 0.25-1.0) throughout cortex (Supplementary Fig. S5C). Importantly, frontoparietal regions showed task decodability comparable to sensory and motor regions, with the lowest decodability observed in default-mode and auditory networks (Supplementary Fig. S5A). When examining representational compression by functional network rather than broad anatomical categories, compression was concentrated in default-mode and

auditory regions, while lateral frontoparietal regions maintained dimensionality values comparable to sensory and motor regions (Supplementary Fig. S5B). This pattern indicates that the compression we observe reflects specific functional network organization rather than a uniform property of all association cortex, and that it corresponds with reduced task decodability (Supplementary Fig. S5D).

Cross-region activity flows transform representational geometries

While 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 those of 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 correlations between observed and predicted activity patterns (averaged across subjects, task conditions, and sessions) ranged from $r = 0.16$ to 0.83 across regions (FDR-corrected p-values: 6.17e-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 1000 group-level null means. The observed group-level means were found to be significantly outside the null distribution ($p_{connperm} < 0.001$ for all target regions, Fig. 3A *bottom*).

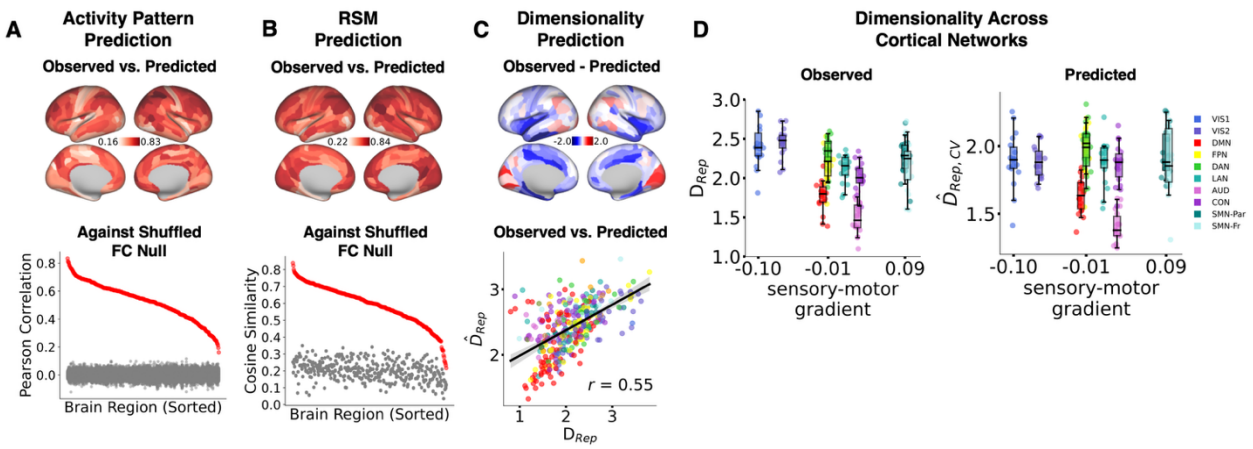

Figure 3: **Activity flow over connectivity generates observed representational geometries. A) Activity flow accurately predicts task-evoked activity patterns.** *Top:* Brain maps showing Pearson correlation between observed and activity flow-predicted activity patterns for each region, averaged across 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 (1000 iterations). All 360 regions show significantly better predictions from intact connectivity than shuffled connectivity (all $p < 0.001$). **B) Activity flow captures representational similarity structure.** *Top:* Brain maps showing cosine similarity between observed and predicted representational similarity matrices (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 $p < 0.001$). **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 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 specialized local computations compress representations beyond connectivity constraints. *Bottom:* Correlation between observed and predicted dimensionality across all 360 regions ($r = 0.55$, $p < 0.001$). Each dot represents one region, colored by network identity. **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 median and interquartile range. For both observed and predicted data, dimensionality is averaged across all regions within each network for each subject individually, then plotted here. Statistical comparisons between networks (VIS1 vs. DMN; DMN vs. 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. (End of Figure 3 caption)

To further validate specificity, we conducted a more conservative permutation test. We shuffled connectivity values only on the source side—randomly permuting connections among vertices within each source region while preserving the target region's connectivity structure. These source-constrained permutations maintained the target region's intrinsic connectivity patterns while disrupting the specific mapping from source regions. Even under these more stringent conditions, 337 out of 360 regions showed significantly better predictions from intact connectivity compared

to source-permuted connectivity ($p_{connperm} < 0.05$), demonstrating that the predictive accuracy reflects genuine structure in the source-to-target connectivity patterns rather than artifacts arising from altered target region connectivity geometry. These results go beyond the prior correlational analyses in demonstrating the hypothesized ability of FC patterns to generate empirically observed task representations.

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, while p_{FDR} values for t-tests against zero ranged from 3.30e-13 to 2.51e-05 (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 $p_{connperm} < 0.001$; source-side shuffling only: $p_{connperm} < 0.05$ 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: $t(17) = 10.94$, $p = 4.07e-09$, Fig. 3C *bottom*; comparison to FC-shuffled null distribution: $p_{connperm} < 0.001$; source-side shuffling only: $p_{connperm} < 0.001$). Finally, we verified that the dimensionality prediction was also consistent across brain regions within each system (sensory: $r = 0.76$, association: $r = 0.53$, motor: $r = 0.66$, all $p<0.0001$).

To identify where connectivity-based predictions succeed and where additional factors beyond connectivity are needed, we examined prediction residuals (observed minus predicted dimensionality; Fig. 3C *top*). The model systematically underpredicted dimensionality in early visual cortex and frontoparietal regions, suggesting these regions employ local dynamics or task-dependent reconfigurations that increase dimensionality beyond connectivity-based flow. Conversely, the model overpredicted dimensionality in late visual areas (inferotemporal cortex), medial temporal lobe, and salience network regions (insula, cingulate), suggesting specialized local computations compress representations beyond what connectivity geometry alone would produce. These functionally interpretable residual patterns highlight where region-specific computational properties operate beyond the connectivity-based baseline, and motivate future work examining how local circuit dynamics and task-dependent modulations interact with connectivity constraints.

We hypothesized that activity flows over observed connectivity patterns also generate the previously observed compression-then-expansion pattern in dimensionality[6] (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. While we used the same dataset as the original study reporting the compression-then-expansion pattern [6], we expanded the number of modeled task conditions from 45 to 96 to better characterize representational dimensionality (see 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 (VIS1) and motor (SMN-

Fr) networks showed significantly higher dimensionality than default-mode (DMN) and auditory (AUD) networks, which exhibited the strongest compression (VIS1 vs. DMN: $t(17) = 9.81$, $p = 2.05e-08$; SMN-Fr vs. DMN: $t(17) = 6.46$, $p = 5.86e-06$; VIS1 vs. AUD: $t(17) = 13.61$, $p = 1.43e-10$; SMN-Fr vs. AUD: $t(17) = 10.02$, $p = 1.49e-08$; Fig. 3D *left*). To test whether the activity flow model could reproduce this pattern, we employed a novel double cross-validated approach (see 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 default-mode and auditory networks (VIS1 vs. DMN: $t(17) = 5.49$, $p = 4.02e-05$; SMN-Fr vs. DMN: $t(17) = 5.11$, $p = 8.65e-05$; VIS1 vs. AUD: $t(17) = 9.57$, $p = 2.93e-08$; SMN-Fr vs. AUD: $t(17) = 13.93$, $p = 1.00e-10$; 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. Further, it remains possible that predicted RSMs of a target region were more similar to the source RSM than the true target RSM. Instead, we hypothesized 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 by computing two key distances: the distance between generated and target RSMs (d_{pred}) and the distance between connectivity-transformed and source RSMs ($\hat{d}_{transform}$) using cosine distance metric (1-cosine) (Supplementary Fig. S6A). The key comparison: is the model-predicted target RSM closer to the observed target RSM (d_{pred}) or closer to the observed source RSM ($\hat{d}_{transform}$)? If $d_{pred} < \hat{d}_{transform}$, this indicates genuine transformation rather than simple transfer. Our analysis revealed three key findings: first, connectivity-based transformations produced non-random effects, with 357 out of 360 regions showing transformation distances significantly different from null models based on randomly shuffled connectivity ($p_{connperm} < 0.001$; Supplementary Fig. S6C). This showed that structure present in the connectivity patterns yielded significantly lower transformation distances compared to random transformations from shuffled connectivity. Second, we validated our distance measure by demonstrating that connectivity-based transformation distances strongly correlated with observed source-to-target distances ($\hat{d}_{transform}$ vs.

$d_{transform}$ across regions: r = 0.62, t(17) = 19.60, p = 4.18e-13; SA-preserving permutation test: $p_{perm} < 0.001$; Supplementary Fig. S6B).

Finally, predicted representational geometries more closely matched target geometries than source geometries in 339 out of 360 regions ($p_{FDR} < 0.05$), with 292 regions showing highly significant differences ($p_{FDR} < 0.0001$; Supplementary Fig. S6D). Supporting our main hypothesis, these results suggest empirical connectivity patterns play a fundamental role in transforming representational geometries across brain 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 association cortex was accompanied by substantial reductions in FC dimensionality (Fig. 2C). What might be the cognitive benefits of such low-dimensional connectivity patterns? While our task battery was not designed to investigate shared cognitive components across tasks, we ran an exploratory analysis to examine how connectivity patterns relate to cross-task representations in the brain. We hypothesized that target regions with lower connectivity dimensionality have greater similarity among task representations, likely reflecting latent factors shared across tasks that are relevant to each region's function. Such shared latent task factors – reflecting reuse of neural representations across diverse contexts – likely support the human brain's ability to perform a wide variety of tasks.

To test this hypothesis, we developed a method to isolate and quantify cross-task similarities while excluding within-task effects. Focusing on representations shared across tasks (rather than between conditions of the same task) helped ensure the identified latent task factors were shared across diverse (rather than same) contexts. We computed mean cosine similarities between task condition combinations for each task pair (among the 16 tasks), including only positive values across subjects to ensure interpretability (Fig. 4A, see Methods). This approach revealed distinct patterns of cross-task relationships that varied by brain region.

