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:
All relevant data will be shared on open access database (like OpenNeuro). 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

The human brain's ability to perform a wide variety of possible tasks raises fundamental questions about how distributed neural circuits transform and integrate diverse task-related information. Using fMRI data from participants performing 16 diverse tasks and at rest, we investigated how intrinsic connectivity patterns shape task representations across the human cortex. We found that cross-region connectivity dimensionality strongly predicts representational transformations between brain regions: low-dimensional connectivity corresponds to representational compression, while high-dimensional connectivity corresponds to expansion. Modeling activity flow over these connections, we determined that task-evoked activity flows over intrinsic connections generate diverse representational geometries along the cortical hierarchy. Critically, in nearly all brain regions, connectivity-based transformations produced activity patterns that more closely matched their targets than their source representations, demonstrating that connectivity patterns actively transform neural representations. Additionally, we found that regions with lower-dimensional connectivity show stronger cross-task similarities, indicative of shared latent task features that generalize neural processing across tasks. In contrast, we found that regions with higher-dimensional connectivity exhibit strong conjunctive coding of task variables, enabling representations for task-specific scenarios. These findings demonstrate how the network geometry of the brain's intrinsic connectivity architecture systematically shapes the transformation of task representations across the cortex to support performance of diverse tasks.

Significance

How can human brains implement such a wide variety of tasks? We tested the possibility that performing multiple tasks depends on how task information is transformed between brain regions. Using fMRI data from participants performing multiple tasks, we found that the dimensionality of functional connectivity between regions systematically predicts how task representations are transformed: low-dimensional connectivity leads to compression of representations, while high-dimensional connectivity leads to expansion. A computational model based on these connectivity patterns recapitulated multitask representations, demonstrating the importance of connectivity-based compression into abstract low-dimensional representations that are reused across tasks. This work demonstrates how the brain's intrinsic functional architecture shapes the processing of multitask information across the cortex, providing insights into the neural basis of cognitive flexibility.

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). Functionally, this suggests that sensory information is initially compressed (generalized) across task conditions to facilitate transferrable representations, before expanding into representations that specify task-specific actions. This compression-then-expansion pattern may be essential for the brain's ability to flexibly perform diverse tasks while maintaining the ability to generalize learning across similar conditions. However, despite these observations, the specific neural circuit mechanisms that form these highly structured yet differentiated neural representations key for flexible cognition remain poorly understood.

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 (7), while others observed coding schemes supporting selective collapse of task-irrelevant information (4, 8, 9) or maintaining abstract representations (10, 11). Other studies point toward representations encoding conjunctive combinations of variables in a non-linear mixed-selective manner (12). Consequently, some brain regions exhibit low-dimensional representations, while others represent task conditions in high-dimensional spaces (13). These variations in representational geometry suggest differing computational functions across brain regions, often balancing between separability and generalizability of task representations (14).

Though these findings have contributed to our understanding of the representational geometry of diverse task sets across the cortical hierarchy, the circuit mechanisms that generate and constrain these representational patterns are only poorly understood. Prior studies have demonstrated that the production of task-specific brain activity can be predicted using estimates of brain network connectivity(15, 16). For example, localized brain function, such as category selectivity, can be predicted through distributed brain network connectivity(17). Building on these ideas, we hypothesize a prominent role of inter-regional connectivity in generating and constraining representational geometries across the cortical hierarchy. Providing a circuit-based explanation for representational differentiation across the brain would provide evidence that the brain’s local representational geometry and capacity is governed by distributed circuit-based computations.

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 (18, 19). 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 empirically-estimated intrinsic (resting-state) connectivity patterns to a target region constrains the dimensionality of the task representations (in the activation space). In particular, we hypothesized that a low-dimensional connectivity pattern to a target region would compress its task representations, while a high-dimensional connectivity pattern would expand its task representations (Fig. 1A). We further hypothesized that connectivity-based compressions emphasize shared features across tasks, while expansions support task-specific information. Establishing a link between connectivity dimensionality and representational transformation would provide a mechanistic explanation for the distinct computational roles of brain regions during flexible task processing.

