UC Berkeley
UC Berkeley Previously Published Works

Title

Orthogonal neural encoding of targets and distractors supports multivariate cognitive control.

Permalink

https://escholarship.org/uc/item/1h78w5n9

Journal

Nature Human Behaviour, 8(5)

Authors

Ritz, Harrison

Shenhav, Amitai

Publication Date

2024-05-01

DOI

10.1038/s41562-024-01826-7

Peer reviewed

Orthogonal neural encoding of targets and distractors supports multivariate cognitive control

Harrison Ritz[*,1,2,3], **Amitai Shenhav**[1,2]

[1]Cognitive, Linguistic & Psychological Science, Brown University, Providence, RI, USA

[2]Carney Institute for Brain Science, Brown University, Providence, RI, USA

[3]Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

Abstract

The complex challenges of our mental life require us to coordinate multiple forms of neural information processing. Recent behavioral studies have found that people can coordinate multiple forms of attention, but the underlying neural control process remains obscure. We hypothesized that the brain implements multivariate control by independently monitoring feature-specific difficulty and independently prioritizing feature-specific processing. During fMRI, participants performed a parametric conflict task that separately tags target and distractor processing. Consistent with feature-specific monitoring, univariate analyses revealed spatially segregated encoding of target and distractor difficulty in dorsal anterior cingulate cortex. Consistent with feature-specific attentional priority, a novel multivariate analysis (Encoding Geometry Analysis) revealed overlapping, but orthogonal, representations of target and distractor coherence in intraparietal sulcus. Coherence representations were mediated by control demands and aligned with both performance and frontoparietal activity, consistent with top-down attention. Together, these findings provide evidence for the neural geometry necessary to coordinate multivariate cognitive control.

Keywords

Introduction

We have remarkable flexibility in how we think and act. This flexibility is enabled by the array of mental tools we can bring to bear on challenges to our goal pursuit [1–6]. For example, someone may respond to a mistake by becoming more cautious, enhancing task-relevant processing, or suppressing task-irrelevant processing [7], and previous work has shown that people simultaneously deploy multiple such strategies at the same time in response to different task demands [3,8–10]. Flexibly coordinating multiple cognitive

[*]*Corresponding author:* hritz@princeton.edu.

processes requires a control system that can monitor multiple forms of task demands and deploy multiple forms of control (also referred to as the necessity for *observability* and *controllability*; [11]). These monitoring and regulation processes are fundamental to control, and are thought to be underpinned by distinct cingulo-opercular and frontoparietal neural systems [12–19]. However, much is still unknown about how multiple forms of control are represented across these domains.

Past research on the neural mechanisms of cognitive control has often sought to identify representations that integrate over multiple different sources of task demands (i.e., represent these different sources in *alignment*). For instance, previous studies has proposed that dorsal anterior cingulate cortex (dACC) tracks integrative features like response conflict, effort, value, error likelihood, and time-on-task [20–27]. Because they integrate over different task features instead of differentiating between them, these forms of 'aligned encoding' (Figure 1a) are ill-suited for carrying out multidimensional control. Multidimensional cognitive control requires independent representations that can track multiple sources of difficulty and regulate multiple cognitive processes (e.g., prioritize multiple sources of information [28]).

An alternative to aligned encoding – one that would allow the brain to separately control multiple processes – is *independent* encoding, which can come in at least two forms. One way the brain can have independent representations is by encoding different task features in spatially segregated neural populations ('segregated encoding'; Figure 1b). For example, past work has shown that different subregions within dACC encode distinct task demands, including various forms of errors and processing conflict [29–34]. The brain can instead have independent representations that are distributed across units within the same population, as has also been observed in dACC [35–37]. Within a shared population, independent encoding of information occurs along a set of orthogonal dimensions or *subspaces* (Figure 1c, 'subspace encoding'; [38–41]). Despite this exciting recent work, it remains unclear to what extent different components of the cognitive control system leverage these aligned, segregated, or orthogonal encoding strategies for monitoring multiple task demands and prioritizing multiple sources of information.