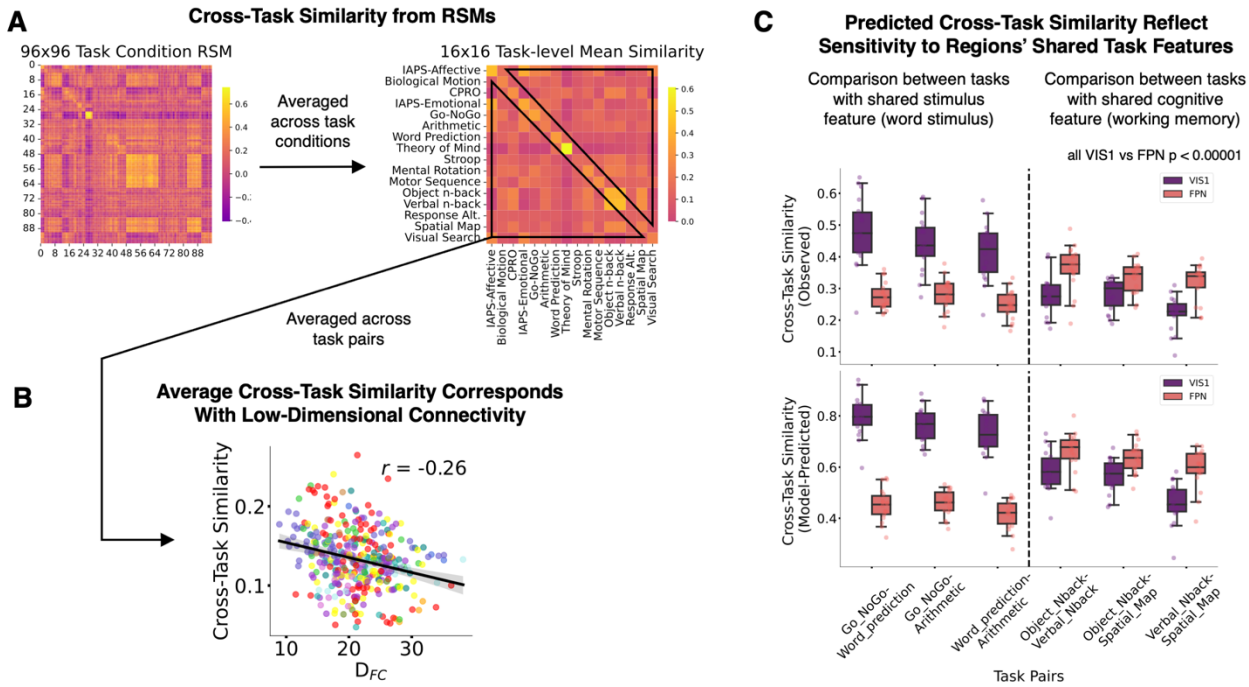

Figure 4: Low-dimensional connectivity generates functionally interpretable cross-task similarities. A) Deriving cross-task similarity from representational similarity matrices. Cross-task similarity quantifies how similar a region's representations are across different tasks, revealing shared latent 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 (e.g., 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), then averaged these positive values within each task pair to create a 16×16 task-level matrix. Off-diagonal elements represent cross-task similarity (e.g., 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.** Correlation between connectivity dimensionality (DFC) and mean cross-task similarity across all 360 regions ($r = -0.26$, $p < 0.001$). Each dot represents one region, colored by network identity. Lower connectivity dimensionality (similar connectivity patterns among target vertices) predicts greater similarity of task representations, suggesting 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 primary visual network (VIS1, purple) and frontoparietal network (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 distribution across subjects (n=18); each dot represents one subject's network-averaged cross-task similarity for that task pair. *Top:* Observed cross-task similarities. *Bottom:* Model-predicted cross-task similarities from activity flow. *Left panel:* Task pairs sharing sensory features (word stimuli): VIS1 shows significantly higher cross-task similarity than FPN across all three task pairs shown (all $p < 0.00001$), reflecting VIS1's sensitivity to shared visual word processing demands. *Right panel:* Task pairs sharing cognitive features (working memory): 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. (End of Figure 4 caption)

Our analysis yielded two key findings. First, we hypothesized that low-dimensional connectivity actively *generates* cross-task similarities—not merely that low-dimensionality 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 ($r = -0.26$, $t(17) = -24.92$, $p = 8.01e-15$; SA-preserving permutation test: $p_{perm} = 0.01$; Fig. 4B).

Validating that these cross-task similarities arise directly from connectivity patterns, activity flow over these same connectivity patterns successfully generated representations that reproduced the observed cross-task similarity structure (comparing vectorized cross-task similarity values) across most brain regions: cosine similarity ranged from 0.13 to 0.99, with $p_{connperm} < 0.05$ in 355 out of 360 regions, Supplementary Fig. S7A).

Finally, to test whether this relationship represents a general property of connectivity-based transformations, we conducted a simple simulation. This simulation used random activity patterns in source regions (serving as task activity patterns, thereby ensuring near-zero cross-task similarity at the source by construction) with activity flow mapping over simulated connectivity matrices with varying dimensionality (measured by participation ratio). The simulation confirmed that activity-flow based propagation over lower-dimensional connectivity matrices consistently resulted in greater average cross-task similarity values (Spearman correlation between connectivity dimensionality and resulting cross-task similarity: $\rho = -0.988$, $p = 1.08e-72$, see Methods for details).

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 twelve cortex-wide FC networks according to an established network parcellation[27]. While our task set was not designed to systematically investigate shared cognitive components, selective examination of two networks revealed potentially meaningful contrasts. For instance, in the primary visual network (VIS1), we observed high cross-task similarity among the *Go-NoGo*, *Word Prediction* and *Arithmetic* tasks, which share visual word processing demands. In contrast, the fronto-parietal 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 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. Cross-task similarities at the network level significantly differed from null models based on shuffled connectivity, for all networks (Supplementary Fig. S7B). Future work with a more systematic task battery, perhaps involving reuse of instruction elements across tasks (e.g., to allow for cross-condition generalization analysis), would be needed to thoroughly investigate these network- and region-specific patterns.

Together, these results demonstrate that low-dimensional connectivity contributes to shared representational structure across tasks. The successful generation of these patterns through activity flow modeling further supports the fundamental role of intrinsic connectivity in specifying task representations throughout the cortex.

Activity flow-driven expansions generate task-variable conjunctions

What role might high-dimensional connectivity patterns play in cognitive processing, particularly in the transformation of task representations? Building on prior literature on conjunctive coding, we hypothesized that connectivity patterns that expand the dimensionality of task representations might specifically support the generation of task-variable conjunctions. Representational conjunctions are task-evoked activity patterns consistently present only when two or more task features are present, and not when any subset of those features are present. For example, a conjunctive representation of the ink color blue and the Stroop task incongruent condition (blue $\cap$ incongruent) would be an activity pattern present only when the ink color is blue (not red, green, etc.) *and* the task condition is incongruent. These high-dimensional representations differ from low-dimensional compositional representations, which are reused across task conditions (e.g., an activity pattern present whenever blue is present across tasks). Conjunctions are important, for instance, for cognitive variables to interact in situation-specific ways during task implementation, such as processing the ink color blue differently for incongruent versus congruent Stroop task conditions. While 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.

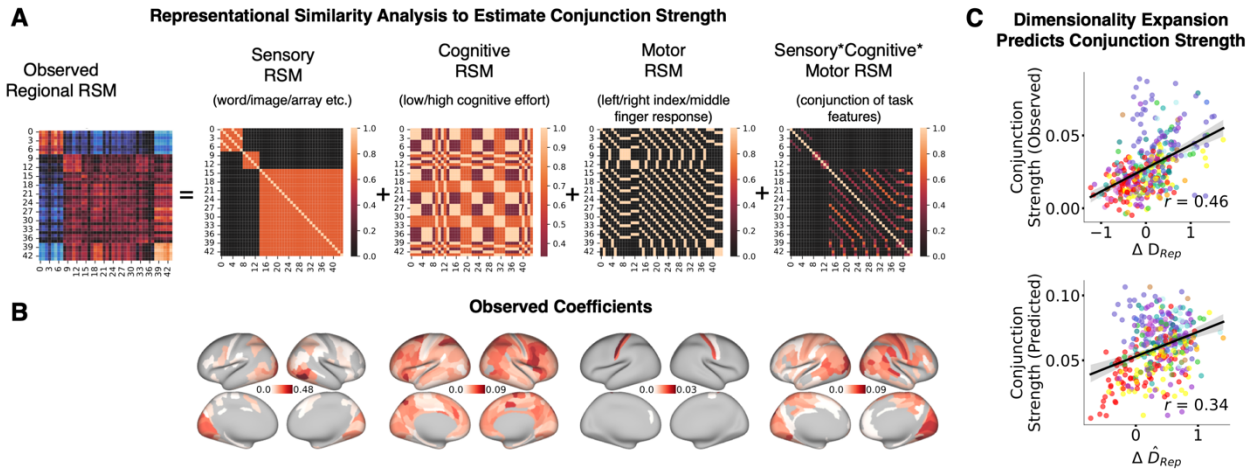

Figure 5: Connectivity-driven dimensionality expansion supports conjunctive encoding of task features. A) Representational similarity analysis to estimate conjunction strength. We identified a subset of 5 tasks (44 task conditions total) with separable sensorimotor and cognitive features: Stroop, Go/NoGo, IAPS-Affective, IAPS-Emotional, and Mental Rotation. For each region, we fit the observed 44×44 task condition RSM (left; example shown from right Area PH in visual cortex) using ridge regression against four template RSMs encoding different feature types. *Sensory template:* similarity based on stimulus type (word, image, spatial array). *Cognitive template:* similarity based on task demand level (low vs. 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 (e.g., "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 displayed (one-sample t-test against zero, FDR corrected). *Sensory:* highest in visual cortex (9/10 highest coefficients). *Cognitive:* highest in association cortex (7/10 highest in frontoparietal and default-mode regions). *Motor:* highest in motor and premotor cortex (8/10 highest). *Conjunction:* more distributed, with highest values in regions showing representational expansion. This anatomical specificity validates the RSA 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 ($r = 0.46$, $p < 0.001$). Each dot represents one region, colored by network. Regions showing expansion (positive ΔD_{Rep}) encode stronger conjunctive representations, while 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} ($r = 0.34$, $p < 0.001$). 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. (End of Figure 5 caption)

This subset included tasks like the Stroop task, where subjects responded to sensorimotor variables (ink color, mapped to specific button presses) under varying cognitive (task demand) conditions (congruent/incongruent). To test which task features drive representational similarity, we constructed four template RSMs: (1) sensory similarity (based on stimulus features like color, shape, words), (2) cognitive similarity (based on task demands like congruency, difficulty), (3) motor similarity (based on required button press responses), and (4) conjunction similarity (multiplicative combination of sensory × cognitive × motor features, capturing non-additive interactions). Using ridge regression, we fit these four vectorized matrices to the vectorized observed neural RSM (Fig. 5A).

These analyses revealed several patterns, all consistent with our hypothesis. First, confirming the efficacy of our template RSM approach, sensory, cognitive, and motor similarities were found to be encoded in their respective brain systems: Nine out of the 10 highest coefficients for the sensory term were observed in sensory regions, seven out of 10 highest coefficients for the cognitive term were in association regions, and seven out of 10 highest coefficients for the motor term were in motor regions (Fig. 5B). The conjunction term coefficients correlated positively with representational expansion ($r_{conj} = 0.46$, $t(17) = 8.13$, $p = 2.92e-07$; Fig. 5C *top*), a relationship stronger than those observed for individual feature terms (except sensory term). For sensory features alone: $r_{sensory} = 0.47$ ($p = 8.41e-10$), not significantly different from conjunctions (paired t-test: $t(17) = -1.94$, $p = 0.07$). For cognitive features: $r_{cognitive} = 0.33$ ($p = 0.013$), weaker than conjunctions ($t(17) = 2.74$, $p = 0.014$). For motor features: $r_{motor} = -0.19$ ($p = 0.002$), significantly weaker than conjunctions ($t(17) = 6.56$, $p = 4.86e-06$). This demonstrates that representational expansion specifically supported conjunctive coding beyond simple feature-based coding. To more directly test our hypothesis about connectivity-driven expansions supporting generation of

task-variable conjunctions, we conducted RSA with model-generated RSMs. We observed a correlation between predicted conjunction strength and predicted representational expansion ($r = 0.34$, $t(17) = 3.52$, $p = 0.006$, Fig. 5C *bottom*). Correlations of predicted representational expansion with the other three terms' coefficients were comparably smaller (sensory: $r = 0.16$, $t(17) = 3.57$, $p = 0.002$; cognitive: $r = 0.13$, $t(17) = 0.482$, $p = 0.636$; motor: $r = -0.15$, $t(17) = -2.23$, $p = 0.039$). This result demonstrates that generation of conjunctive representations specifically was supported by connectivity-based expansion: the activity flow model, using only connectivity geometry and source activity, reproduces the relationship between expansion and conjunction strength observed in empirical data. This suggests that high-dimensional connectivity architecture itself supports the neural capacity for conjunctive coding.

These results suggest a link between connectivity-driven expansions and the generation of task-variable conjunctions. While high-dimensional connectivity appears to support task-specific processing through conjunctive representations, we hypothesized that maintaining higher-dimensional representations might also relate to cognitive flexibility more broadly. To test this possibility, we analyzed cognitive performance data from the Human Connectome Project dataset ($n=352$ unrelated subjects with low motion).

Consistent with our hypothesis, region-averaged representational dimensionality showed a significant positive correlation with HCP general fluid intelligence composite scores (combining fluid cognition measures and matrix reasoning performance; see Methods; $r=0.29$, $p=1.505e-08$, Supplementary Fig. S8A). Region-specific correlations between dimensionality and general fluid intelligence, after correction for multiple comparisons, revealed several regions with significant positive correlations (Supplementary Fig. S8B). To rigorously assess system-specific contributions, we employed cross-validation using a multiple regression model that separated dimensionality contributions from sensory, association, and motor regions. Using leave-one-subject-out cross-validation, predicted and actual cognitive flexibility scores showed strong correspondence ($r=0.278$, $p=1.17e-07$, Supplementary Fig. S8C). While coefficients averaged across iterations were positive for all three region types, sensory and association regions showed stronger predictive power than motor regions (sensory: 0.943, association: 0.853, motor: 0.142, all pairwise comparisons are significant, $p<0.00001$, Supplementary Fig. S8D).