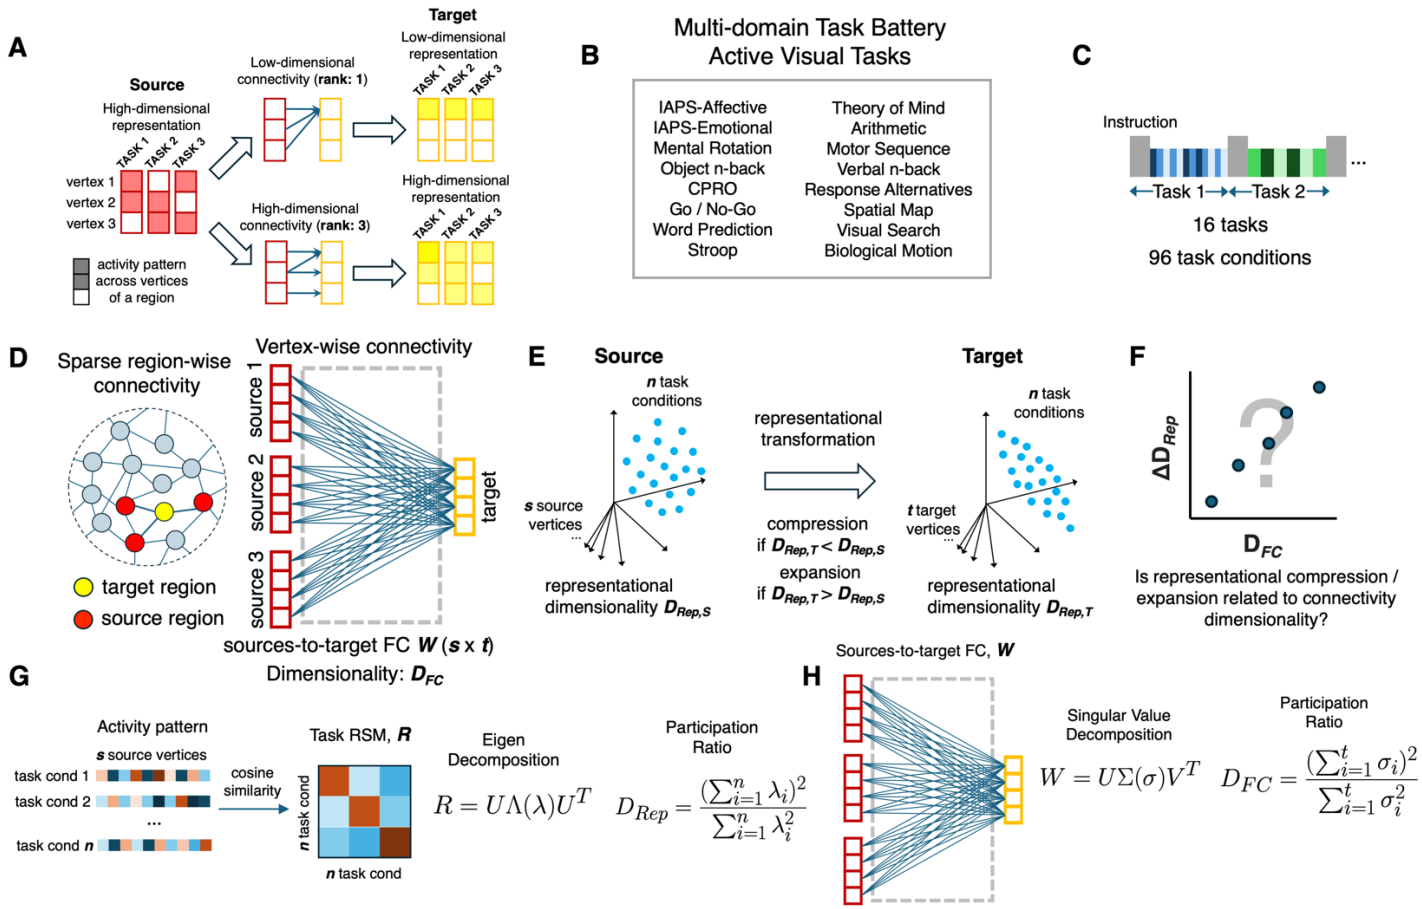

Figure 1: **Testing the role of connectivity patterns in generating representational transformations**. A) Using activity flow modeling to investigate the relationship between connectivity dimensionality and representational compression/expansion. If task activity 'flowing' over intrinsic connectivity explains activity in other regions, we hypothesize that the dimensionality of the connectivity projection of source region to target region predicts the dimensionality of representations in target region. B) List of 16 active visual tasks used in the study, a subset from (20), C) Task structure: Each task was comprised of a 5-second instruction period followed by 30 seconds of the task. Different hued bars represent different task conditions within each task that are either interleaved or blocked by design of the original study. A total of 96 task conditions (excluding instructions) were estimated using a regularized general linear model. D) We used a two-step, regularized approach to modeling intrinsic (resting-state) FC to estimate projection maps between pairs of regions (see Methods for further details). E) Schematic of representational transformation between pairs of regions. Task conditions are embedded in the space of each region's vertices. An increase in dimensionality from sources to target represents expansion, whereas a decrease represents compression. F) Schematic of the hypothesis: a positive relationship between representational expansion and connectivity dimensionality, $\mathbf{D}_{FC}$. G) Representational dimensionality was measured using participation ratio of the eigenvalue spectrum of task activity patterns. H) Connectivity dimensionality, measured for each target region, was measured in a similar manner using participation ratio of the singular value spectrum of the vertex-level intrinsic connectivity values between the sources and target. (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. 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 cognitive flexibility. 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 posits that a region's task-evoked activity (a_t) is determined by the connectivity-specified flow ($a_s c_{ts}$) of task-evoked activity from other regions ($a_t = f(\sum a_s c_{ts})$) (19). Consistent with this, our hypothesis predicted that there should be a relationship between cross-region representational dimensionality (compression or expansion) and connectivity dimensionality. We expected to observe that low connectivity dimensionality between regions compressed task representations, and that high connectivity dimensionality expands task representations.

We examined representational transformations between connected regions using a two-step approach to FC estimation (Fig. 1D; 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 and principal components regression. 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) (Fig. 1G). The participation ratio provides a continuous measure of dimensionality by quantifying how quickly the eigenvalue spectrum decays – a slower decay indicates higher dimensionality, reflecting more distinct representational components. Intuitively, more compressed activity patterns could be described using fewer components, meaning they were simpler and likely changed less across task conditions. We defined compression as the difference between the representational dimensionality of the target region and the average representational

dimensionality of its source regions (negative 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 (Fig. 1H). Just as in representational dimensionality, this provides a continuous measure where higher values indicate more distinct components in the connectivity structure between source and target regions.

Using these metrics, we tested whether changes in representational dimensionality from sources to each target correlated with the dimensionality of their FC projection. Building on prior work(6), we plotted transformations along an FC gradient from sensory cortex to association cortex to motor cortex (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 both task-evoked activity and intrinsic connectivity dimensionality along the FC gradient (Fig. 2B,C). As targets, sensory regions typically showed expansion, association regions exhibited compression, and motor regions expansion, forming a U-shaped trajectory of representational compression/expansion. To assess the generalizability of this compression-then-expansion pattern, we analyzed an independent dataset from the Human Connectome Project(21, 22) (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 (Fig. S1E). 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 (Fig. S1F). 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: $t(17) = 16.94$, $p = 4.44e\text{-}12$, Fig. 2D). This positive relationship was consistent across sensory, association, and motor systems (all $p < 0.001$, Fig. S2), demonstrating that the relationship between representational transformations and connectivity is highly general.

This relationship between connectivity and representational dimensionality is non-trivial for several reasons. First, connectivity and representational dimensionality are based on fundamentally distinct sources of variance: connectivity dimensionality is based on cross-region covariance within a single brain state (resting state), while representational dimensionality is based on within-region variance between diverse states (96 task conditions across 16 tasks). 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 theoretically could have supported representational expansion (rather than compression) through non-linear transformations. Finally, high-dimensional connectivity might have simply reflected noisy or random connectivity patterns,

bearing no relationship to representational dimensionality. Given that dimensionality scales with the size of the neural space (e.g., number of target region vertices), we controlled for region size by using region size as a covariate and retesting the relationship. The relationship remained significant after using region size as a covariate (r = 0.28, group t-test against zero: t(17) = 7.98, p = 3.80e-07). Together, these results support our hypothesis that as FC becomes more constrained (lower dimensional), neural representations tend to become more compressed.