To gain new insight into the representations supporting cognitive control, we drew upon two key innovations. First, we leveraged an experimental paradigm we developed to tag multiple control processes [10]. Building on prior work [3,30,41,42], this task incorporates elements of perceptual decision-making (discrimination of a target feature) and inhibitory control (overcoming a salient and prepotent distractor). We have previously shown that we can separately tag target and distractor processing from participants' performance on this task, and that target and distractor processing are independently controlled. For example, participants adjust target and distractor sensitivity in response to distinct task demands (e.g., previous conflict or incentives; [10]). In conjunction with this process-tagging approach, our second innovation was to develop a novel multivariate fMRI analysis for measuring relationships between feature encoding (i.e., *encoding geometry*). Extending recent statistical approaches in systems neuroscience [35,43,44], we combined the strengths of multivariate encoding analyses and representation similarity analyses into a method we call 'Encoding Geometry Analysis' (EGA). We used EGA to characterize whether putative

markers of monitoring and prioritization leverage independent representations for targets and distractors.

In brief, we found that key nodes within the cognitive control network use orthogonal representations of target and distractor information to support cognitive control. In the dorsal anterior cingulate cortex (dACC), encoding of target and distractor difficulty was spatially segregated and arranged along a rostrocaudal gradient. By contrast, in the intraparietal sulcus (IPS), encoding of target and distractor coherence was encoded along orthogonal neural subspaces. These regional distinctions are consistent with hypothesized roles in planning and implementing (multivariate) attentional policies [12,17]. Furthermore, we found that coherence encoding depended on control demands, and was aligned with both task performance and frontoparietal activity, consistent with these coherence representations playing a critical role in cognitive control (e.g., feature prioritization). Together, these results suggest that cognitive control uses representational formats that allow the brain to monitor and control multiple streams of information processing.

Results

Task overview

Twenty-nine human participants performed the Parametric Attentional Control Task (PACT [10]) during fMRI. On each trial, participants responded to an array of colored moving dots (colored random dot kinematogram; Figure 1d). In the critical condition (Attend-Color), participants respond with a left/right keypress based on which of two colors were in the majority. In alternating scanner runs, participants instead responded based on motion (Attend-Motion), which was designed to be less control-demanding due to the (Simon-like) congruence between motion direction and response hand [3,10]. Across trials, we independently and parametrically manipulated target and distractor information across five levels of target coherence (e.g., percentage of dots in the majority color, regardless of which color) and distractor congruence (e.g., percentage of dots moving either in the congruent or incongruent direction relative to the correct color response; Figure 1e). This task allowed us to 'tag' participants' sensitivity to each dimension by measuring behavioral and neural responses to independently manipulated target and distractor features. Unlike a similar task used to study post-error adjustments [3], our parametric manipulation of target and distractor coherence allows us to better measure feature-specific representations. Unlike similar tasks used to study contextual decision-making [30,41,45], this task pits more control-demanding responses (towards color) against more automatic responses (towards motion), allowing comparisons between Attend-Color and Attend-Motion tasks to isolate the contributions of cognitive control [46,47].

Behavior

Participants had overall good performance on the task, with a high level of accuracy (median Accuracy = 89%, IQR = [84% - 92%]), and a low rate of missed responses (median lapse rate = 2%, IQR = [0% - 5%]). We used mixed effects regressions to characterize how target coherence and distractor congruence influenced participants' accuracy and log-transformed correct reaction times. Replicating previous behavioral findings using this task, participants

were sensitive to both target and distractor information [10]. When target coherence was weaker, participants responded slower ($t_{(27.6)} = 16.1$, $p = 1.60 \times 10^{-15}$) and less accurately ($t_{(28)} = -8.90$, $p = 1.19 \times 10^{-9}$; Figure 1f). When distractors were more incongruent, participants also responded slower ($t_{(28.8)} = 5.09$, $p = 2.15 \times 10^{-5}$) and less accurately ($t_{(28)} = -4.66$, $p = 6.99 \times 10^{-5}$; Figure 1g). Also replicating prior findings with this task, interactions between targets and distractors were not significant for reaction time ($t_{(28.2)} = 0.143$, $p = .887$) and had a weak influence on accuracy ($t_{(28)} = 2.36$, $p = .0257$), with model omitting target-distractor interactions providing a better complexity-penalized fit (RT ΔAIC = 17.7, Accuracy ΔAIC = 1.38).

Segregated encoding of target and distractor difficulty in dACC

Past work has separately shown that the dACC tracks task demands related to perceptual discrimination (induced in our task when target information is weaker) and related to the need to suppress a salient distractor (induced in our task when distractor information is more strongly incongruent with the target [12,30–32,48]). Our task allowed us to test whether these two sources of increasing control demand are tracked within common regions of dACC (reflecting an aggregated representation of multiple sources of task demands), or whether they are tracked by separate regions (potentially reflecting a specialized representation according to the nature of the demands).