These findings complement our observations about conjunctive coding by suggesting two distinct ways that high-dimensional representations support flexible cognition. First, as demonstrated in our conjunctive coding analysis, high-dimensional connectivity enables the integration of multiple task variables for specific cognitive demands. Second, maintaining higher-dimensional representations, particularly in sensory and association regions, may enhance an individual's capacity for cognitive flexibility by enabling better discrimination between task contexts and reducing interference between competing task representations. 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.

Discussion

We identified a systems-level organizing principle by which the human brain transforms task representations across diverse cognitive tasks. Specifically, the dimensionality of intrinsic connectivity patterns directs the cross-region compression and expansion of neural representations. Low-dimensional connectivity leads to compressed representations that show shared similarities across tasks, while high-dimensional connectivity expands representations to integrate multiple task variables for specific cognitive demands. This connectivity-based mechanism generates representational transformations across sensory, association, and motor systems, as demonstrated by our activity flow modeling approach. The model not only generates observed representational geometries but also captures the broader pattern of representational complexity across the cortical hierarchy. Critically, regions with lower-dimensional connectivity show stronger cross-task similarities, indicating shared latent features that likely enable neural processing to generalize across diverse tasks. Meanwhile, high-dimensional regions support the creation of rich, conjunctive representations that situate this generalized information for specific task contexts. The robustness of this mechanism is further demonstrated by our replication of the compression-then-expansion pattern in an independent Human Connectome Project dataset, where we also found that individual differences in representational dimensionality predict general fluid intelligence across diverse tasks. Together, these findings reveal how the brain's intrinsic connectivity architecture enables cognitive flexibility through systematic transformation of task-related information.

Our findings reveal a specific computational regime in which the human brain possibly operates. Through simulation, we demonstrate that strong non-linear local computations can fully decouple connectivity dimensionality from representational transformations (Supplementary Fig. S4), yet empirically we observe robust coupling ($r=0.44$ across 360 regions). This positions 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: the brain leverages network geometry as a computational substrate rather than relying exclusively on complex local computations in every region. Connectivity patterns, once established during development, provide stable constraints that systematically shape representations across diverse tasks without requiring task-dependent reconfiguration of local circuits. The prediction residuals reveal where this constraint-based approach proves insufficient—early visual cortex and frontoparietal regions show underprediction, while inferotemporal and medial temporal regions show overprediction—indicating a principled partitioning of computational labor between distributed architecture and specialized local dynamics. Importantly, this framework generates testable predictions: diseases affecting white matter connectivity should produce predictable distortions in representational geometry, and interventions targeting connectivity should have systematic effects on downstream representational transformations.

Our findings provide empirical validation of theoretical predictions from computational neuroscience regarding how connectivity structure constrains representational geometry, broadly [18–20,31–33]. While these principles have been extensively characterized in artificial neural networks, their empirical validation in biological circuits has remained to be demonstrated [34]. Our results bridge this gap by demonstrating that connectivity-based constraints on representational geometry

operate as fundamental organizing principles in the human brain during diverse cognitive tasks. Our approach using resting-state functional connectivity provides a functionally defined measure of connectivity constraints complementary to the structural connectivity emphasized in theoretical work, enabling empirical tests of these principles in the human brain during diverse cognitive tasks.

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 Rules Operation (C-PRO) paradigm (one of the tasks included in the present study's task battery) demonstrated that task rule information transfers between brain regions via resting-state connectivity [22]. Here, we show this principle generalizes far beyond systematic rule combinations to a diverse set of common cognitive tasks that lack such structured relationships. This broader demonstration supports connectivity-based information transfer as a general principle of neural computation. Our findings build upon the activity flow modeling framework, which established that resting-state functional connectivity can predict the movement of task activations between neural populations [21], by showing how these connectivity patterns actively transform task representations between brain regions. While 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 prior findings showing that pairwise similarity between regional RSMs correlates with resting-state connectivity at the parcel level [6]. 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 generalization [35]. Through our analysis of cross-task similarities across the cortex, we found that activity flows over intrinsic connectivity patterns generate widespread shared representations. This aligns with previous work showing that cognitive encoding models can explain variance in 86% of cortical voxels [2], as our connectivity-based model similarly achieves significant predictions in 355 out of 360 regions, but through direct voxel-to-voxel mappings mediated by connectivity rather than predefined cognitive factors. While expert-annotated cognitive factors can predict activation patterns for novel tasks [5], our approach leverages a brain-based ontology grounded in intrinsic connectivity patterns [36]. This connectivity-based framework remains open to discovering novel latent cross-task factors that emerge from how the brain naturally routes information, even when these factors do not correspond directly to existing cognitive constructs. Future work using task-state FC could further illuminate these brain-based latent factors, potentially revealing additional principles of neural information routing during task performance[37].

Our findings complement a growing body of work in animal models examining how inter-regional connectivity shapes neural computation[38–40]. Studies in non-human primates have demonstrated that functional connectivity between sensory and association areas constrains how information is communicated and transformed across regions[41], with communication subspaces between areas shaping what information can be transmitted[42–44]. Recent work in mice similarly shows that inter-regional coordination structures cortical dynamics and supports flexible behavior[45]. 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 key conceptual point warrants clarification regarding our use of "compression" and "expansion" to describe representational transformations. Activity flow modeling as implemented in this study involves convergence in the literal sense: we include all functionally connected regions as sources (though alternative implementations may restrict sources to specific regions of interest[23]), such that many source vertices (tens of thousands across all connected regions) project to fewer target vertices (tens to hundreds within a region). However, we define compression and expansion relative to the *average dimensionality of individual source regions* rather than the pooled dimensionality of all source vertices combined. This distinction is important because it allows us to characterize region-to-region transformations — whether a given brain region represents tasks more or less distinctly than a typical source region — while acknowledging that the mechanistic reality involves simultaneous integration across many sources. Notably, this framework reveals that convergent connectivity architectures can produce dimensionality expansion through a plausible linear mechanism: when source regions represent tasks in different low-dimensional subspaces, and target vertices receive diverse weighted mixtures of these sources (high connectivity dimensionality), the target can occupy a higher-dimensional representational space spanning combinations of source subspaces. For example, if one source region primarily encodes left/right distinctions and another encodes go/nogo distinctions, diverse linear mixing at the target can produce representations that distinguish all four combinations (left-go, left-nogo, right-go, right-nogo), yielding higher dimensionality than either source alone. This mechanism may help explain our conjunctive coding results (Fig. 5): regions with high-dimensional connectivity patterns that permit diverse source integration show stronger conjunction-like representations, potentially because linear mixing of diverse sources naturally creates feature combinations that approximate conjunctive coding. This hypothesis generates testable predictions – for instance, regions showing connectivity-driven expansion should receive inputs from source regions with lower cross-source representational similarity (i.e., sources occupying distinct representational subspaces) – and suggests that the geometric structure of connectivity patterns fundamentally shapes not only whether representations compress or expand, but also the computational operations (generalization versus specialization) that different cortical regions implement.

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 frontoparietal networks. These findings align with the neural reuse framework[46,47], 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 RSA 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 flows 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.

Our analysis of the HCP dataset revealed that higher representational dimensionality, particularly in sensory and association regions, predicts general fluid intelligence across individuals. Notably, our use of a composite score across a variety of cognitive flexibility tests resulted in a measure of general fluid intelligence, which reflects a highly general cognitive ability shared across a wide variety of tasks[48]. Finding that higher representational dimensionality predicts general fluid intelligence appears to present a paradox when considered alongside the robust compression-then-expansion pattern we observed in both datasets: if compression in association regions is a fundamental organizing principle, why would higher dimensionality in these regions predict better cognitive performance? We propose that these findings reflect distinct but complementary aspects of neural information processing. The compression-then-expansion pattern likely represents a canonical within-subject organization that supports efficient task processing: compression in association regions enables abstraction of task-general features and reduction of redundant information, while subsequent expansion in motor regions supports task-specific response implementation. However, the absolute level of dimensionality maintained within this organizational framework may determine an individual's capacity for flexible cognition. Higher dimensionality in sensory and association regions, rather than motor regions, was most predictive of fluid intelligence. This suggests that the critical bottleneck for flexible behavior may lie in the initial perception and abstract representation of tasks, rather than in keeping motor responses separate. Maintaining higher-dimensional representations in these early and intermediate processing stages may enable better discrimination between similar task contexts and reduce interference before the motor execution stage.

In summary, our study reveals a fundamental organizing principle by which the brain's intrinsic connectivity architecture enables cognitive flexibility through systematic transformation of task representations. This connectivity-based control of representational dimensionality across the cortex helps explain how a single system – the human brain – can perform such diverse cognitive tasks. This mechanism provides a framework for understanding both healthy cognitive function and dysfunction. Future work can investigate how these connectivity-based representational mechanisms adapt during learning and development, and whether their disruption contributes to cognitive deficits in neurological and psychiatric conditions.

Acknowledgements

This work was supported by the US National Science Foundation (NSF) under award 2219323. Multi-domain Task Battery data were provided by Maedbh King, Carlos Hernandez-Castillo, Russell A. Poldrack, Jonathon Walters, Richard B. Ivry, Jörn Diedrichsen via OpenNeuro (doi:10.18112/openneuro.ds002105.v1.1.0). HCP data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for

Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. 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. The authors thank the members of Cole Neurocognition Lab for helpful conversations during the course of the research. The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

Materials and Methods

Multi-domain task battery dataset

Portions of this section are paraphrased from the original data study [49].

We used the multi-domain task battery (MDTB) dataset [49], 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 females, 8 males; mean age 23.8 ± 2.6 years) at Western University. Participants provided informed consent under an approved experimental protocol. Detailed participant selection criteria are outlined in the original study.

The MDTB 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 one year for the others. Each set involved two 10-minute imaging runs. For 18 out of 24 participants, an additional resting-state scan (two 10-minute 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. Detailed task and condition information is available in the supplementary materials of the original publication. Tasks were performed once per imaging run using an interleaved block design. Each task block consisted of a 5-second instruction screen followed by 30 seconds of continuous task performance, totaling 35-second 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 index or middle fingers 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) for subsequent analyses, decomposing them into 96 task conditions based on stimulus type, motor response, and context variables (Table 1). Note that the decomposition of task conditions into 96 distinct conditions is unique to this study, and was carefully performed by scrutinizing the original experimental data. By focusing on visual tasks, we treated visual regions as sensory regions in the sensory-association-motor hierarchy. Detailed information about fMRI acquisition parameters can be found in the original study.