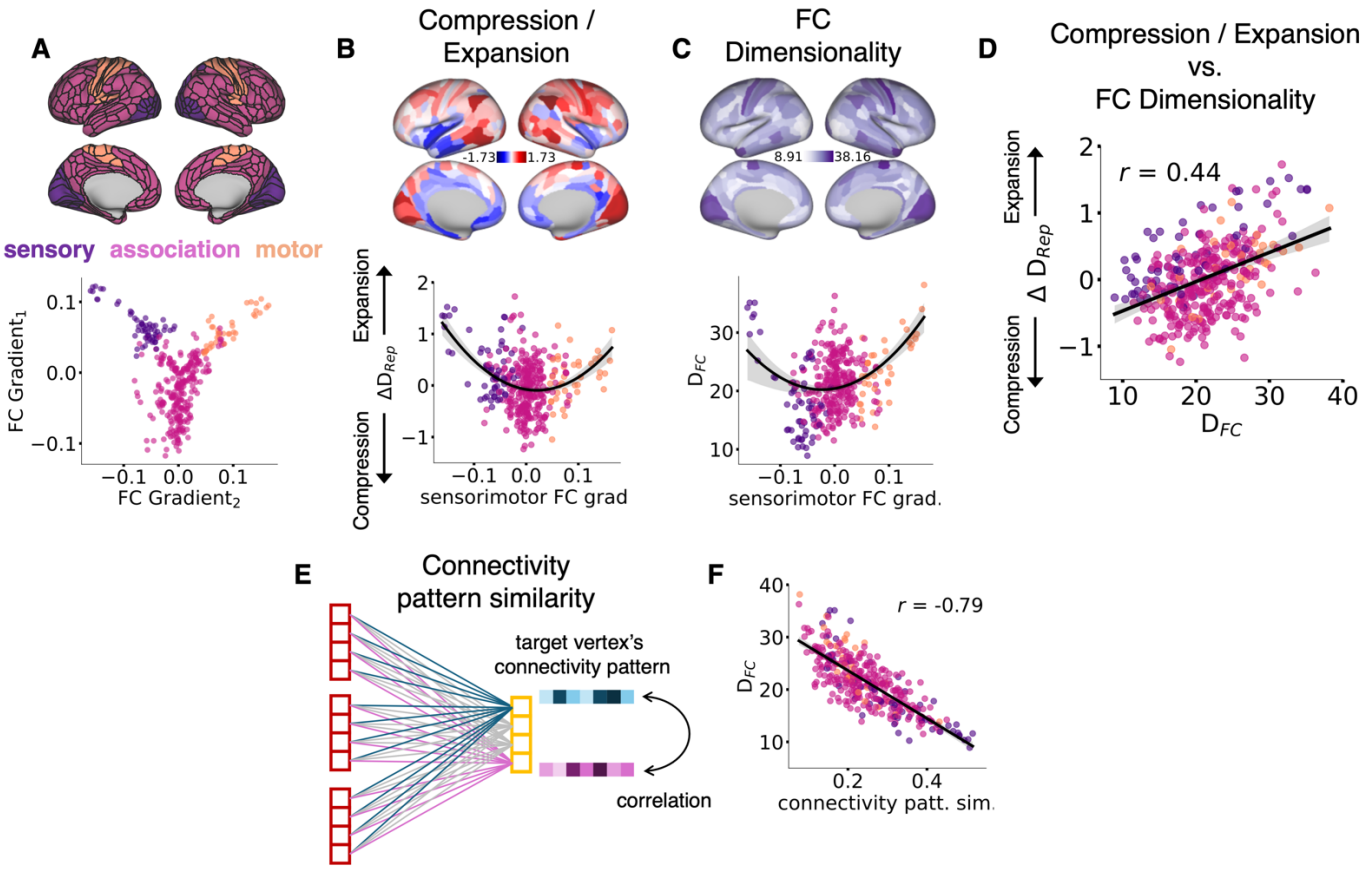

Figure 2: The dimensionality of a region's representations is determined by its connectivity dimensionality. A) Discrete and continuous divisions between sensory, association, and motor brain systems. (Top) Discrete divisions were obtained from a functional network atlas(23). (Bottom) Continuous FC gradients were derived from performing principal components analysis on sparse parcel-level connectivity graph. The second principal component runs from sensory to association to motor regions, a proxy for cortical hierarchy, referred to as sensorimotor FC gradient henceforth. B) (Top) Brain plot showing degree of compression/expansion of representations in each region compared to the corresponding source regions. (Bottom) Depicting the regions along the sensorimotor FC gradient reveals the previously reported compression then expansion pattern along the cortical hierarchy. C) (Top) Brain plot showing the connectivity dimensionality of each region as a target. (Bottom) Connectivity dimensionality also shows a U-shaped pattern along the cortical hierarchy. D) Higher values of connectivity dimensionality correspond with representational expansion whereas lower values with representational compression. E) Connectivity dimensionality can be understood in terms of similarity between connectivity patterns of the target vertices. F) The greater

the similarity between target vertices' connectivity patterns, lower is the connectivity dimensionality. (End of Figure 2 caption)

To improve the interpretability of our abstract measure of connectivity dimensionality in terms of specific connectivity properties, we examined the similarity (correlation) among connectivity patterns of target vertices (Fig. 2E). While there is a straightforward intuition for how representational dimensionality relates to generalization across tasks (lower dimensionality implies more cross-task generalization), the concept of connectivity dimensionality benefits from a concrete illustration. Mathematically, the dimensionality of a connectivity matrix is determined by the similarity structure among target vertices' connectivity patterns. Note that the mean value of the connectivity pattern similarity is independent of the number of target vertices (region size), precluding its influence on connectivity pattern similarity explaining connectivity dimensionality. Analysis revealed a strong negative relationship between connectivity dimensionality and mean connectivity pattern similarity ($r = -0.79$, group t-test against zero: $t(17) = -58.76$, $p = 4.44e-21$, Fig. 2F), consistent across sensory ($r = -0.86$), association ($r = -0.74$), and motor ($r = -0.84$) systems (all $p < 0.001$). These results suggest that connectivity controls the degree of cross-region representational compression via the similarity (versus diversity) of large-scale connectivity patterns.

Cross-region activity flows transform representational geometries

While our analyses revealed a robust correlation between connectivity dimensionality and representational compression, our main hypothesis remains untested: that low connectivity dimensionality *generates* representational compression. We tested this hypothesis using activity flow modeling (15, 16), testing whether simulated flows over empirical connectivity patterns can generate 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.

We projected activity estimates from source regions to target regions using vertex-resolution intrinsic FC, then compared model-generated activity with observed data (Fig. 3A). Note that model weights consisted of task-independent (resting-state) FC values, substantially reducing any chance of overfitting to task-evoked representations. The model-generated activity patterns were significantly predictive in all regions. Pearson correlation (averaged across subjects) between observed and predicted activity patterns (averaged across task conditions and sessions): the range of r across regions was 0.16 to 0.83, while t-tests against zero had a range of FDR-corrected p-values (p_{FDR}) across regions of 6.17e-14 to 0.041 (Fig. 3B (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 ($p_{permutation} < 0.001$ for all target regions, Fig. 3B (bottom)). These results go beyond the prior correlational analyses in demonstrating the hypothesized ability of FC patterns to generate empirically observed task representations.

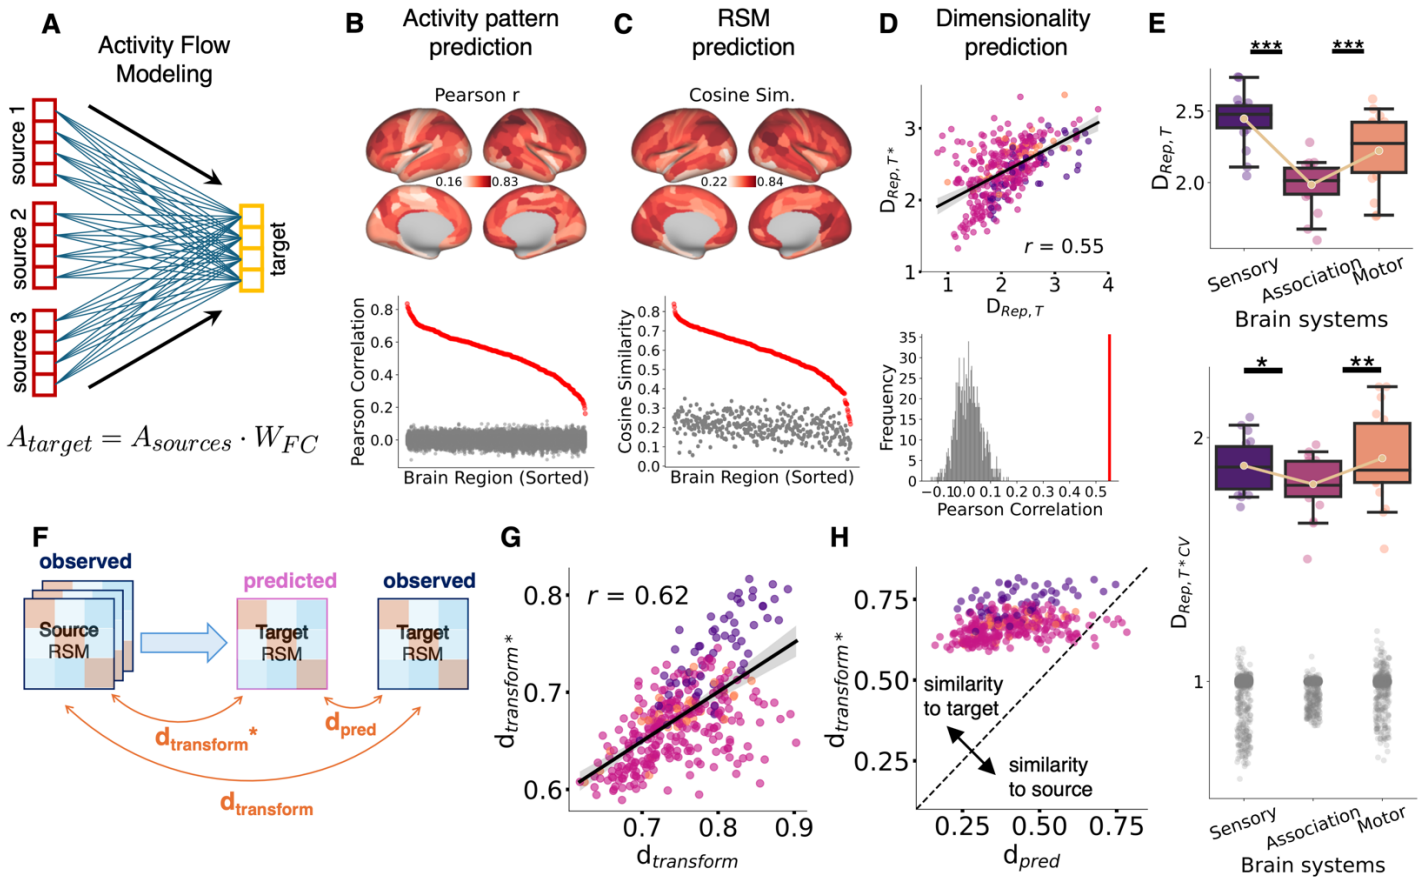

Figure 3: Activity flows over intrinsic connectivity generate observed aspects of representational geometry and transformation. A) Schematic of activity flow modeling approach. Vertex-wise activity vectors for each task condition were projected from source regions to target region over the empirically-estimated connectivity graph. B) (Top) Pearson correlation between observed and model-predicted activity patterns averaged across subjects, sessions and task conditions. (Bottom) Activity pattern predictions remain significant against null distribution generated by random shuffling of sources-to-target connectivity values C) (Top) Cosine similarity between the observed and model-predicted RSMs (averaged across subjects). (Bottom) Significance against null distribution of RSM prediction values generated by FC shuffling D) (Top) Predicted values of representational dimensionality (averaged across subjects) also show significant correspondence with the observed values. (Bottom) Observed Pearson correlation against null distribution of correlation values generated from FC shuffling E) (Top) Compression-then-expansion pattern of representational dimensionality across sensory, association and motor brain systems (averaged across regions in each subject). (Bottom) Dimensionality estimated from cross-validating between observed and model-predicted activity patterns reveal a similar compression-then-expansion pattern, against null values of dimensionality estimated via FC shuffling. F) Schematic of measures for representational transformation (cosine distance, i.e., 1-cosine) from sources to target. G) Correspondence between observed vs. predicted transformation distance. H) Comparison of predicted transformation distance against prediction distance, showing that predicted targets are closer to the observed targets compared to observed sources in most of the brain regions. (End of Figure 3 caption)