Targeting a large region of dACC – a conjunction of a cortical parcellation with a meta-analytic mask for 'cognitive control' (see 'fMRI univariate analyses' in Methods) – we found spatially distinct signatures of target difficulty and distractor congruence within dACC. In caudal dACC, we found significant clusters encoding the parametric effect of target difficulty (Figure 2a; negative effect of target coherence in green), and in more rostral dACC we found clusters encoding parametric distractor incongruence (negative effect of distractor congruence in blue). Supporting this dissociation, the spatial patterns of target and distractor regression weights were uncorrelated across dACC voxels ($t_{(28.0)} = 1.32$, $p = .197$, logBF = −0.363). These analyses control for omission errors, and additionally controlling for commission errors produced the same whole-brain pattern at a reduced threshold (see Supplementary Figure 1). While distractor congruence was marginally significant when correcting for multiple comparisons across the entire brain (one-sided $p = .08$, whole-brain TFCE), extensive previous research predicts congruence effects in this ROI and in this direction [12,20,48], suggesting that a whole-brain corrected estimate is overly conservative.

To further quantify how feature encoding changed along the longitudinal axis of dACC, we used principal component analysis to extract the axis position of dACC voxels (see 'dACC longitudinal axis analyses' in Methods), and then regressed target and distractor beta weights onto these axis scores. We found that targets had stronger difficulty coding in more caudal voxels ($t_{(27.9)} = 3.74$, $p = .000840$), with a quadratic trend ($t_{(26.5)} = 4.48$, $p = .000129$; Figure 2b). In line with previous work on both perceptual and value-based decision-making [30,49–52], we found that signatures of target discrimination difficulty (negative correlation with target coherence) in caudal dACC were paralleled by signals of target discrimination *ease* (positive correlation with target coherence) within the rostral-most extent of our dACC ROI (Supplementary Figure 2). In contrast to targets, distractors had stronger incongruence

coding in more rostral voxel ($t_{(28.0)} = -3.26$, $p = .00294$), without a significant quadratic trend. We used participants' random effects terms to estimate the gradient location where target and distractor coding were at their most negative, finding that the target minimum was significantly more caudal than the distractor minimum (signed-rank test, $z_{(28)} = 2.41$, $p = .0159$). Target and distractor minima were uncorrelated across subjects ($r_{(27)} = .0282$, $p = .880$, logBF = −0.839), again consistent with independent encoding of targets and distractors.

As additional evidence that target-related and distractor-related demands have a dissociable encoding profile, we found that the crossover between target and distractor encoding in dACC occurred at the boundary between two well-characterized functional networks [53–55]. Whereas distractor-related demands were more strongly encoded rostrally in the Control Network (particularly within regions of dACC and insula corresponding to the 'Control C' Sub-Network; [54,56]), target-related demands were more strongly encoded caudally within the 'Salience / Ventral Attention (SVA)' Network (Figure 2C–D). Including network membership alongside long axis location predicted target and distractor encoding better than models with either network membership or axis location alone (BIC > 1675).

Subspace encoding of target and distractor coherence in intraparietal sulcus

We found that dACC appeared to dissociably encode target and distractor difficulty through spatially segregated encoding, consistent with a role in monitoring different task demands and/or specifying different control signals [12]. To identify neural mechanisms for the implementation of this control through the prioritization targets versus distractors, we next tested for regions that encode target and distractor coherence (the amount of information in a feature, regardless of which response it supports). Based on previous research, we might expect to find this form of selective attention in posterior parietal cortex [17,57,58]. We explored whether target and distractor coherence share a common neural code (e.g., as a global index of spatial salience), compared to where these features are encoded distinctly (e.g., as separate targets of control).

An initial whole-brain univariate analysis showed that overlapping regions throughout occipital, parietal, and prefrontal cortices track the feature coherence (proportion of dots in the majority category) for both targets and distractors (Figure 3a; conjunction in orange). These regions showed elevated responses to lower target coherence and higher distractor coherence, potentially reflecting the relevance of each feature for task performance. Note that in contrast to distractor congruence, distractor *coherence* had an inconsistent relationship with task performance (RT: $t_{(27.0)} = 2.08$, $p = .048$; Accuracy: $t_{(28)} = -0.845$, $p = .406$), suggesting that these neural responses are unlikely to reflect task difficulty per se.