Task	# Conditions	Task Condition	Key	stimType1	stimType2	stimType3
IAPS-Affective	2	Affective_pleasant	L_I	image	scene	
		Affective_unpleasant	L_M			
IAPS-Emotional	2	Emotional_happy	R_I		face	
		Emotional_sad	R_M			
Mental Rotation	6	mentalRotation_doesnotmatch_0_deg	R_M		objects	
		mentalRotation_doesnotmatch_50_deg	R_M			
		mentalRotation_doesnotmatch_150_deg	R_M			
		mentalRotation_match_0_deg	R_I			
		mentalRotation_match_50_deg	R_I			
		mentalRotation_match_150_deg	R_I			
Object n-back	12	nBackPic_notpossible2backfork.jpg	none			fork image
		nBackPic_possible2backfork.jpg	R_I			
		nBackPic_notpossible2backhydrant.jpg	none			hydrant image
		nBackPic_possible2backhydrant.jpg	R_I			
		nBackPic_notpossible2backlamp.jpg	none			lamp image
		nBackPic_possible2backlamp.jpg	R_I			
		nBackPic_notpossible2backplate.jpg	none			plate image
		nBackPic_possible2backplate.jpg	R_I			
		nBackPic_notpossible2backwhistle.jpg	none			whistle image
		nBackPic_possible2backwhistle.jpg	R_I			
		nBackPic_notpossible2backzip.jpg	none			zip image
		nBackPic_possible2backzip.jpg	R_I			
Concrete Permuted Rule Operations (CPRO)	8	CPRO_doesnotfollowRule_buttonOne	L_I	[image, text]	[object, word]	press button One
		CPRO_followsRule_buttonOne	L_M			
		CPRO_doesnotfollowRule_buttonTwo	L_M			press button Two'
		CPRO_followsRule_buttonTwo	L_I			
		CPRO_doesnotfollowRule_buttonThree	R_M			press button Three'
		CPRO_followsRule_buttonThree	R_I			
		CPRO_doesnotfollowRule_buttonFour	R_I			press button Four'
		CPRO_followsRule_buttonFour	R_M			
Go / No-Go	2	GoNoGo_Go	L_I	text	word	
		GoNoGo_NoGo	none			
Word Prediction	3	Prediction_random	L_M			
		Prediction_valid	L_I			
		Prediction_violation	L_M			
Stroop	8	cong_color1	L_M			ink color 1

		incong_color1	L_M			
		cong_color2	L_I			ink color 2
		incong_color2	L_I			
		cong_color3	R_I			ink color 3
		incong_color3	R_I			
		cong_color4	R_M			ink color 4
		incong_color4	R_M			
Theory of Mind	4	ToM_ToM_expected_false	L_M		sentence	
		ToM_ToM_expected_true	L_I			
		ToM_control_expected_false	L_M			
		ToM_control_expected_true	L_I			
Arithmetic	4	DigitJudgement_target_absent	R_M		letter-digit	
		DigitJudgement_target_present	R_I			
		Math_equation_correct	R_I			
		Math_equation_incorrect	R_M			
Motor Sequence	5	motorSequence_complex	all			
		motorSequence_simple_digit1	L_M			
		motorSequence_simple_digit2	L_I			
		motorSequence_simple_digit3	R_I			
		motorSequence_simple_digit4	R_M			
Verbal n-back	6	nBack_notpossible2backA	none			letter 'A'
		nBack_possible2backA	L_I			
		nBack_notpossible2backB	none			letter 'B'
		nBack_possible2backB	L_I			
		nBack_notpossible2backC	none			letter 'C'
		nBack_possible2backC	L_I			
Response Alternatives	12	respAlt_1_prime_targPos1	L_M	[text, array]	letter-digit	
		respAlt_1_prime_targPos2	L_I			
		respAlt_1_prime_targPos3	R_I			
		respAlt_1_prime_targPos4	R_M			
		respAlt_2_primes_targPos1	L_M			
		respAlt_2_primes_targPos2	L_I			
		respAlt_2_primes_targPos3	R_I			
		respAlt_2_primes_targPos4	R_M			
		respAlt_4_primes_targPos1	L_M			
		respAlt_4_primes_targPos2	L_I			
		respAlt_4_primes_targPos3	R_I			
		respAlt_4_primes_targPos4	R_M			
	12	spatialMap_1_element_targNum1	L_M			

Spatial Map		spatialMap_1_element_targNum2	L_I			
		spatialMap_1_element_targNum3	R_I			
		spatialMap_1_element_targNum4	R_M			
		spatialMap_4_elements_targNum1	L_M			
		spatialMap_4_elements_targNum2	L_I			
		spatialMap_4_elements_targNum3	R_I			
		spatialMap_4_elements_targNum4	R_M			
		spatialMap_7_elements_targNum1	L_M			
		spatialMap_7_elements_targNum2	L_I			
		spatialMap_7_elements_targNum3	R_I			
		spatialMap_7_elements_targNum4	R_M			
Visual Search	6	visualSearch_target_absent_setsize12	L_M			
		visualSearch_target_absent_setsize4	L_M			
		visualSearch_target_absent_setsize8	L_M			
		visualSearch_target_present_setsize12	L_I			
		visualSearch_target_present_setsize4	L_I			
		visualSearch_target_present_setsize8	L_I			
Biological Motion	4	BiologicalMotion_happy	R_I	video	point light	
		BiologicalMotion_sad	R_M			
		ScrambledMotion_fast	R_I			
		ScrambledMotion_slow	R_M			

Table 1. Task condition split of 16 active, visual tasks in Multi-domain Task Battery Dataset. For more information on task description, see the supplementary information of the original paper [49] [https://static-content.springer.com/esm/art%3A10.1038%2Fs41593-019-0436-x/MediaObjects/41593_2019_436_MOESM1_ESM.pdf]. Response key index: L_I: left index finger button, L_M: left middle finger button, R_I: right index finger button, R_M: right middle finger button, all: response using all four buttons was required, none: no response required.

fMRI preprocessing

This section is paraphrased from the preprocessing details of Ito & Murray (2023) [6], which used the same preprocessing pipeline.

The preprocessing of the fMRI data in this study utilized the Human Connectome Project (HCP) preprocessing pipeline[50], which was executed through the Quantitative Neuroimaging Environment & Toolbox [QuNex, version 0.61.17] [51]. 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 employed a modified version of the single-subject beta series regression technique to analyze the fMRI task data [52]. 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 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 SPM canonical hemodynamic response function to account for temporal lags in hemodynamic responses [53]. 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 to data from each task run using ridge regression in Python (version 3.8.5), utilizing the 'ridge_corr' function (https://github.com/alexhuth/ridge/blob/master/ridge.py). Ridge regression was chosen to address potential multicollinearity issues arising from closely spaced task conditions within each task. In each run, the optimal regularization parameter (alpha) was selected using four-fold cross-validation, testing a range of alpha values from 1 to 1000 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 research [54]. The nuisance parameters included 24 motion-related regressors (six motion parameters, their time-derivatives, and quadratic terms) and eight 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 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). While 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-minute resting-state runs for each subject to measure FC between brain regions when participants were not engaged in any task. Following preprocessing, the timeseries were subjected to nuisance regression, as described in the section above, and the residuals from that step were used for this analysis. The resting-state timeseries from the two runs were concatenated for each subject before estimating FC in two steps, as detailed below.

First, we estimated sparse FC between mean parcel timeseries from the Glasser parcellation (see section '*Discrete brain systems and continuous FC gradient estimation*' below for justification of Glasser parcellation) [55] using a regularized partial correlation approach with L1 penalty (graphical lasso). This method has been shown to improve both validity against structural connectivity and repeat reliability at the individual level for unconfounded FC estimates [56]. The estimation was performed using the 'GGLasso' package [57]. This two-step approach—sparse parcel-level FC followed by dense vertex-level FC—was designed to balance statistical power with spatial specificity. Parcel-level graphical lasso identifies which regions are functionally connected, providing a principled graph structure. Vertex-level principal component regression then estimates fine-grained connectivity patterns within this graph structure, enabling analysis of connectivity geometry at high spatial resolution while avoiding the intractability of estimating full vertex-to-vertex connectivity across the entire cortex (which would involve millions of connections). The regularization parameter (alpha) was optimized using a range from 0.001 to 1 in steps of 10, with the optimal alpha chosen based on the minimum negative log-likelihood value after averaging across folds.

The sparse FC at the parcel level for each individual provided the graph to identify connected neighbors of each region. Dense FC at the vertex level was then estimated between vertices of each target region (analyzed one at a time) and all vertices of its connected neighbors (source regions, as determined by the parcel-level graph). This produced a rectangular connectivity matrix for each target region of size (number of source vertices) × (number of target vertices), capturing the specific pattern of connectivity weights from each source vertex to each target vertex. For each vertex in the target region, source region vertices were treated as regressors in a cross-validated principal component regression. Principal component regression was chosen to address a key challenge: estimating connectivity for each target vertex from potentially thousands of source vertices using only ~1,200 timepoints. By projecting source time series onto their principal components, we reduce dimensionality while retaining the major patterns of covariance, preventing overfitting while maintaining biological plausibility of connectivity estimates. Note that principal component regression has been used successfully for FC estimation in multiple prior studies [21,22,58]. The number of principal components was determined by 10-fold cross-validation, with the optimal number chosen per subject based on the lowest average mean squared error across all vertices. This resulted in a mean of 57.1 principal components across subjects (range: 51 - 71). Importantly, source-region vertices falling within 10 mm of the target region's boundaries were excluded to prevent circularity due to spatial autocorrelation. This approach provided dense FC matrices between each target region and its corresponding sources. As a quality control measure, we calculated the goodness of fit (r-squared) of the principal component regression to the resting-state fMRI data used for model training. On average across subjects, the range of r-squared across regions varied from 0.24 to 0.66.

RSM estimation

We followed an approach similar to that described in [6] to estimate RSMs. For each region (parcel), we measured the shared similarity structure between different task conditions' activity estimates. Specifically, we calculated the cosine similarity between task condition activity vectors in each parcel, cross-validated between two sessions (e.g., cosine similarity between activity of task condition A from session 1 and task condition B from session 2). Cosine similarity was chosen

because it captures both the pattern of activation across vertices (comparing directions of two vectors in vertex space) and the overall magnitude of activation, unlike Pearson correlation which only considers the pattern of deviations from the mean. Cross-validating between sessions yielded a non-trivial diagonal, revealing its test-retest reliability. This process resulted in a 96x96 RSM for each of the 360 parcels from the Glasser parcellation in each subject, corresponding to the 96 task conditions derived from the 16 tasks as described in the 'Multi-domain task battery dataset' section.

Representational and connectivity dimensionality

We measured representational dimensionality using the participation ratio of the eigenvalue spectrum obtained from the eigen decomposition of each RSM. The participation ratio quantifies the rate of decay in the eigenvalue spectrum and serves as a continuous measure of dimensionality. It is calculated as:

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

where λ_i are the eigenvalues of the RSM.

Previous research has shown that this measure corresponds to the number of principal components required to explain approximately 80% of the explained variance [59]. The participation ratio was chosen rather than this alternative as it provides a continuous measure of dimensionality, allowing for more nuanced comparisons across brain regions and subjects.

We related the metric of representational dimensionality to multi-task decoding of 96 task conditions using minimum-distance classifier based on cosine distance (cross-validated across two sessions). This can be intuited directly from the RSM such that a task condition is decodable if the diagonal value is higher compared to rest of the values in the row/column corresponding to the task condition.

For connectivity dimensionality, we sought a measure that maximally corresponded to our measure of representational dimensionality, but that appropriately took the unique aspects of connectivity data (which forms a rectangular matrix connecting two distinct sets of vertices: those from source regions and those from target regions) into account. To estimate connectivity dimensionality we first performed singular value decomposition on the dense FC matrix between vertices of each target region and its corresponding 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:

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

Where s_i are singular values following singular value decomposition.

These dimensionality measures were calculated separately for each subject, providing individual-level estimates of representational and connectivity complexity. In the context of representational dimensionality, higher values indicate that task conditions are represented in a highly specialized manner, with distinct patterns for different conditions. Lower representational dimensionality suggests a possibly shared structure among task conditions, which can be probed to identify patterns of abstraction and latent factors. For connectivity dimensionality, lower values reflect convergent, homogeneous input patterns from source regions, whereas higher values reflect more distributed or diverse input patterns..

As part of control analyses, we tested for potential parcel-specific properties that vary across the cortex and influence both representational and connectivity dimensionality. We measured the following confounds: target region size (number of vertices per parcel), signal-to-noise ratio (measured by averaging the values along the diagonal of each RSM), mean inter-vertex distance (mean geodesic distance from a vertex to its neighboring vertices connected by a mesh edge[60]), temporal variance of resting-state activity (standard deviation of the time series signal in each vertex across time, averaged across vertices in each parcel), and spatial homogeneity of activity (standard deviation across vertices within each parcel at each time point, averaged across time). For correlation between connectivity and representational dimensionality, these confounds were first regressed out of connectivity dimensionality and the residuals were used for correlation with representational dimensionality. For across-parcel correlations, spatial autocorrelation preserving permutation tests were performed using BrainSMASH toolbox[61] by generating 1000 null surrogate maps for comparison.