Next, we tested the central hypothesis as to how these generated activity patterns relate to representational geometry, particularly the similarity structure between different task conditions

as 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. 3C (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 $p_{permutation} < 0.001$, Fig. 3C (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. 3D (top); comparison to FC-shuffled null distribution, $p_{permutation} < 0.001$). 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 $p<0.0001$).

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. Notably, 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). Given this large expansion of the number of included task conditions, we began by replicating this compression-then-expansion pattern in our data's task condition split (paired t-tests against zero: sensory vs. association regions: $t(17) = 12.065$, $p = 9.26e-10$, and association vs. motor regions: $t(17) = 6.03$, $p = 1.34e-05$, Fig. 3E (top)). 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 (paired t-test against zero: sensory vs. association: $t(17) = 2.465$, $p = 0.0245$, association vs. motor regions: $t(17) = 3.13$, $p = 0.006$, Fig. 3E (bottom)). This finding demonstrated that our model not only predicted individual regional representations but also captured the pattern of compression-then-expansion of representations across the cortical hierarchy. 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 true 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 target 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 ($d_{transform*}$) using cosine distance metric (1-cosine) (Fig. 3F). 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 ($p_{permutation} < 0.001$; Fig. S3). This showed that structure present in the connectivity patterns yielded significantly lower transformation distances compared to random transformations afforded by shuffled FC. Second, we validated our distance measure by demonstrating that connectivity-based transformation distances strongly correlated with observed source-to-target distances ($d_{transform*}$ vs. $d_{transform}$ across regions: $r = 0.62$, $t(17) = 19.60$, $p = 4.18e-13$; Fig. 3G).

Finally, and most critically, 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 ($p_{FDR} < 0.0001$; Fig. 3H). Supporting our main hypothesis, these results suggest empirical connectivity patterns play a fundamental role in transforming representational geometries across brain regions.

Low-dimensional connectivity shapes cross-task representational similarity

We saw substantial reductions in FC dimensionality in the transition from sensory cortex to association cortex (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 diverse 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 similar) contexts. We computed mean cosine similarities between task

condition combinations for each task pair (among the 16 tasks), including only reliably 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.

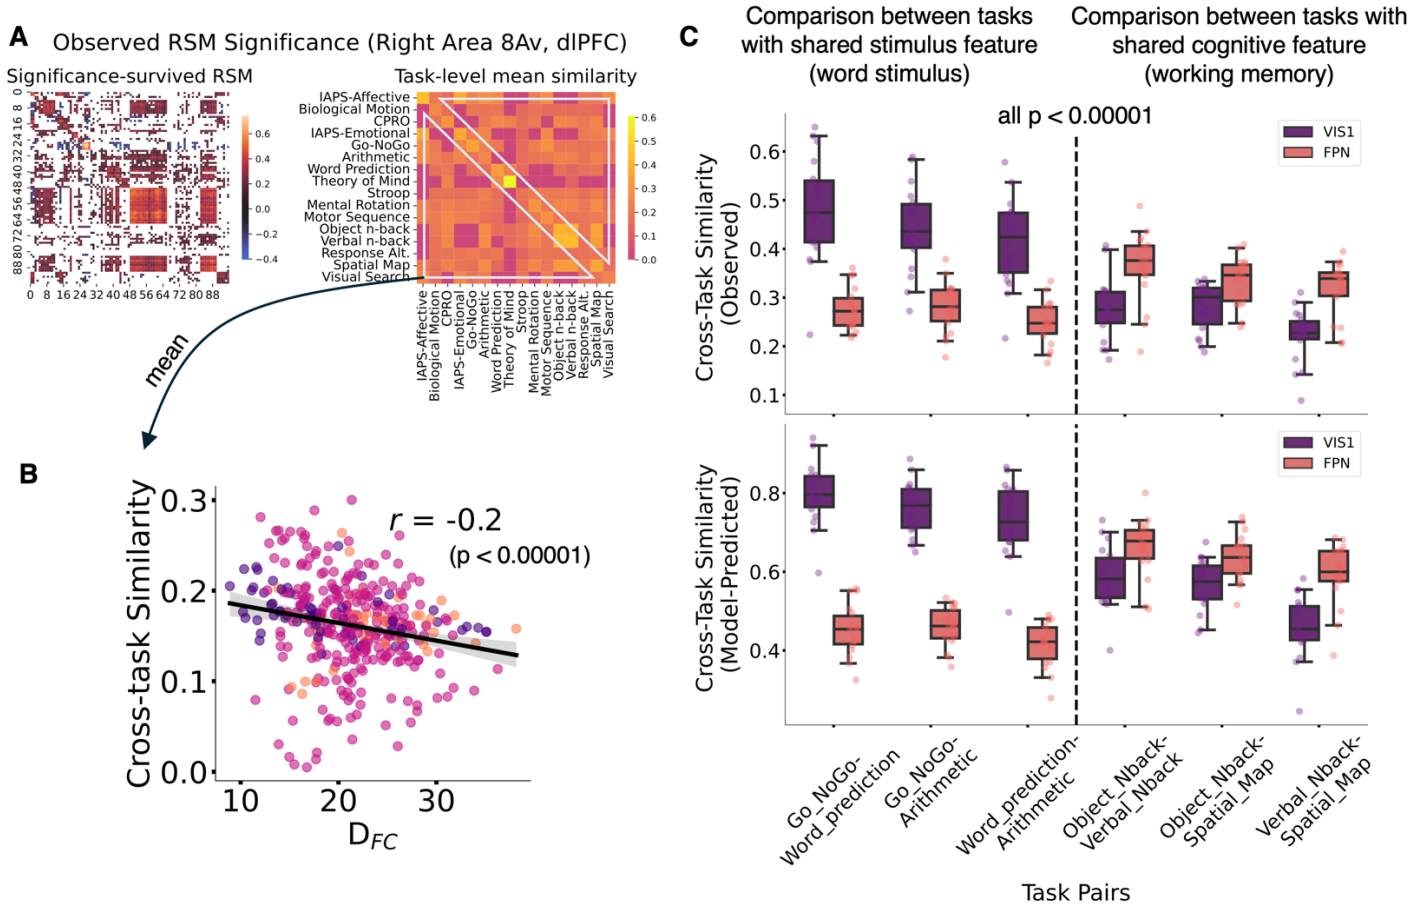

Figure 4: Shared cross-task cognitive and sensory features are generated by connectivity-based representational compressions. A) In an example region: (Left) RSM with surviving cells (with mean values across subjects) following one-sample t-test and FDR correction. (Right) Mean across positive values of the significance-survived RSM for each pair of tasks, which is a reduced form of RSM at the level of tasks instead of task conditions. The off-diagonal elements represent cross-task similarity values B) Mean cross-task similarity (averaged across all task pairs) at the region level correlates with connectivity dimensionality: higher cross-task similarity was observed when its connectivity with other regions is low-dimensional. C) Observed and predicted network-level cross-task similarities for example tasks in visual and frontoparietal networks: [Top] Cross-task similarities among tasks with shared sensory feature (usage of word stimuli) is higher in VIS1 compared to FPN regions. Opposite is true for tasks with shared cognitive feature (employment of working memory). [Bottom] Similar patterns are observed when connectivity-based predictions were analyzed. P-values correspond to comparison between VIS1 and FPN cross-task similarity values. (End of Figure 4 caption)

Our analysis yielded two key findings that together demonstrate how connectivity dimensionality shapes cross-task representations: First, according to our activity flow mapping framework, we hypothesized that propagation over low-dimensional connectivity would increase the similarity

between representations of different tasks. Empirical results supported our hypothesis, as we found that regions with lower-dimensional connectivity showed stronger average cross-task representational similarities across all task pairs ($r = -0.2$, $t(17) = -16.85$, $p = 4.81e-12$, 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 $p_{permutation} < 0.05$ in 355 out of 360 regions, Fig. S4A).