While these univariate activations point towards widespread and coarsely overlapping encoding of the feature coherence (potentially consistent with aligned encoding; Figure 1a), they lack information about how these features are encoded at finer spatial scales. To interrogate the relationship between target and distractor encoding, we developed a multivariate analysis that combines multivariate encoding analyses with pattern similarity analyses, which we term Encoding Geometry Analysis (EGA). Whereas pattern similarity analyses typically quantify relationships between representations of specific stimuli or

responses (e.g., whether they could be classified, [59]), EGA characterizes relationships between encoding subspaces (patterns of contrast weights) across different task features, consistent with recent analyses trends in systems neuroscience [35,36,43,60–62]. A stronger correlation between encoding subspaces (either positive or negative) indicates that features are similarly encoded (i.e., that their representations are aligned and thus confusable by a linear decoder; Figure 1a), whereas weak correlation indicate that these representations are orthogonal (and thus distinguishable by a linear decoder; [59]). In contrast to standard pattern similarity, the sign of these relationships is interpretable in EGA, reflecting how features are coded relative to one another. Compared to standard encoding analyses, EGA is less sensitive to noise (Supplementary Figure 3). We estimated this encoding alignment within each parcel, correlating unsmoothed and spatially pre-whitened patterns of parametric regression betas across scanner runs to minimize spatiotemporal autocorrelation [63–65]. This cross-validated similarity further allowed us to anchor our analysis on the measurement reliability of encoding profiles (i.e., the self-correlation of encoding patterns across cross-validation folds [66,67]).

Focusing on regions that encoded both target and distractor information (parcels where both group-level $p < .001$), EGA revealed clear dissociations between regions that represent these features in alignment versus orthogonally. Within visual cortex and the superior parietal lobule (SPL), target and distractor representations demonstrated significant negative correlations (Figure 3b, red), reflecting (negatively) aligned encoding. In contrast, early visual cortex and intraparietal sulcus (IPS; see Figure 3c for anatomical boundaries) demonstrated target-distractor correlations near zero (Figure 3b, black), suggesting encoding along orthogonal subspaces.

To bolster our interpretation of the latter findings as reflecting orthogonal (i.e., uncorrelated) representations rather than merely small but non-significant correlations, we employed Bayesian t-tests at the group level to estimate the relative (log-10) likelihood that these encoding dimensions were orthogonal or correlated. Consistent with our previous analyses, we found strong evidence for correlation (positive log bayes factors) in more medial regions of occipital and posterior parietal cortex (e.g., SPL), and strong evidence for orthogonality (negative log bayes factors) in more lateral regions of occipital and posterior parietal cortex (e.g., IPS; Figure 3D). Control analyses confirmed that coherence orthogonality was not due to encoding reliability, as a similar topography was observed with disattenuated correlations (normalizing correlations by their reliability; see Supplementary Figure 4). Further supporting these results, our Bayes factor analyses were robust to the choice of priors (see Supplementary Figure 5).

While our analyses support independent encoding of targets and distractors within the same parcel, we further explored whether feature information is reflected in overlapping voxels (i.e., voxel-level mixed selectivity [40]). Simulations revealed that the alignment between absolute encoding weights can differentiate between pure and mixed selectivity, and parietal coherence representations bore this signature of voxel-level mixed selectivity (Supplementary Figure 6), consistent with the subspace encoding hypothesis.

These results have focused on the coherence of different features regardless of the response they support, demonstrating that SPL exhibits aligned representations of target and distractor coherence. Past decision-making research has separately demonstrated that SPL tracks the amount of evidence supporting specific response [42,68,69], which we found was also true for our task. In addition to encoding target and distractor coherence, SPL and visual cortex also tracked target and distractor 'evidence' (proportion of dots supporting a rightward vs leftward response; Figure 3e). EGA revealed orthogonal evidence representations between targets and distractors, in the same areas with aligned coherence representations (compare Figure 3d and 3e), consistent with previous observations of multiple decision-related signals in SPL [68].

We complemented our whole-brain analyses with ROI analyses in areas exhibiting reliable encoding of key variables, focusing on core frontal regions linked with cognitive control (dACC and lateral PFC [lPFC]), and parietal regions linked with decision-making and attention (SPL and IPS; [12,15]). Consistent with our analyses above, we found that target and distractor coherence encoding was aligned in SPL, but not in IPS (Figure 4a, compare to Figure 3d), whereas SPL encoded target and distractor evidence. Directly comparing these regions (see Table 1), we found stronger encoding of target evidence in SPL, stronger encoding of target coherence in IPS, and stronger alignment between target-distractor coherence alignment in SPL. Unlike our univariate results, we did not find distractor congruence encoding in dACC (though this was found in lPFC and IPS). Instead, dACC showed multivariate encoding of target coherence and evidence.