Discrete brain systems and continuous FC gradient estimation

To delineate the sensory, association, and motor regions that constitute the cortical information processing hierarchy, we employed the Glasser parcellation and the Cole-Anticevic Brain Network Partition (CAB-NP) as functional network definition [27]. We used the Glasser et al. (2016) multi-modal parcellation (HCP-MMP1.0) to define cortical regions. This parcellation was selected because it overcomes the limitations of unimodal approaches by identifying areas where multiple independent neurobiological properties converge. Specifically, it utilizes a gradient-based approach that incorporates cortical architecture (myelin content and thickness), function (task-fMRI activations), functional connectivity, and topographic organization. This multi-modal integration allows for a more precise delineation of 180 distinct areas per hemisphere, providing significantly improved spatial localization compared to traditional volume-based or resting-state-only parcellations [55,62]. The CAB-NP provided a systematic approach for delineating known functional networks using the Glasser parcellation, dividing it into 12 distinct functional networks. We categorized the primary and secondary visual network regions as sensory (restricted to visual networks given our focus on visual tasks). For the somatomotor network, we separated regions into somatosensory (SMN-Parietal, posterior to the central sulcus) and motor (SMN-Frontal, anterior to the central sulcus), and the remaining regions as association areas.

To derive continuous connectivity gradients from sensory to association to motor regions, we analyzed the parcel-level sparse FC graph. Consistent with previous studies [63], the parcel-level FC (averaged across subjects) was first thresholded to retain the top 20% of its values. Subsequently, we applied principal component analysis (PCA) to the thresholded graph. As previously reported, the first principal component typically runs from unimodal (sensory and motor) regions to

transmodal (association) regions. In contrast, the second principal component exhibited continuous variation from sensory to association to motor regions, which is thought to represent the broad cortical hierarchy of information processing. We refer to this as the sensorimotor FC gradient. This pattern has been widely replicated in previous studies, and we visually confirmed its presence in our data. This sensorimotor FC gradient provided a continuous measure of cortical organization that complements the discrete categorization of brain regions.

Connectivity pattern similarity and Sparsity

To ground the abstract measure of connectivity dimensionality in interpretable terms, we analyzed the connectivity patterns of target region vertices. These patterns are represented by vectors of connectivity weights from source region vertices to a target region vertex. We aimed to understand the similarity in how the set of source region vertices are connected to the vertices within each corresponding target region. Our focus on target vertices' patterns stems from the activity flow framework, which is a theoretical framework regarding how activity is generated in a given (target) neural population. Moreover, the connectivity dimensionality derived from singular value decomposition is ideally capped at the target region size, given that it is the smaller of the two dimensions in the dense connectivity matrix for each region as target. For each target region, we computed the Pearson correlation between the connectivity patterns of its vertices, resulting in a correlation matrix of size t × t, where t is the number of vertices in the region. We then derived the mean correlation, providing a measure of overall similarity among connectivity patterns.

Sparsity analysis. To test whether connectivity dimensionality reflects sparsity (proportion of near-zero connection weights) rather than geometric structure, we computed approximate sparsity of the dense FC matrices. Since principal component regression produces dense matrices (all weights are non-zero), we defined approximate sparsity by setting thresholds near zero (±0.0001, ±0.00005, ±0.00001) and computing the proportion of weights falling within these thresholds. We then examined whether this sparsity measure correlated with connectivity dimensionality across regions.

Activity flow modeling

We implemented activity flow modeling, an approach to model network flow-based generation of task-evoked activity adapted from previous studies [21,22]. For each subject and target region, we identified connected neighbors from the parcel-level FC graph as source regions, using the method described in the 'FC estimation' section. This approach allowed us to relate fine-grained representations across vertices in each brain region with dense FC estimated between the vertices of connected brain regions. In brief, we projected activity estimates from all source region vertices onto the target region using the dense FC graph. This was accomplished by performing matrix multiplication between the activity matrix (size: task conditions × source vertices) and the FC graph (size: source vertices × target vertices) in each session. The matrix multiplication summarizes series of multiplicative weighting of each task condition's source activity by corresponding connectivity weight, summed to generate predicted activity in each target. More precisely, a target vertex's task-evoked activity (a_t) was generated by projecting the source vertex activity (a_s) over the source-target connectivity weight (c_{ts}), [with the product construed as the connectivity-specified flow ($a_s c_{ts}$) of task-evoked activity from that specific source vertex], summed over vertices of all connected (source) regions ($a_t = \sum_{s=1}^{n} a_s c_{ts}$), where n is the total

number of source vertices. To evaluate the model, we computed the Pearson correlation between the flow-generated activity matrix (size: task conditions × target vertices) and the empirically observed task condition activity matrix in each target region, after averaging over task conditions and two sessions. This correlation quantified the correspondence between the observed and model-predicted estimates.

Using the model generated activity estimates, we then derived estimates of model-predicted RSMs and representational dimensionality. For comparison of observed and model-generated measures, we compared model-predicted RSMs against the actual observed values using cosine similarity in each region. We also assessed the relationship between model-predicted and observed representational dimensionality values across regions using Pearson correlation. This approach tests whether the geometric structure of FC constrains task-evoked representational geometry. High correlations in the model evaluation and comparison steps would suggest that the source activity and FC structure accurately generate task-evoked activity patterns and representational geometries. Regions where predictions are less accurate point to where additional factors—such as local circuit dynamics, task-dependent neuromodulation, or attention—operate beyond the baseline constraints provided by intrinsic connectivity architecture. Examining prediction residuals (observed minus predicted values) can reveal systematic patterns indicating which regions rely more heavily on local computations versus distributed connectivity-based transformations.

Significance of the predictions is estimated by running the projection of the source activities over 100 iterations of randomly shuffled connectivity values in each subject, from which 1000 group-level means were generated via bootstrapping [64]. For more stringent test, the shuffling of the source-to-target connectivity matrices was constrained to source side only. That is, the connectivity values were randomly permuted among vertices within each source region. This preserved spatial distribution properties unique to each source by preventing mixing between vertices of different sources. More importantly, it preserved the connectivity properties of the target vertices and their spatial distribution as they are left unaltered. Significance against such null distributions demonstrates that specific source-to-target connectivity patterns, rather than general connectivity properties (e.g., spatial distribution of target vertex degree), drive the predictive accuracy of activity flow modeling.

Double cross-validated RSMs for cortical dimensionality pattern comparison

To compare observed and model-predicted dimensionality patterns along the cortical hierarchy, we employed a novel approach inspired by cross-decoding techniques. We ran this comparison to verify whether model-predicted dimensionality patterns replicate the compression-then-expansion pattern of multi-task representations across cortical regions along the sensorimotor FC gradient [6].

We developed double-cross-validated RSMs to test the ability of activity flow-predicted task estimates to capture the representational geometry of observed activity patterns. This approach involved cross-validating between observed and model-predicted estimates, in addition to the previously defined cross-validation between two identical sessions from which task activity estimates were computed. Specifically, we computed cosine similarity between observed task activity estimates of condition A from session 1 and model-predicted task activity estimates of

condition B from session 2. This process yielded double-cross-validated RSMs (size: task conditions × task conditions) that captured similarities between sessions and between observed and model-predicted estimates. We then repeated this procedure for the opposite combination (model-predicted activity estimates from session 1 against observed activity estimates from session 2). From these double-cross-validated RSMs, we computed representational dimensionality values for each combination and averaged them within each cortical region. This approach allowed us to validate the model's predictions by assessing their similarity to observed data, and identify robust, shared patterns that persist across both observed and model-predicted estimates, while reducing the influence of noise and session-specific artifacts. This method provided a stringent test of the model's predictive power and helped isolate the most reliable aspects of the neural representations.

RSM-based transformation index

To quantify the degree of representational transformation across diverse task conditions predicted by our model, we developed an RSM-based transformation index. This index measures the extent to which the model-generated representations in a target region differ from the observed representations in the source regions. For each target region, we computed three measures: (a) observed transformation distance ($d_{transform}$): the cosine distance (1 - cosine similarity) between each observed source RSM and the observed target RSM, averaged across all source-target pairs, (b) predicted transformation distance ($d_{transform*}$): the cosine distance between each observed source RSM and the predicted target RSM, averaged across all source-target pairs, (c) prediction distance (d_{pred}): the cosine distance between the predicted target RSM and the observed target RSM, serving as a reference measure. This approach allowed us to assess the significance and validity of representational change from source to target regions as predicted by the model, by benchmarking against the model's predictive error.

Cross-task representational similarity

To better understand cross-task relationships, we aimed to strictly isolate representations of similarities across tasks, distinct from similarities across task conditions within each task. Firstly, we performed a one-sample t-test for each cell of the RSM across subjects, applying false discovery rate correction for multiple comparisons. This yielded significance-survived RSMs, revealing more reliable task-condition combinations. From these modified RSMs, we computed the average positive similarity values among the survived task condition combinations. We restricted the averaging to positive values only, as negative cosine similarity values are difficult to interpret. This measure provided insights into the nature and strength of relationships between different tasks across brain regions.

To investigate how connectivity dimensionality influences cross-task similarity patterns in a more general manner, we simulated activity propagation from source to target regions across connectivity matrices with varying dimensionality. Source and target region sizes were randomly sampled from uniform distributions of (100,150) and (50,100) vertices respectively, reflecting the typically larger number of source vertices compared to target vertices. For each source region, we generated 100 random task activity patterns (drawn from a standard normal distribution) to propagate through the connectivity matrix. Since connectivity dimensionality is constrained by the smaller target region size, we created low-rank connectivity matrices at factors of [0.01, 0.025, 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1.0] times the full rank (target region size), also populated using values

from standard normal distribution. For each rank factor, we simulated 10 different source-target pairs with varying sizes. After propagating random activity patterns through these connectivity matrices, we computed cross-task similarity by measuring the cosine similarity between target task activity patterns and averaging the positive values in the upper-triangular similarity matrix.

Representational similarity analysis in multiple tasks with shared features

We identified a subset of five tasks comprising 44 task conditions that had especially interpretable shared task features. This subset allowed us to identify regions encoding conjunctive combinations of features and assess whether such coding is relevant to observed and predicted values of high dimensionality (representational expansion). We constructed ground truth similarity matrices representing possible similarity structures across task conditions based on shared sensory, cognitive, and motor features, as well as their multiplicative conjunction.

Sensory and motor features were assessed through careful examination of stimuli and expected responses in each task condition (Table 1). Stimuli were categorized hierarchically, including higher-order features (text, images, videos, spatial arrays) and more specific elements (e.g., scenes, faces, objects within images). We constructed a graded theoretical stimulus RSM, coding similarity levels based on shared features: one for same stimulus type 1 (e.g., images), two for same stimulus type 2 (e.g., objects), and three for identical stimuli reused across conditions.

For motor responses, we leveraged the fact that participants practiced the tasks to sufficiently high accuracy (>85%) prior to scanning to create a ground truth motor similarity matrix based on expected motor responses in each task condition [49]. Task conditions with similar motor responses (e.g., specific button presses) were coded as ones, others as zeros.

Cognitive (task demand) similarity was computed for each task by assigning demand levels to each cognitive factor manipulation. Most tasks had three levels (easy, medium, hard), coded as continuous values of one, two, and three respectively, with Stroop task using one for congruent and three for incongruent conditions. Similarity between task conditions was calculated based on difference (due to its continuous nature) as:

$$sim_{cognitive}(cond_A, cond_B) = 3 - |td(cond_A) - td(cond_B)|$$

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

The three matrices were normalized for values to range between zero and one, then multiplied to produce a theoretical similarity conjunction term. In each cortical region, the observed RSM among the 44 task conditions, averaged across subjects, was fit to the four ground truth similarity matrices using ridge regression with 5-fold cross-validation, in each region, in each subject.