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) 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(23). 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).

Interestingly, 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. Cross-task similarities at the network level significantly differed from null models based on shuffled connectivity, for all networks (Fig. S4B). 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 ∩ 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 (Fig. 5A).

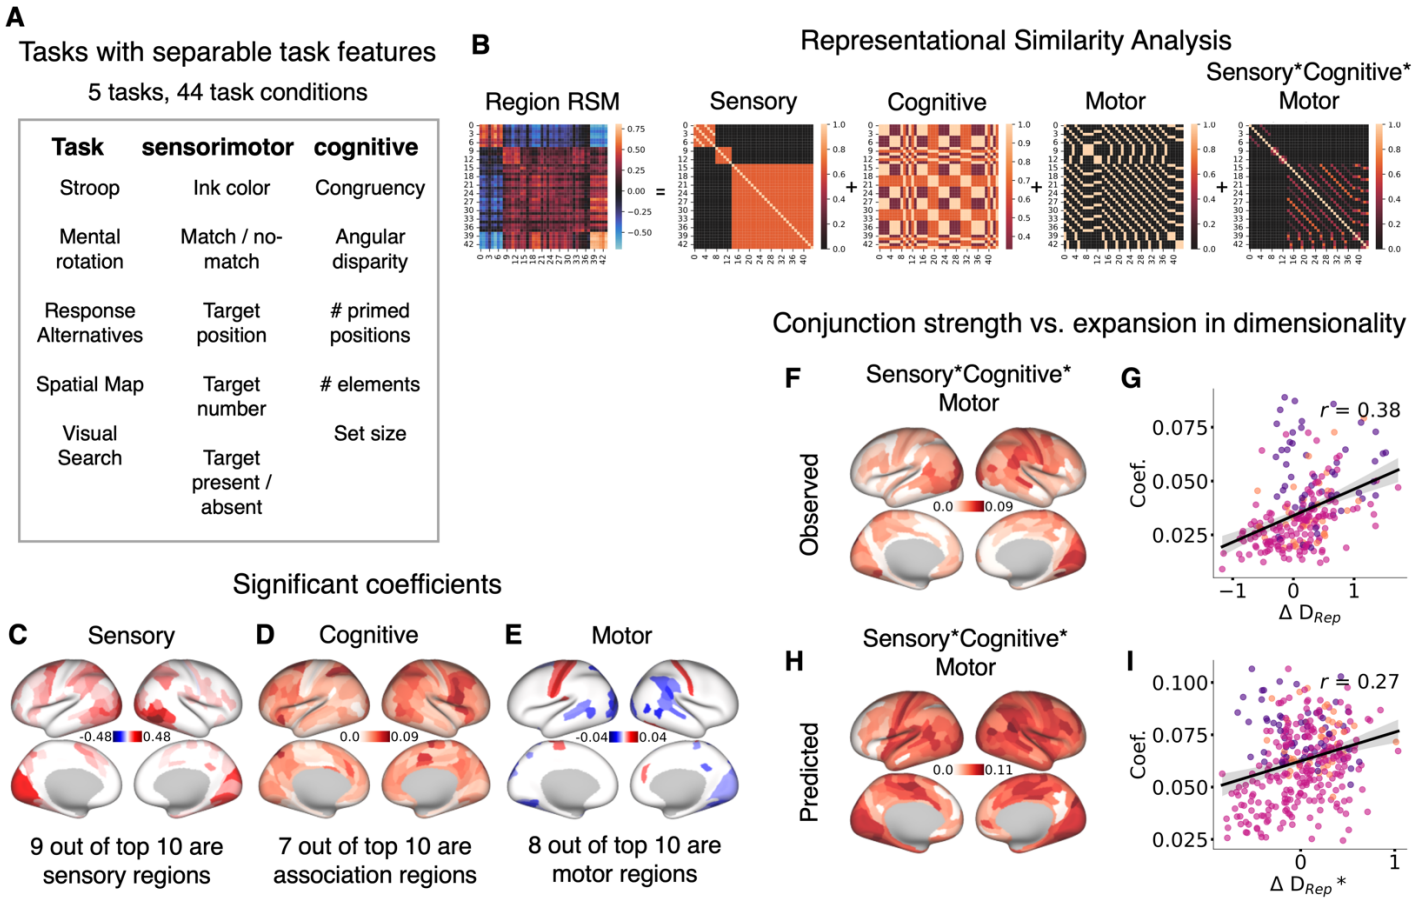

Figure 5: Conjunctive encoding of shared task features supported by connectivity-based expansion in dimensionality. A) List of a subset of five tasks with shared, separable sensorimotor and cognitive features across the task conditions. B) Task representational similarity analysis (RSA): regressing each observed RSM against theoretical RSMs based on shared task features: sensory, cognitive and motor features and their multiplicative conjunction. Example Region RSM is of the region Right Area PH, part of MT+Complex visual areas, a higher-order visual processing region. C-E) Brain plots showing significant coefficients of sensory, cognitive and motor terms (one-sample t-test against zero for significance, followed by FDR correction). F) Brain plot showing significant coefficients of the conjunction term. G) Conjunction term coefficients show significant relationship against representational compression/expansion. H) Brain plot of conjunction term coefficient by performing RSA on model-predicted RSMs. I) Predicted conjunction term coefficients also show significant relationship with predicted representational compression/expansion values.

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 probe potential relationships between these variables, we constructed theoretical similarity matrices (template RSMs) based on sensory, cognitive, and motor variables, as well as their multiplicative conjunctions. Using ridge regression, we fit these four vectorized matrices to the vectorized observed neural RSM (Fig. 5B).

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 eight out of 10 highest coefficients for the motor term were in motor regions (Fig. 5C-E). Consistent with our hypothesis, the relationship between conjunctive term coefficients and representational expansion was positive (Pearson r between values averaged across subjects: r_{conj} = 0.38, t-test against zero: t(17) = 6.81, p = 3.04e-06, Fig. 5F-G). This relationship was stronger than similar comparisons involving sensory, cognitive, and motor terms' coefficients against representational expansion: against sensory: $r_{sensory}$ = 0.15, t(17) = 2.63, p = 0.017, r_{conj} vs. $r_{sensory}$ (paired t-test following Fisher transformation): t(17) = 2.63, p = 0.018; against cognitive: $r_{cognitive}$ = 0.28, t(17) = 2.01, p = 0.061, r_{conj} vs. $r_{cognitive}$: t(17) = 2.43, p = 0.026; against motor: r_{motor} = 0.04, t(17) = 1.41, p = 0.178), r_{conj} vs. r_{motor} : t(17) = 2.69, p = 0.016.

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 = 0.27, t(17) = 3.11, p = 0.006, Fig. 5H-I). Correlations of predicted representational expansion with the other three terms' coefficients did not reach significance (against sensory: t(17) = -0.48, p = 0.637; against cognitive: t(17) = 0.48, p = 0.638; against motor: r = t(17) = 1.43, p = 0.171).

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).

Consistent with our hypothesis, region-averaged representational dimensionality showed a significant positive correlation with HCP cognitive flexibility composite scores (r=0.29, p=1.505e-08, Fig. S1A). Region-specific correlations between dimensionality and cognitive flexibility, after correction for multiple comparisons, revealed several regions with significant positive correlations (Fig. S1B). 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, Fig. S1C). 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, Fig. S1D).