To further characterize how feature coherence and evidence are encoded across these regions, we performed multidimensional scaling over each regions task representations (Figure 4b; [64,70]). Briefly, this method allows us to visualize – in a non-parametric manner – the relationships between representations of different feature levels (e.g., levels of target coherence), by estimating each feature level separately within a GLM and then using singular value decomposition to project these patterns into a 2D space (see Methods for additional details). We found that coherence and evidence axes naturally emerge in the top two principal components in this analysis within dACC, SPL, and IPS. Coherence axes (light to dark shading) are parallel between left (blue) and right (brown) responses, suggesting a response-independent encoding. In these components, evidence encoding appeared to be binary, in contrast to parametric coherence encoding (we found similar whole-brain encoding maps for binary-coded evidence; see Supplementary Figure 7). Critically, whereas coherence encoding axes within SPL were aligned between targets (circles) and distractors (diamonds; confirming aligned encoding), in IPS these representations form perpendicular lines (confirming orthogonal encoding). When we visualized higher dimensions, we found that IPS did appear to have weak encoding alignment between target and distractor coherence in higher dimensions (Supplementary Figure 8). Nevertheless, the orthogonal encoding in the first two principal components is sufficient for a downstream region to have an independent read-out of feature-specific coherence. These analyses both help to visualize cross-region dissociations in encoding profiles and validate that task features are encoded in a monotonic fashion.

Finally, to explore the divisions between SVA and Control networks evident in the univariate analyses, we split up our two prefrontal ROIs by their network membership (Supplementary Figure 9). In dACC, we found that SVA parcels tended to have stronger feature encoding than Control parcels. Interestingly, in these SVA parcels several features were aligned with the target evidence dimension, consistent with recent human electrophysiology findings [35]. In lPFC, we found that Control parcels, but not SVA parcels, encoded distractor congruence (Control: $t_{(28)} = 3.60$, two-tailed $p = .0012$, logBF = 1.45; SVA: $t_{(28)} = 0.57$, $p = .57$, logBF = −0.64; Control − SVA: $t_{(28)} = 3.27$, $p = .0029$, logBF = 1.12). This distractor congruence encoding was present in lPFC both in 'Control A/B' parcels ($t_{(28)} = 3.66$, $p = .001$) and marginally in 'Control C' parcels ($t_{(28)} = 1.86$, $p = .073$). This network-selective encoding of congruence is consistent with the univariate results in dACC (see Figure 2).

Control demands dissociate coherence and evidence encoding

Our findings thus far demonstrate two sets of dissociations within and across brain regions. In dACC, we find that distinct regions encode the control demands related to discriminating targets (caudal dACC) versus overcoming distractor incongruence (rostral dACC). In posterior parietal cortex, we find that overlapping regions track the coherence of these two stimulus features, but that distinct regions represent these features in alignment (SPL) versus orthogonally (IPS). While these findings suggest that this set of regions was involved in translating between feature information and goal-directed responding, they only focus on the information that was presented to the participant on a given trial. To provide a more direct link between feature-specific encoding and control, we examined how the encoding of feature coherence differed between matched task that placed stronger or weaker demands on cognitive control. So far, our analyses have focused on conditions in which participants needed to respond to the color feature while ignoring the motion feature (Attend-Color task), but on alternating scanner runs participants instead responded to the motion dimension and ignored the color dimension (Attend-Motion task). These tasks were matched in their visual properties (identical stimuli) and motor outputs (left/right responses), but critically differed in their control demands. Attend-Motion was designed to be much easier than Attend-Color, as the left/right motion directions are compatible with the left/right response directions (i.e., Simon facilitation; [3,10]). Comparing these tasks allows us to disambiguate bottom-up attentional salience from the top-down contributions to attentional priority [47,71–73].

Consistent with previous work [10], performance on the Attend-Motion task was better overall (mean RT: 565ms vs 725ms, sign-rank $p = 2.56 \times 10^{-6}$; mean Accuracy: 93.7% vs 87.5%, sign-rank $p = .000318$). Unlike the Attend-Color task, performance was not impaired by distractor incongruence (i.e., color distractors; RT: $t_{(28)} = −1.39$, $p = .176$; Accuracy: $t_{(28)} = 0.674$, $p = .506$). To investigate these task-dependent feature representations, we fit a GLM that included both tasks. To control for performance differences across tasks, we only analyzed accurate trials and included trial-wise RT as a nuisance covariate, concatenating RT across tasks.