Human Connectome Project: neural and behavioral data analysis

We analyzed data from 352 unrelated subjects from the HCP '1200 Subjects' release, selected based on low motion during scanning. We computed representational dimensionality from HCP task data[28,29] following procedures analogous to our primary dataset analysis. After applying the same preprocessing pipeline, we obtained activity estimates for 24 task conditions across seven tasks

using GLM-based multiple regression. Separate regressors were used for each task condition across two identical, independent runs that differed only in phase encoding direction. We constructed cross-validated RSMs by measuring cosine similarity between the two runs and derived dimensionality metrics from these matrices. To quantify the relative compression in association regions compared to sensory and motor regions, we computed a compression ratio:

$$Compression\ Ratio = \frac{(D_{Rep,Sensory} + D_{Rep,Motor})}{2 * D_{Rep,Association}}$$

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

For behavioral analysis, we used a general fluid intelligence composite score following previous work[65]. This score combined the z-scored fluid cognition composite (aggregating performance on Dimensional Change Card Sort, Flanker, Picture Sequence Memory, List Sorting, and Pattern Comparison tests[66,67]) with z-scored correct responses from the Penn Matrix Reasoning Test (a simplified version of Raven's Progressive Matrices). This composite score measures general fluid intelligence—a broad cognitive ability reflecting performance across diverse domains including working memory, cognitive control, and reasoning[48].

References

1. Miller, E. K. & Cohen, J. D. An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience* **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. *Nature neuroscience* **26**, 306–315 (2023).
7. Badre, D., Bhandari, A., Keglovits, H. & Kikumoto, A. The dimensionality of neural representations for control. *Current Opinion in Behavioral Sciences* **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., Woolrich, M. W., Behrens, T. E. & Klein-Flügge, M. C. Projections of non-invasive human recordings into state space show unfolding of spontaneous and over-trained choice. 2020.02.24.962290 Preprint at https://doi.org/10.1101/2020.02.24.962290 (2020).
10. Brincat, S. L., Siegel, M., von Nicolai, C. & Miller, E. K. Gradual progression from sensory to task-related processing in cerebral cortex. *Proceedings of the National Academy of Sciences* **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. Preprint at http://arxiv.org/abs/2209.07431 (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. *Proceedings of the National Academy of Sciences* **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. *bioRxiv* 2023.06.09.544428 (2023) doi:10.1101/2023.06.09.544428.
17. Fusi, S., Miller, E. K. & Rigotti, M. Why neurons mix: high dimensionality for higher cognition. *Current Opinion in Neurobiology* **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 in Cognitive Sciences* **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 Computational Biology* **20**, e1012507 (2024).
24. Ito, T., Hearne, L., Mill, R., Cocuzza, C. & Cole, M. W. Discovering the Computational Relevance of Brain Network Organization. *Trends in Cognitive Sciences* **24**, 25–38 (2020).

25. Cole, M. W. Network modeling: The explanatory power of activity flow models of brain function. in *Computational and Network Modeling of Neuroimaging Data* 81–117 (Elsevier, 2024). doi:10.1016/B978-0-443-13480-7.00004-1.
26. Cole, M. W. Cognitive flexibility as the shifting of brain network flows by flexible neural representations. *Current Opinion in Behavioral Sciences* **57**, 101384 (2024).
27. Ji, J. L. *et al.* Mapping the human brain's cortical-subcortical functional network organization. *NeuroImage* **185**, 35–57 (2019).
28. Barch, D. M. *et al.* Function in the human connectome: Task-fMRI and individual differences in behavior. *NeuroImage* **80**, 169–189 (2013).
29. Van Essen, D. C. *et al.* The WU-Minn Human Connectome Project: An overview. *NeuroImage* **80**, 62–79 (2013).
30. 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).
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 Computational Biology* **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 Computational Biology* **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 (2024).

36. Buzsáki, G. *The Brain from inside Out*. xvii, 441 (Oxford University Press, New York, NY, US, 2019). doi:10.1093/oso/9780190905385.001.0001.

37. Cole, M. W., Ito, T., Cocuzza, C. & Sanchez-Romero, R. The Functional Relevance of Task-State Functional Connectivity. *J. Neurosci.* **41**, 2684–2702 (2021).

38. 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).

39. Higley, M. J. & Cardin, J. A. Spatiotemporal dynamics in large-scale cortical networks. *Current Opinion in Neurobiology* **77**, 102627 (2022).

40. Wang, X.-J. Theory of the Multiregional Neocortex: Large-Scale Neural Dynamics and Distributed Cognition. *Annual Review of Neuroscience* **45**, 533–560 (2022).

41. 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).

42. Semedo, J. D. *et al.* Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nat Commun* **13**, 1099 (2022).

43. Srinath, R., Ruff, D. A. & Cohen, M. R. Attention improves information flow between neuronal populations without changing the communication subspace. *Current Biology* **31**, 5299-5313.e4 (2021).

44. Stoll, F. M. & Rudebeck, P. H. Decision-making shapes dynamic inter-areal communication within macaque ventral frontal cortex. *Current Biology* **34**, 4526-4538.e5 (2024).

45. MacDowell, C. J. *et al.* Multiplexed subspaces route neural activity across brain-wide networks. *Nat Commun* **16**, 3359 (2025).

46. Anderson, M. L. Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences* **33**, 245–266 (2010).

47. Anderson, M. L. Evolution of Cognitive Function via Redeployment of Brain Areas. *Neuroscientist* **13**, 13–21 (2007).

48. Gray, J. R., Chabris, C. F. & Braver, T. S. Neural mechanisms of general fluid intelligence. *Nat Neurosci* **6**, 316–322 (2003).

49. 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).

50. Glasser, M. F. *et al.* The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* **80**, 105–124 (2013).

51. Ji, J. L. *et al.* QuNex–A Scalable Platform for Integrative Multi-Modal Neuroimaging Data Processing and Analysis. *bioRxiv* (2022).

52. Rissman, J., Gazzaley, A. & D'Esposito, M. Measuring functional connectivity during distinct stages of a cognitive task. *NeuroImage* **23**, 752–763 (2004).

53. Friston, K. J. *et al.* Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping* **2**, 189–210 (1994).

54. Lindquist, M. A., Geuter, S., Wager, T. D. & Caffo, B. S. Modular preprocessing pipelines can reintroduce artifacts into fMRI data. *Human Brain Mapping* **40**, 2358–2376 (2019).

55. Glasser, M. F. *et al.* A multi-modal parcellation of human cerebral cortex. *Nature* **536**, 171–178 (2016).

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. 2023.09.16.558065 Preprint at https://doi.org/10.1101/2023.09.16.558065 (2023).
57. 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).
58. Mill, R. D. *et al.* Network modeling of dynamic brain interactions predicts emergence of neural information that supports human cognitive behavior. *PLOS Biology* **20**, e3001686 (2022).
59. Gao, P. *et al.* A theory of multineuronal dimensionality, dynamics and measurement. 214262 Preprint at https://doi.org/10.1101/214262 (2017).
60. Jeganathan, J. *et al.* Spurious correlations in surface-based functional brain imaging. *Imaging Neuroscience* **3**, imag_a_00478 (2025).
61. Burt, J. B., Helmer, M., Shinn, M., Anticevic, A. & Murray, J. D. Generative modeling of brain maps with spatial autocorrelation. *NeuroImage* **220**, 117038 (2020).
62. 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 U S A* **115**, E6356–E6365 (2018).
63. Margulies, D. S. *et al.* Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences* **113**, 12574–12579 (2016).

64. 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).
65. Podvalny, E., Sanchez-Romero, R. & Cole, M. W. Functionality of arousal-regulating brain circuitry at rest predicts human cognitive abilities. *Cerebral Cortex* **34**, bhae192 (2024).
66. 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. *Journal of the International Neuropsychological Society* **25**, 857–867 (2019).
67. Heaton, R. K. *et al.* Reliability and Validity of Composite Scores from the NIH Toolbox Cognition Battery in Adults. *Journal of the International Neuropsychological Society* **20**, 588–598 (2014).

Supplementary Figures

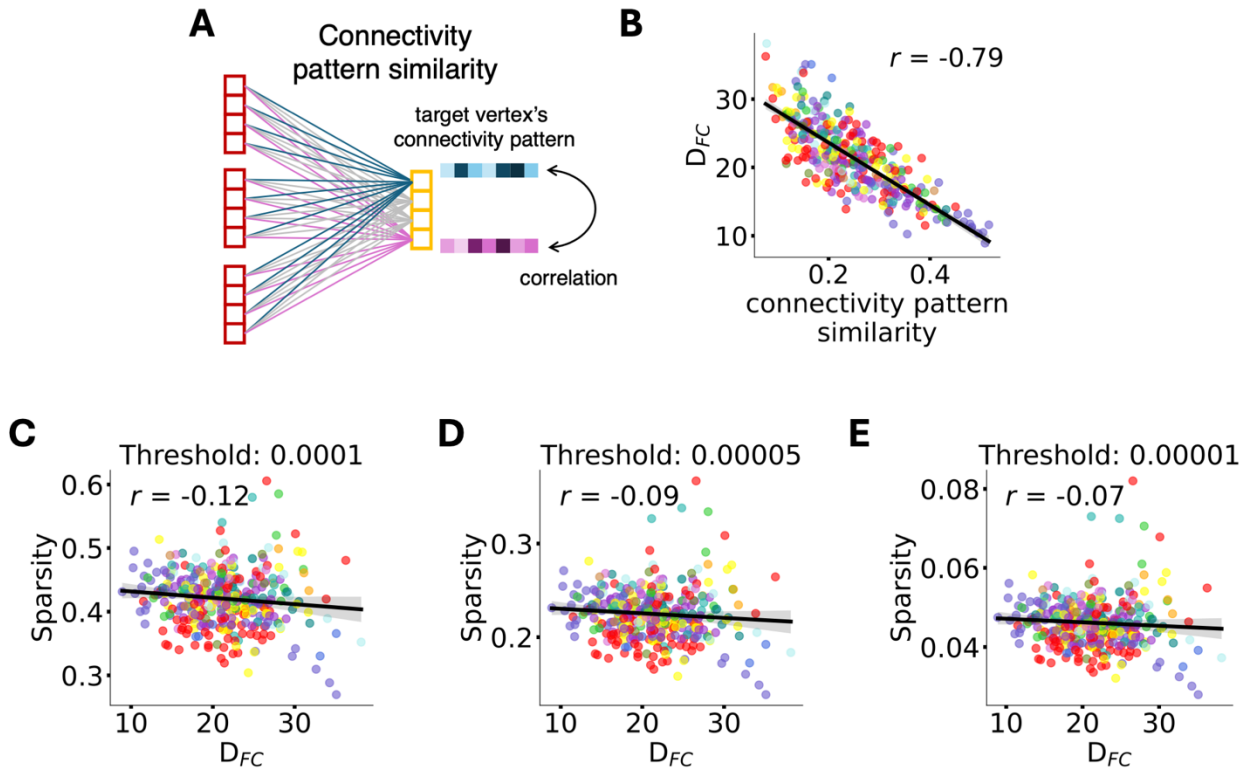

Supplementary Figure S1: Connectivity dimensionality reflects similarity of connectivity patterns, not sparsity. A) Schematic of connectivity pattern similarity. Each target vertex (yellow and purple boxes on right) receives weighted inputs from all source vertices (red boxes on left), forming a connectivity 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 (DFC) 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.