These findings complement our observations about conjunctive coding by suggesting two distinct ways that high-dimensional representations support cognitive flexibility. 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. While limited by our current task set, 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 circuit-based mechanism by which the human brain likely 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 cognitive flexibility across diverse tasks. Together, these findings reveal how the brain's intrinsic connectivity architecture enables cognitive flexibility through systematic transformation of task-related information.

The present study expands previous connectivity-based models of task processing from structured rule combinations to diverse cognitive demands. Earlier work using the 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 (16). 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 validates connectivity-based information transfer as a general mechanism 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 (15), 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 explains how these representations arise.

Our study realizes a key theoretical framework for understanding cognitive flexibility: examining how multiple tasks mutually constrain neural activity patterns to reveal mechanisms of generalization (24). 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 (25). This connectivity-based framework remains open to discovering novel latent cross-task factors that emerge from how the brain naturally routes information, even if these factors don't map cleanly onto 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.

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 (7, 10, 12, 13). 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. Complementarily, high-dimensional connectivity enables the integration of multiple task variables into conjunctive representations, which we measured using established RSA methods (26) 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 of abstraction through compression and specificity through expansion provides a fundamental basis for cognitive flexibility, allowing the brain to both generalize across tasks and implement task-specific processing requirements.

Our analysis of the HCP dataset revealed that higher representational dimensionality, particularly in sensory and association regions, predicts enhanced cognitive flexibility 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(27). Finding that higher representational dimensionality predicts general cognitive flexibility 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 cognitive flexibility. Intriguingly, higher dimensionality in sensory and association regions, rather than motor regions, was most predictive of cognitive flexibility. 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 all, our study reveals a fundamental mechanism 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 novel framework for understanding both healthy cognitive function and potential dysfunction. Future studies may investigate how these connectivity-based representations adapt during learning and development, and whether alterations in this mechanism contribute to cognitive deficits in neurological and psychiatric disorders. Understanding how intrinsic connectivity shapes cognitive representations opens new avenues for interventions that leverage the brain's network architecture to enhance cognitive function or mitigate dysfunction.

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 (20).

We used the multi-domain task battery (MDTB) dataset (20), 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 (20) [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 (6), which used the same preprocessing pipeline.

The preprocessing of the fMRI data in this study utilized the Human Connectome Project (HCP) preprocessing pipeline(28), which was executed through the Quantitative Neuroimaging Environment & Toolbox [QuNex, version 0.61.17] (29). 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 (30). 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 (31). 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 (32). 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 (33) 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 (34). The estimation was performed using the 'GGLasso' package (35). 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 region (referred to as the target region, one at a time) and vertices of its connected neighbors (referred to as source regions). 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 as a regularization method to address the noisiness in estimating dense connectivity, reducing the complexity of the regression problem given the number of samples (timepoints) and thus reducing overfitting to noise. Note that principal component regression has been used successfully for FC estimation in multiple prior studies (15, 16, 36). 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 (37). 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.

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 point towards constrained connectivity with possibly convergent pathways and relatively sparse connections, whereas higher dimensionality might indicate more distributed or random connectivity patterns.

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 functional network definition proposed by (23). The Glasser parcellation was selected for its improved delineation of somatotopic and visuotopic areal organization, which are not accounted for in atlases defined solely on resting-state fMRI. Specifically, this parcellation incorporates prior knowledge (e.g., previously published retinotopic maps) to guide the division of areal/parcel boundaries, offering

advantages over purely data-driven approaches (33). The functional network definition 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), the somatomotor network regions as motor, 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 (38), 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. It offered a framework for investigating the transformation of neural representations across the cortical hierarchy, revealing insights into how information is processed and integrated across different functional domains.

Connectivity pattern similarity

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 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 $t \times t$, 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.

Activity flow modeling

We implemented activity flow modeling, and approach to model network flow-based generation of task-evoked activity adapted from previous studies (15, 16). 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 the connectivity-specified flow ($a_s c_{ts}$) of task-evoked activity from vertices of other 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 modeling approach contributes to our understanding of the relationship between FC and task-evoked activity patterns by providing a mechanistic link between the two. 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. Conversely, low correlations might indicate that other factors beyond the measured FC (and/or source task-evoked activity) play a significant role in shaping target task-evoked representations (e.g., local computations within each target region).

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 multitask 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 (20). 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 computed representational dimensionality from HCP task data(21, 22) 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 cognitive flexibility composite score following previous work(39). 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(40, 41)) with z-scored correct responses from the Penn Matrix Reasoning Test (a simplified version of Raven's Progressive Matrices). This composite score is a measure of general fluid intelligence, which indexes cognitive ability across a wide variety of cognitive tasks.

References

1. E. K. Miller, J. D. Cohen, An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience* **24**, 167–202 (2001).

2. T. Nakai, S. Nishimoto, Quantitative models reveal the organization of diverse cognitive functions in the brain. *Nat Commun* **11**, 1142 (2020).

3. T. Nakai, S. Nishimoto, Representations and decodability of diverse cognitive functions are preserved across the human cortex, cerebellum, and subcortex. *Commun Biol* **5**, 1–13 (2022).

4. T. Flesch, K. Juechems, T. Dumbalska, A. Saxe, C. Summerfield, Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* **110**, 1258-1270.e11 (2022).

5. J. Walters, *et al.*, Predicting brain activation maps for arbitrary tasks with cognitive encoding models. *NeuroImage* **263**, 119610 (2022).

6. T. Ito, J. D. Murray, Multitask representations in the human cortex transform along a sensory-to-motor hierarchy. *Nature neuroscience* **26**, 306–315 (2023).

7. V. Mante, D. Sussillo, K. V. Shenoy, W. T. Newsome, Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).

8. Y. Takagi, L. Hunt, M. W. Woolrich, T. E. Behrens, M. C. Klein-Flügge, Projections of non-invasive human recordings into state space show unfolding of spontaneous and over-trained choice. [Preprint] (2020). Available at: https://www.biorxiv.org/content/10.1101/2020.02.24.962290v1 [Accessed 8 November 2024].

9. S. L. Brincat, M. Siegel, C. von Nicolai, E. K. Miller, Gradual progression from sensory to task-related processing in cerebral cortex. *Proceedings of the National Academy of Sciences* **115**, E7202–E7211 (2018).

10. S. Bernardi, *et al.*, The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell* **183**, 954-967.e21 (2020).

11. T. Ito, *et al.*, Compositional generalization through abstract representations in human and artificial neural networks. [Preprint] (2022). Available at: http://arxiv.org/abs/2209.07431 [Accessed 5 November 2024].

12. M. Rigotti, *et al.*, The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).

13. S. Fusi, E. K. Miller, M. Rigotti, Why neurons mix: high dimensionality for higher cognition. *Current Opinion in Neurobiology* **37**, 66–74 (2016).

14. D. Badre, A. Bhandari, H. Keglovits, A. Kikumoto, The dimensionality of neural representations for control. *Current Opinion in Behavioral Sciences* **38**, 20–28 (2021).

15. M. W. Cole, T. Ito, D. S. Bassett, D. H. Schultz, Activity flow over resting-state networks shapes cognitive task activations. *Nat Neurosci* **19**, 1718–1726 (2016).

16. T. Ito, *et al.*, Cognitive task information is transferred between brain regions via resting-state network topology. *Nat Commun* **8**, 1027 (2017).

17. C. V. Cocuzza, *et al.*, Distributed network flows generate localized category selectivity in human visual cortex. *PLOS Computational Biology* **20**, e1012507 (2024).

18. T. Ito, L. Hearne, R. Mill, C. Cocuzza, M. W. Cole, Discovering the Computational Relevance of Brain Network Organization. *Trends in Cognitive Sciences* **24**, 25–38 (2020).

19. M. W. Cole, "Network modeling: The explanatory power of activity flow models of brain function" in *Computational and Network Modeling of Neuroimaging Data*, (Elsevier, 2024), pp. 81–117.