Whereas the encoding of both color and motion coherence was widespread during the Attend-Color task (Figure 3), coherence encoding was consistently weaker during the less demanding Attend-Motion task (Figure 5A). Coherence encoding was weaker during

Attend-Motion whether classifying according to goal-relevance (comparing targets or distractors) or the features themselves (comparing motion or color). Task-relevant ROIs revealed that coherence encoding was effectively absent during the easy Attend-Motion task (Figure 5B), suggesting that they depend on the control demands of the Attend-Color task [47,74].

In contrast to these stark task-related differences in coherence encoding, we found that neural encoding of the target evidence (color evidence in the Attend-Color task and motion evidence in the Attend-Motion task) was preserved across tasks, including within dACC, lPFC, SPL, and IPS (Figure 5B). Consistent with previous experiments examining context-dependent decision-making [36,41,42,45,73,75,76], we found stronger target evidence encoding relative to distractor evidence encoding, in our case in the evidence-encoding SPL (Attend-Color: $t_{(28)} = 4.26$, one-tailed $p = 0.0001$; Attend-Motion: $t_{(28)} = 2.37$, one-tailed $p = 0.0124$). We also found that target evidence encoding during Attend-Motion was aligned with Attend-Color, both for *motion* evidence encoding ('stimulus axis'; SPL: one-tailed $p = .0236$, IPS: one-tailed $p = .0166$) and *target* evidence encoding ('decision axis'; SPL: one-tailed $p = 1.29 \times 10^{-6}$; IPS: one-tailed $p = .0005$), again in agreement with these previous experiments. Whereas our experiment replicates previous observations of the neural representations supporting contextual decision-making, we now extended these findings to understand how putative attention signals (i.e., feature coherence) are encoded in response to the asymmetric inference that is characteristic of cognitive control [77].

Aligned encoding dimensions for feature coherence and task performance

Feature coherence encoding (i.e., feature strength, regardless of response or congruence) depends on task demands, consistent with a role in cognitive control. To further understand this relationship between coherence encoding and control, we next explored how coherence encoding was related to task performance. We tested this question by determining whether feature coherence representations were aligned with performance representations (i.e., alignment between stimulus and behavioral subspaces [78]). Specifically, we included trial-level reaction time and accuracy in our first-level GLMs. Encoding of performance was itself highly robust: most parcels encoded reaction time and accuracy, with the strongest encoding in cognitive control regions (Supplementary Figure 10). Across cortex, reaction time and accuracy were negatively correlated, again most prominently across the cognitive control network. To explore the behavioral relevance of coherence representations, we tested whether coherence encoding was aligned with the voxel patterns encoding task performance.

We found that the encoding of target and distractor coherence was aligned with performance across frontoparietal and visual regions (Figure 6a–b). If a regions' encoding of target coherence reflects how sensitive the participant was to target information on that trial (e.g., due to top-down priority), we would expect target encoding to be positively aligned with performance on a given trial, such that stronger target coherence encoding is associated with better performance and weaker target coherence encoding is associated with poorer performance. We would also expect distractor encoding to demonstrate the opposite pattern – stronger encoding associated with poorer performance and weaker encoding associated with better performance. We found evidence for both patterns of feature-performance

alignment across visual and frontoparietal cortex: target encoding was aligned with better performance (faster RTs and higher accuracy; Figure 6a), whereas distractor encoding was aligned with worse performance (slower RTs and lower accuracy; Figure 6b).

Next, we examined whether performance-coherence alignment reflected individual differences in participants' task performance in our main task-related ROIs (see Figures 3–4). In particular, we tested whether the alignment between features and behavior reflects specific relationships with speed or accuracy, or whether they reflected overall increases in evidence accumulation (e.g., faster responding and higher accuracy). Within each ROI, we correlated feature-RT alignment with feature-accuracy alignment across subjects. We found that in dACC and IPS, participants showed the negative correlation between performance alignment measures predicted by an increase in processing speed (Figure 6c). People with stronger alignment between target coherence and shorter RTs tended to have stronger alignment between target coherence and higher accuracy, with the opposite found for distractors. While these between-participant correlations were present within targets and distractors, we did not find any significant correlations across features (between-feature: all $ps > .10$), again consistent with feature-specific processing. These analyses were qualitatively similar after partialing out the reliability of coherence and performance encoding, albeit with dACC and lPFC now showing marginal correlations for target coherence (see Supplementary Table 1). While between-participant analyses using small sample sizes warrant a note of caution, these findings are consistent across features and regions. In conjunction with our within-participant evidence that feature coherence representations are aligned with performance efficiency, these findings support a role for coherence encoding in adaptive control.