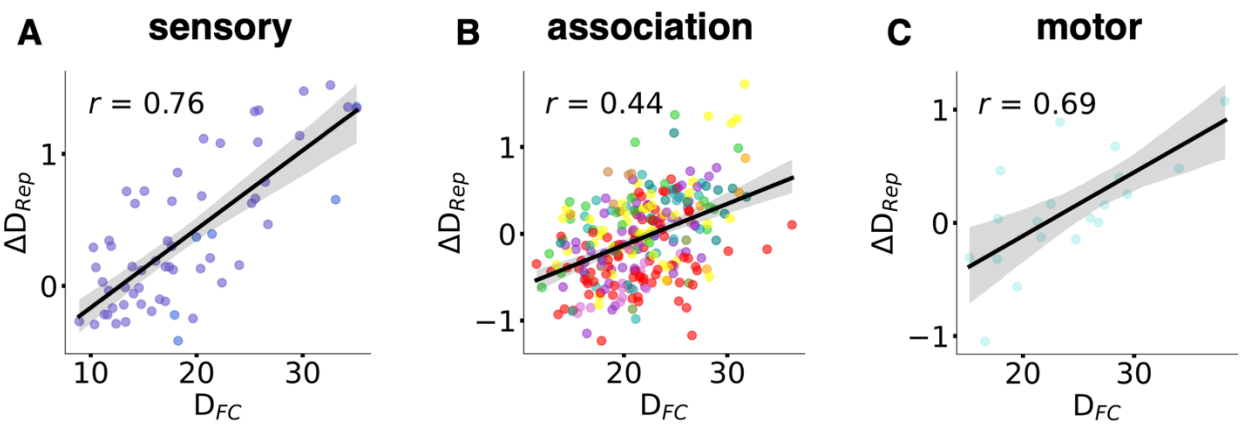

Supplementary Figure S2: Connectivity dimensionality predicts compression/expansion within sensory, association, and motor systems. The relationship between connectivity dimensionality (DFC) and representational compression/expansion (ΔDRep) holds consistently across all three major cortical divisions. **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.

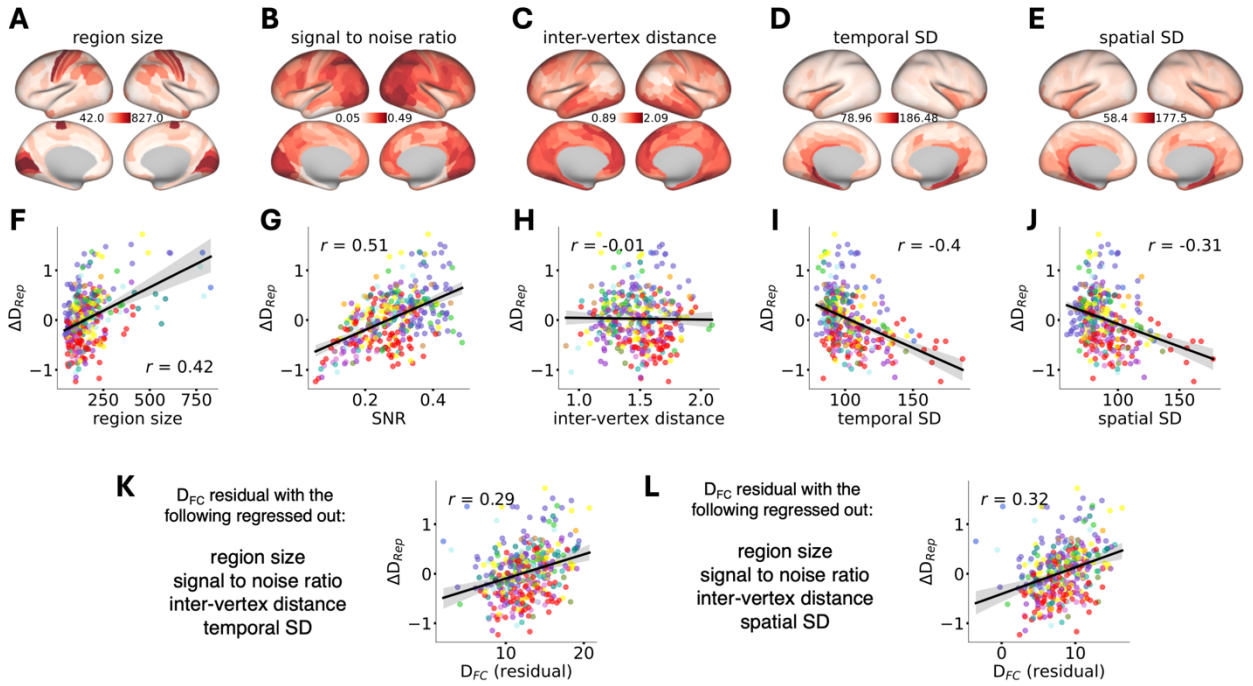

Supplementary Figure S3: Relationship between connectivity dimensionality and representational compression/expansion survives control for potential confounds. A-E) Spatial distribution of potential confounding variables across cortex. Brain maps showing five region-specific properties that vary across cortex and could potentially explain both connectivity and representational dimensionality. **A) Region size** (vertex count; range 42-827). **B) Signal-to-noise ratio** (mean task condition reliability; range 0.05-0.49). **C) Inter-vertex distance** (mean geodesic distance on cortical surface; range 0.89-2.09 mm). **D) Temporal variance** (standard deviation of resting-state BOLD signal across time, averaged across vertices; range 78.96-186.48). **E) Spatial variance** (standard deviation across vertices at each timepoint, averaged across time; range 58.4-177.5). **F-J) Individual confound relationships with compression/expansion.** Each panel shows correlation between one confound and ΔDRep across all 360 regions. Region size (r = 0.42), SNR (r = 0.51), temporal SD (r = -0.40), and spatial SD (r = -0.31) show significant relationships; inter-vertex distance shows no relationship (r = -0.01). These correlations demonstrate that several region-specific properties do vary with compression/expansion, necessitating control analyses. **K-L) Central result remains significant after controlling for confounds.** To test whether connectivity dimensionality independently predicts compression/expansion, we regressed confounds from DFC and examined residuals. **K)** After regressing out region size, SNR, inter-vertex distance, and temporal SD from DFC, the correlation between DFC residuals and ΔDRep remains significant ($r = 0.29$, $p < 0.001$, spatial-autocorrelation preserving permutation test). **L)** After regressing out region size, SNR, inter-vertex distance, and spatial SD from DFC, the correlation remains significant ($r = 0.32$, $p < 0.001$). Temporal and spatial SD were analyzed separately (panels K vs. L) because they are highly correlated with each other. These results demonstrate that connectivity dimensionality predicts representational transformations independent of region size, signal quality, geometry, and activity variance properties.

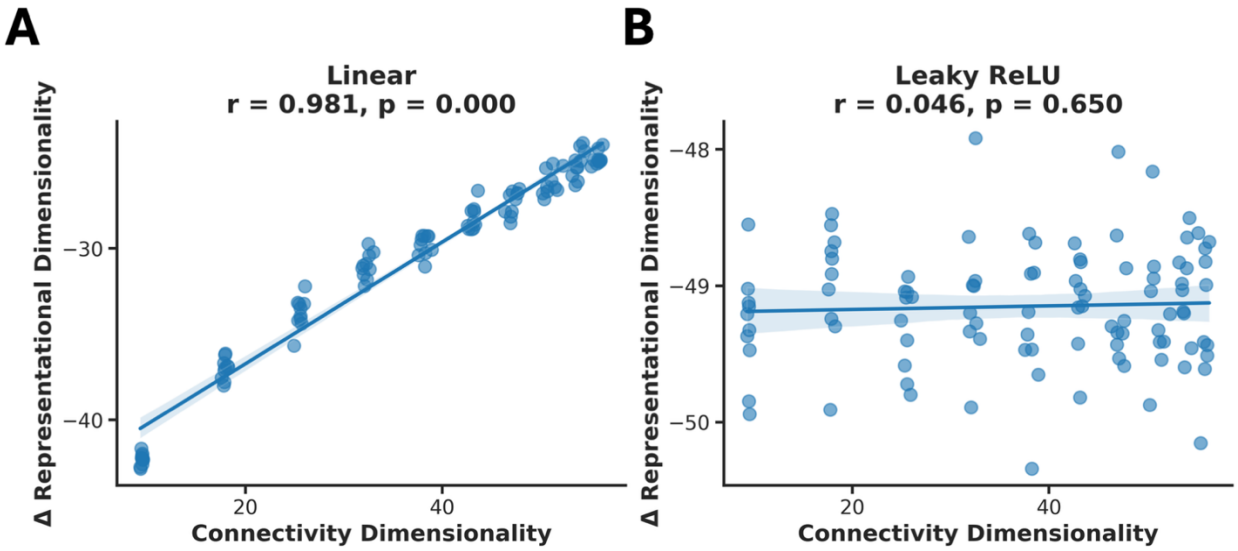

Supplementary Figure S4: Non-linear local computations can decouple connectivity dimensionality from representational dimensionality change. To test whether the relationship between connectivity dimensionality and representational dimensionality change is mathematically inevitable, we simulated source 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 the 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 connectivity dimensionality level, across 9 dimensionality levels). Lower connectivity dimensionality produces compression (negative Δ representational dimensionality); higher connectivity dimensionality produces less compression or expansion. This demonstrates that under purely linear transformations, connectivity dimensionality strongly constrains representational dimensionality. **B) Non-linear local computations eliminate the relationship.** We performed the following simulation to simulate nonlinear transformation at the target region level: connectivity-projected activity additionally undergoes a leaky ReLU non-linearity (alpha = 0.01) with region-specific and task-dependent bias terms before representational dimensionality is computed. Specifically, target vertex activity $y = f(T \cdot W + b_{region} + b_{task})$, 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, $b_{region} \sim N(400, 50)$, and task-specific biases with smaller variance on top, $b_{task} \sim N(0, 25)$. 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. This exaggerated scenario demonstrates that the relationship between connectivity dimensionality and representational dimensionality change is not mathematically inevitable — strong non-linear local computations in target regions can decouple the two. The empirical observation that connectivity dimensionality robustly predicts representational transformations across cortex (Fig. 2E, r = 0.44) therefore suggests that while local computations

certainly contribute, connectivity geometry provides systematic constraints that shape representational transformations throughout the brain.

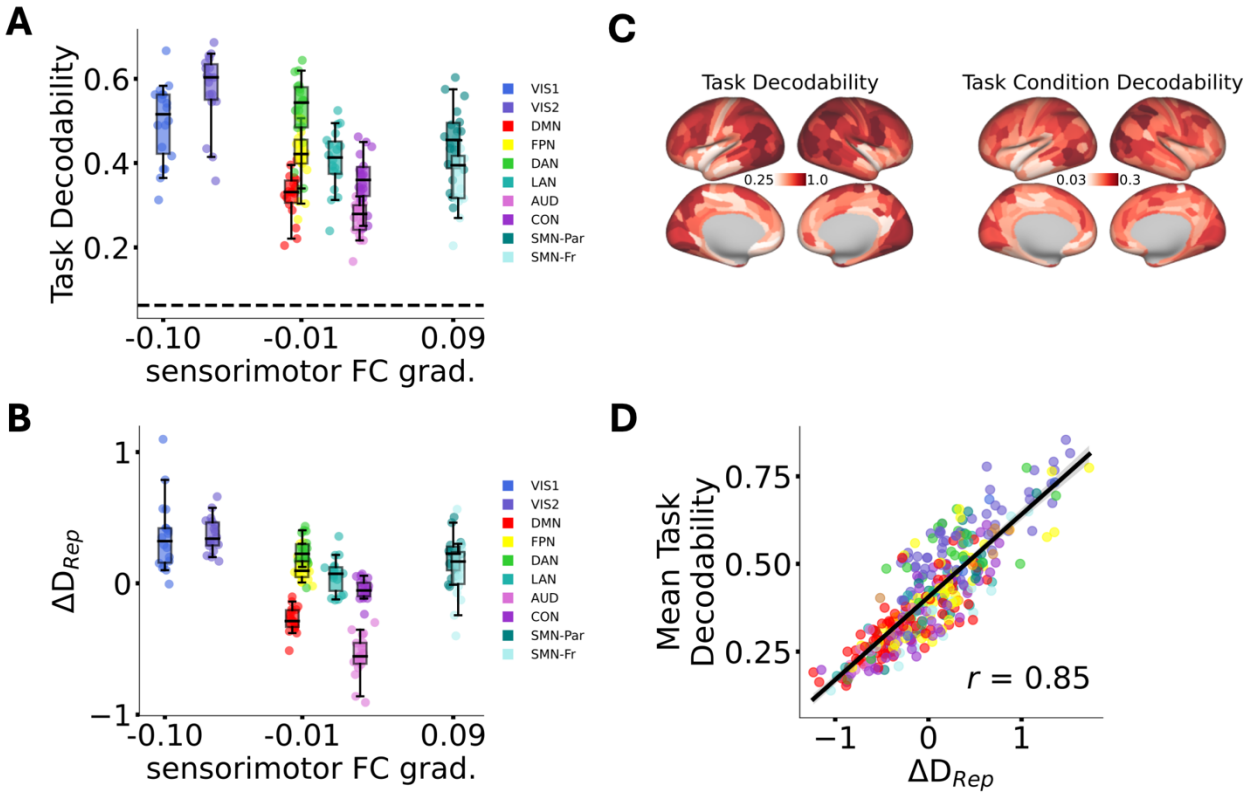

Supplementary Figure S5: Representational compression corresponds to reduced task decodability across cortical networks. A) Task decodability varies across cortical networks. Decodability of 16 individual tasks from representational similarity matrices (2-fold minimum distance classification based on 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 ΔDRep 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. This tight relationship validates that representational dimensionality changes reflect functionally meaningful transformations in how task information is encoded.