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

21. D. M. Barch, *et al.*, Function in the human connectome: Task-fMRI and individual differences in behavior. *NeuroImage* **80**, 169–189 (2013).

22. D. C. Van Essen, *et al.*, The WU-Minn Human Connectome Project: An overview. *NeuroImage* **80**, 62–79 (2013).

23. J. L. Ji, *et al.*, Mapping the human brain's cortical-subcortical functional network organization. *NeuroImage* **185**, 35–57 (2019).

24. M. Nau, A. C. Schmid, S. M. Kaplan, C. I. Baker, D. J. Kravitz, Centering cognitive neuroscience on task demands and generalization. *Nat Neurosci* **27**, 1656–1667 (2024).

25. G. Buzsáki, *The brain from inside out* (Oxford University Press, 2019).

26. A. Kikumoto, U. Mayr, Conjunctive representations that integrate stimuli, responses, and rules are critical for action selection. *Proceedings of the National Academy of Sciences* **117**, 10603–10608 (2020).

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

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

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

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

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

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

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

34. K. L. Peterson, R. Sanchez-Romero, R. D. Mill, M. W. Cole, Regularized partial correlation provides reliable functional connectivity estimates while correcting for widespread confounding. [Preprint] (2023). Available at: https://www.biorxiv.org/content/10.1101/2023.09.16.558065v1 [Accessed 5 August 2024].

35. F. Schaipp, C. L. Müller, O. Vlasovets, GGLasso -- a Python package for General Graphical Lasso computation. [Preprint] (2021). Available at: http://arxiv.org/abs/2110.10521 [Accessed 5 August 2024].

36. R. D. Mill, *et al.*, Network modeling of dynamic brain interactions predicts emergence of neural information that supports human cognitive behavior. *PLOS Biology* **20**, e3001686 (2022).

37. P. Gao, *et al.*, A theory of multineuronal dimensionality, dynamics and measurement. [Preprint] (2017). Available at: https://www.biorxiv.org/content/10.1101/214262v2 [Accessed 5 August 2024].

38. D. S. Margulies, *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).

39. E. Podvalny, R. Sanchez-Romero, M. W. Cole, Functionality of arousal-regulating brain circuitry at rest predicts human cognitive abilities. *Cerebral Cortex* **34**, bhae192 (2024).

40. E. P. Scott, A. Sorrell, A. Benitez, 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).

41. R. K. Heaton, *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

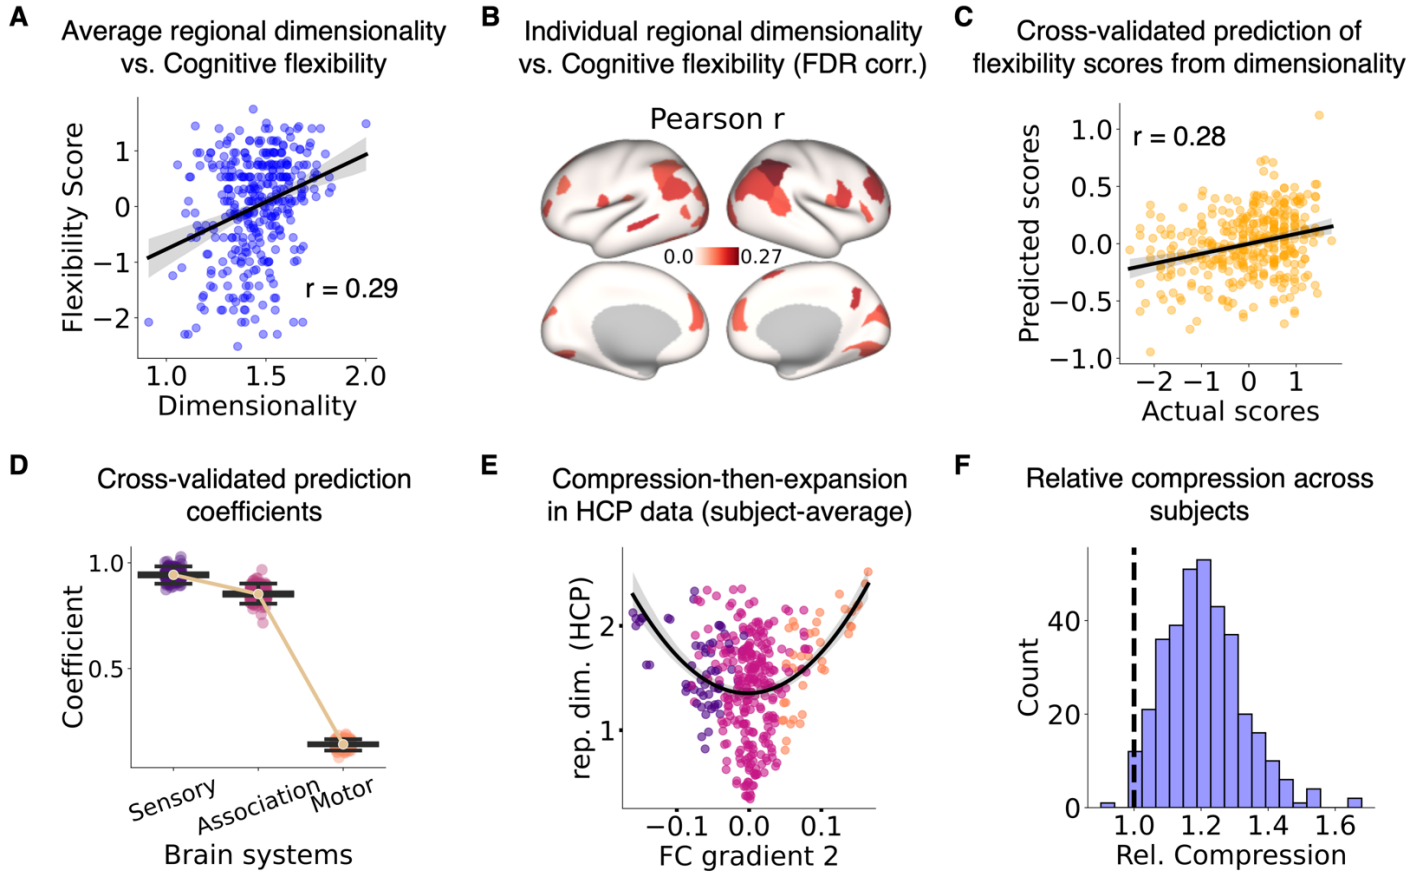

Supplementary Figure 1: Representational dimensionality is predictive of behavioral measures of cognitive flexibility. A) Average representational dimensionality across regions was positively correlated with cognitive flexibility scores across subjects. B) Similar relationship in individual regions (following 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.

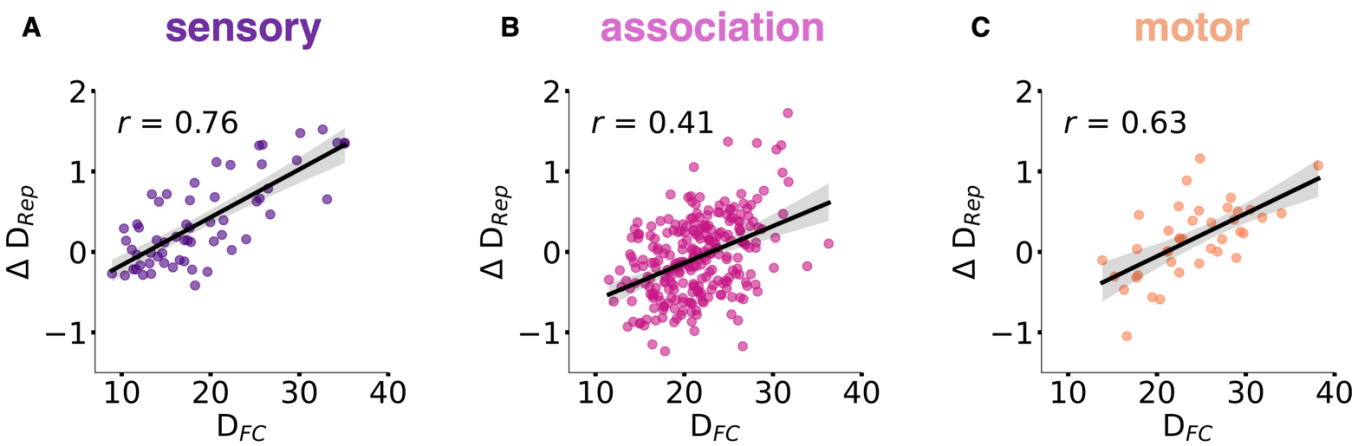

Supplementary Figure 2: Representational compression/expansion was significantly correlated with connectivity dimensionality across the regions of all three brain systems.