Coherence encoding aligns with frontoparietal activity

Across frontal, parietal, and visual cortex, encoding of target and distractor coherence depended on task demands and was aligned with performance. Since this widespread encoding of task information likely reflects distributed network involvement in cognitive control [77,79,80], we sought to understand how frontal and parietal systems interact. We focused our analyses on IPS and lateral PFC (lPFC), linking the core parietal site of orthogonal coherence encoding (IPS) to an prefrontal site previous work suggests provides top-down feedback during cognitive control [58,79,81,82]. Previous work has found that IPS attentional biases lower-level stimulus encoding in visual cortices [83,84], and that IPS mediates directed connectivity between lPFC and visual cortex during perceptual decision-making [42]. Here, we extended these experiments to test how IPS mediates the relationship between prefrontal feedback and stimulus encoding.

To investigate these putative cortical interactions, we developed a novel multivariate connectivity analysis to test whether coherence encoding was aligned with prefrontal activity, and whether this lPFC-coherence alignment was mediated by IPS. We first estimated the voxel-averaged residual timeseries in lPFC (SPM12's eigenvariate), and then included this residual timeseries alongside task predictors in a whole-brain regression analysis (Supplementary Figure 11). This analysis can be schematized as:

$$\beta_{seed} = GLM(X, Y_{seed})$$

(1)

$$e_{seed} = PCA(Y_{seed} - X\beta_{seed})$$

(2)

$$\beta_{\text{all}} = GLM([X, e_{seed}], Y_{all})$$

(3)

The GLM function performs regression using design matrix X and multivariate voxel timeseries Y, and the PCA function extracts the first principal component of the residuals. Finally, we used EGA to test whether there was alignment between patterns encoding lPFC functional connectivity (i.e., betas from the residual timeseries predictor e_seed) and patterns encoding target and distractor coherence. Note that while these results reflect functional connectivity, all correlational measures are subject to potential confounding [85].

We found that lPFC connectivity patterns were aligned with coherence-encoding patterns in visual cortex (Figure 7A). Stronger prefrontal functional connectivity was aligned with weaker target coherence and stronger distractor coherence, consistent with prefrontal recruitment during difficult trials. Notably, IPS connectivity was also aligned with target and distractor coherence in overlapping parcels, even when controlling for lPFC connectivity. These effects were liberally thresholded for visualization, as significant direct and indirect effects are not necessary for significant mediation [86].

Our critical test was whether IPS mediated the relationship between lPFC activity and coherence encoding. We compared regression estimates between a model that only included lPFC residuals ('solo' model) to a model that included both lPFC and IPS residuals ('both' model). Comparing the strength of lPFC-coherence alignment with and without IPS is a test of whether parietal cortex mediates lPFC-coherence alignment (MacKinnon et al., 2007). These models can be schematized as:

$$\beta_{solo} = GLM([X, e_{lPFC}], Y_{all})$$

(4)

$$\beta_{both} = GLM([X, e_{lFPC}, e_{IPS}], Y_{all})$$

(5)

We found that this mediation was strongest in early visual cortex, where the alignment between lPFC and feature coherence was reduced in a model that included IPS relative to a model without IPS (Figure 7B). The negatively correlated target-lPFC relationship

became more positive when IPS was included (top), and the positively correlated distractor-PFC relationship became more negative when IPS was included (bottom). Critically, we found that IPS reduced prefrontal-coherence alignment in early visual cortex more than lPFC reduced parietal-coherence alignment (Figure 7B inset; Supplementary Figure 12A–B), consistent with frontal-to-parietal directed connectivity in previous research [42,81]. Looking within color- and motion-sensitive parcels, determined using task-free localizer runs (see Methods), we found this mediation was robustly significant in color-sensitive cortex and marginally significant in motion-sensitive cortex. The opposite relationship, lPFC mediation of IPS connectivity, appeared in higher-level visual cortex for distractor coherence (Supplementary Figure 12C–D), though these effects were not reliable in explicit contrasts and may reflect projections from both regions. Note that we did not see any significant mediation of first-order target or distractor coherence encoding by IPS.