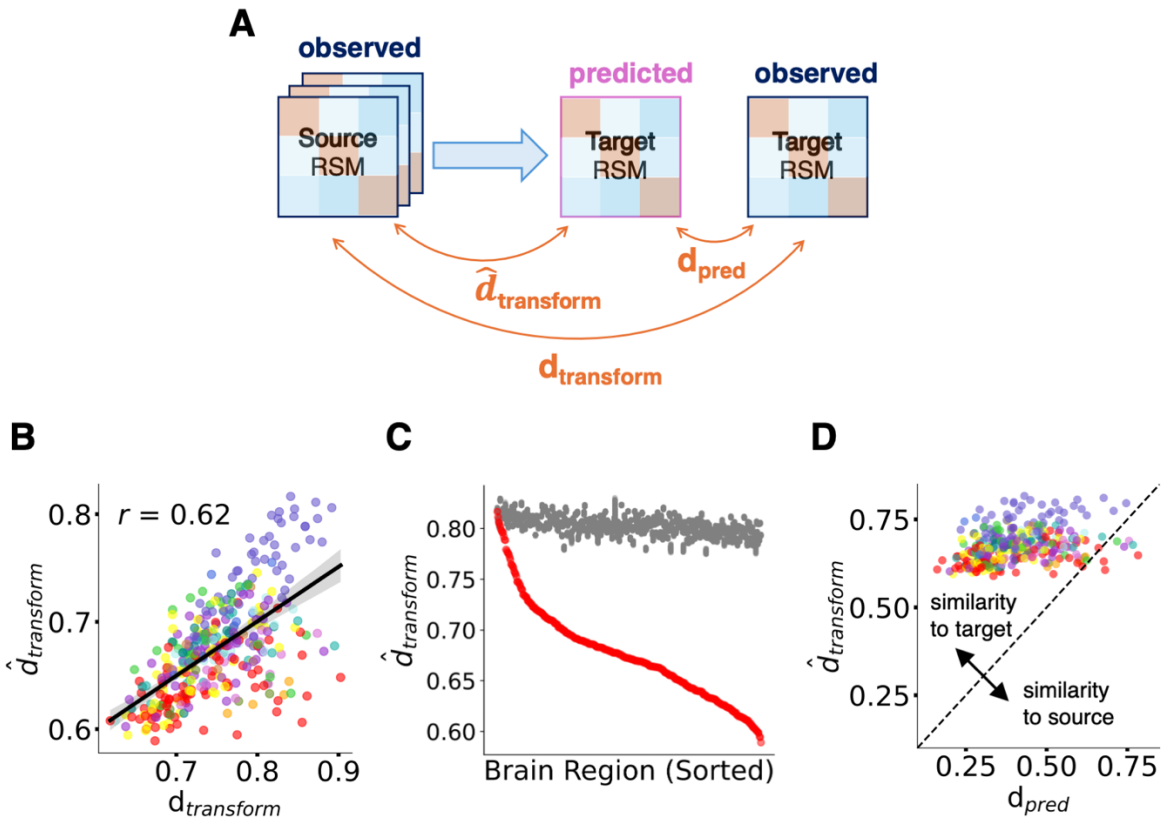

Supplementary Figure S6: Activity flow over connectivity generates genuine representational transformations, not simple transfer. A) Measuring transformation distance. To test whether activity flow implements true representational transformations rather than simply transferring source representations to targets, we computed three distances using cosine distance (1 - cosine similarity) between RSMs. $\hat{\mathbf{d}}_{\mathbf{transform}}$: distance between observed source RSM and predicted target RSM, quantifying how much the model transforms representations. $\mathbf{d}_{\mathbf{transform}}$: distance between observed source RSM and observed target RSM, quantifying the actual transformation. $\mathbf{d}_{\mathbf{pred}}$: distance between predicted target RSM and observed target RSM, quantifying prediction error. **B) Predicted transformation distance correlates with observed transformation distance.** Correlation between $\hat{d}_{transform}$ (model-predicted source-to-target transformation) and $d_{transform}$ (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 ($\hat{d}_{transform}$, 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 ($\hat{d}_{transform}$, similarity to source). Points below the diagonal (339/360 regions, 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.

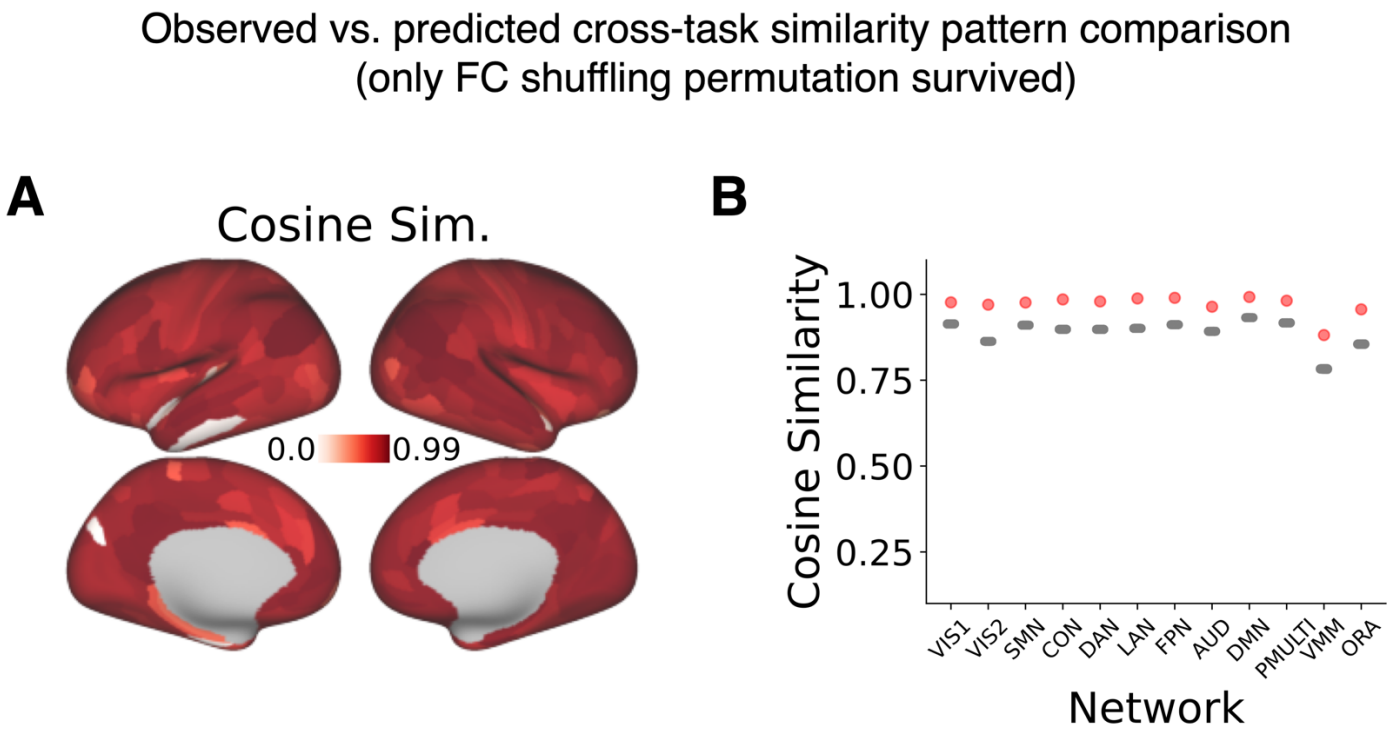

Supplementary Figure S7: Cross-task similarity predictions at the level of individual regions and large-scale brain networks. (A) Cosine similarity between vectorized cross-task similarity values from observed (empirically-estimated) and from connectivity-predicted estimates. Regions with no values did 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 [27]: 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.

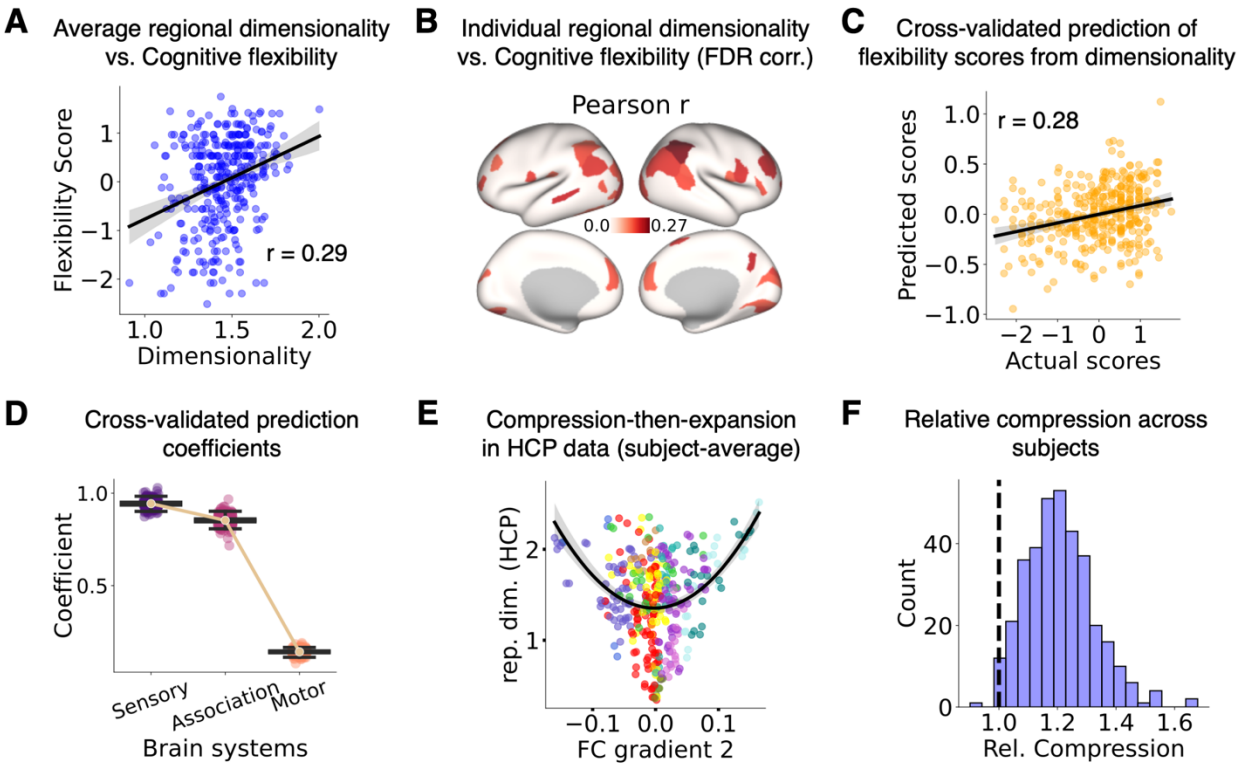

Supplementary Figure S8: Representational dimensionality is predictive of behavioral measures of cognitive flexibility. A) Average representational dimensionality across regions was positively correlated with cognitive flexibility scores across subjects. B) Similar relationship in individual regions (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.