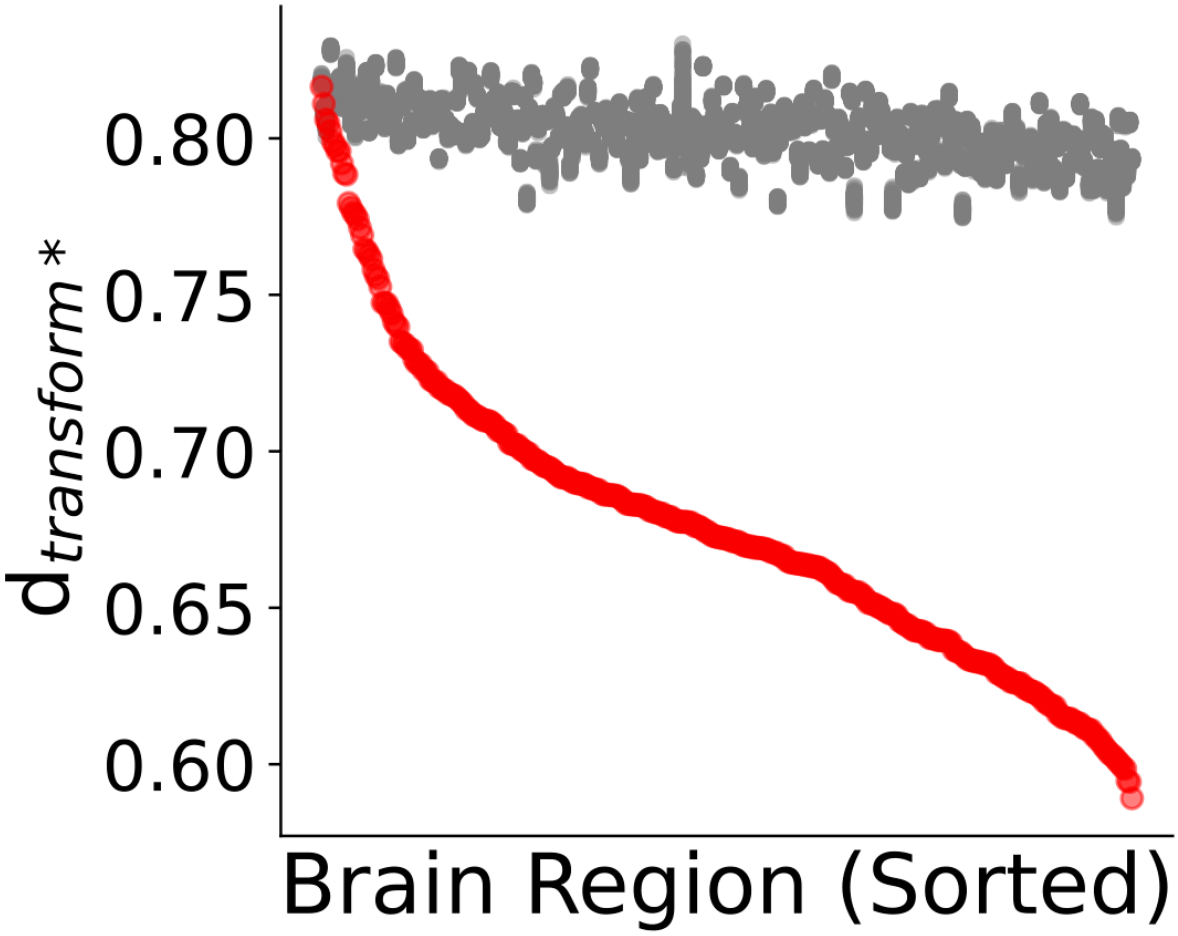

Supplementary Figure 3: Model-predicted transformation distance against null values determined from FC shuffling. Showing that structure present in the connectivity patterns yields significantly lower transformation distance compared to random transformation afforded by shuffled FC.

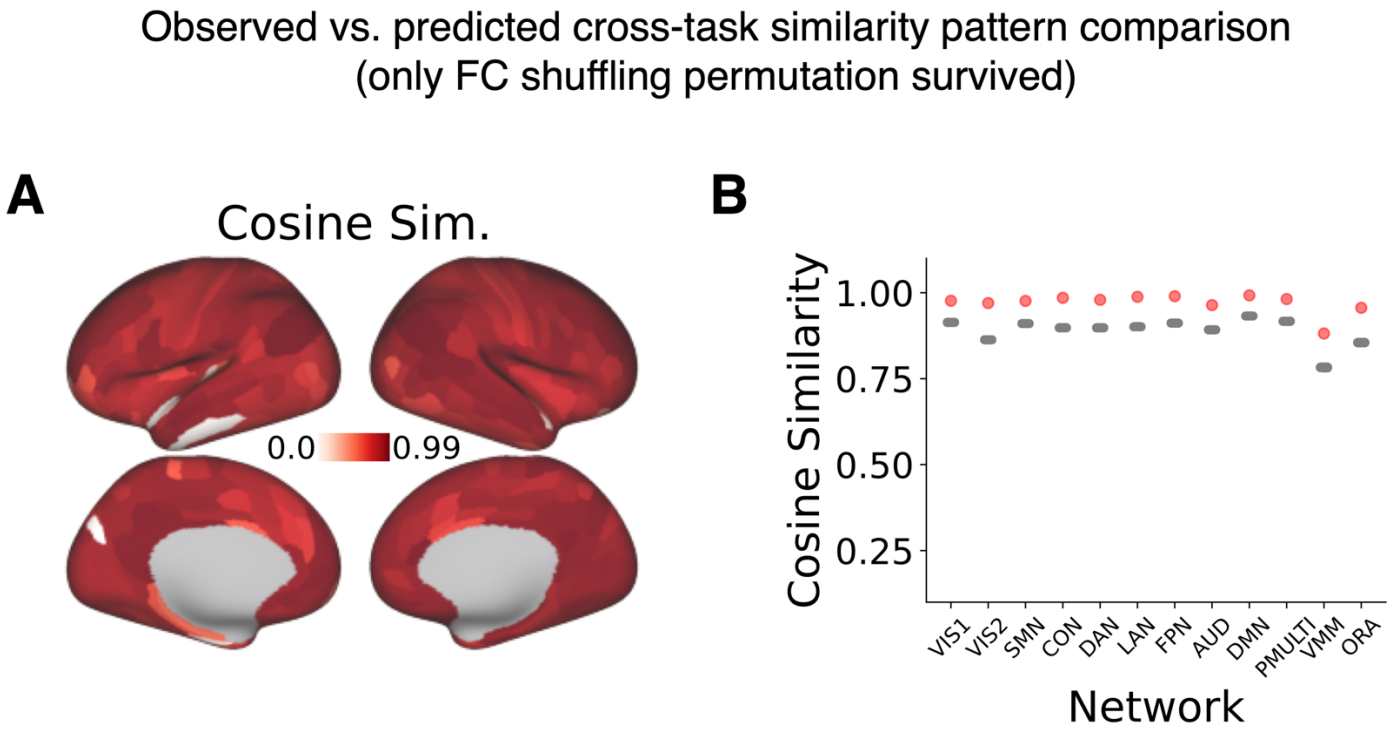

Supplementary Figure 4: Cross-task similarity predictions at the level of individual regions and large-scale brain networks. (A) Cosine similarity between vectorized cross-task similarity values from observed (empirically-estimated) and from connectivity-predicted estimates. Regions with no values did 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 (23): 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.