While we were primarily interested in alignment with lPFC due to previous work implicating these regions in top-down control (for reviews, see [12,87]), for completeness we also examined how different subnetworks in both lPFC and dACC aligned with coherence encoding. In lPFC, we found that SVA and Control subnetworks had similar patterns of alignment (Supplementary Figure 13). In dACC we found that the SVA subnetwork had a qualitatively similar profile of coherence alignment as lPFC, but this alignment was absent in the Control subnetwork. Whereas this seed-coherence alignment was similar across lPFC and SVA dACC, unlike lPFC we found that SVA dACC failed to demonstrate strong evidence for mediation by IPS (Supplementary Figure 14).

A final set of analyses examined whether SPL and IPS demonstrated different patterns of task-related functional connectivity with other regions, given that we found that these regions differentially encoded evidence and coherence. When seeding our connectivity analyses with SPL activity, we found that SPL activity aligned with evidence encoding in bilateral motor cortex (Supplementary Figure 15). In contrast, IPS activity did not significantly align with evidence encoding, and this seed-evidence alignment in motor cortex was stronger for SPL than IPS, consistent with a putative role for SPL in response selection [68].

Discussion

In this experiment, we explored whether neural control systems use representations with the same dimensionality as the processes they regulate [2,5,11]. Inspired by behavioral evidence that participants can independently control their sensitivity to targets and distractors [10], we set out to understand whether the neural correlates of monitoring and prioritization leverage independent encoding for feature-selective control (Figures 1a–c). We found that key nodes of canonical cognitive control networks had orthogonal neural representations of targets and distractors. Within dACC, orthogonal representations of target and distractor difficulty arose from segregated encoding along a rostrocaudal axis. Within IPS, orthogonal representations of target and distractor coherence arose from orthogonal subspaces in overlapping voxels. Consistent with a role in attentional priority, coherence representations depended on control demands, task performance, and frontoparietal activity. Together, these results reveal a

neural mechanism for how cognitive control prioritizes multiple streams of information during decision-making.

Neurocomputational theories have proposed that dACC is involved in planning control across multiple levels of abstraction [12,88–90]. Past work has found that control abstraction is hierarchically organized along dACC's rostrocaudal axis, with more caudal dACC involved in lower-level action control, and more rostral dACC involved in higher-level strategy control [30–32,34], an organization that may reflect a more general hierarchy of abstraction within PFC [31,91–93]. Consistent with this account, we found that caudal dACC tracked the coherence of the target and distractor dimensions, especially within the SVA network. In contrast, more rostral dACC tracked incongruence between targets and distractors, especially within the Control network. Speculatively, our results are consistent with caudal dACC tracking the first-order difficulty arising from the relative salience of feature-specific information, and more rostral dACC tracking the second-order difficulty arising from cross-feature (in)compatibility [92], the latter of which may require additional disengagement from distractor-dependent attentional capture.

Whereas dACC encoded feature difficulty (e.g., distractor incongruence), in parietal cortex we found overlapping representations of feature coherence (e.g., distractor coherence). In SPL, features had correlated coherence encoding (similarly representing low target coherence and high distractor coherence), consistent with this region's transient and non-selective role in attentional control [94–99]. In contrast, IPS had orthogonal representations of feature coherence, consistent with selective prioritization of task-relevant information [47,71–73,81,83,94–96,99,100]. While IPS primarily encoded features orthogonally (i.e., in the largest components of our multidimensional scaling analysis), the total coherence across features could also be read out at higher dimensions. The ability of IPS to communicate both orthogonal and aligned coherence representations is consistent with the diverse roles of IPS in attentional control.

Our previous work has demonstrated behavioral evidence for independent control over target and distractor attentional priority in this task [10], with different task variables selectively enhancing target or distractor sensitivity (see also [4,101]). Orthogonal feature representation in IPS may offer a mechanism for this feature-selective control, consistent with theoretical accounts of IPS implementing a priority map that combines stimulus- or value-dependent salience with goal-dependent feedback from PFC [17,57,58,80,102].

In dACC, we found that target and distractor difficulty encoding was consistent with the segregated encoding hypothesis, with features evoking univariate responses in distinct but adjacent regions. Interestingly, we did not find corresponding encoding of distractor congruence in our multivariate analyses within dACC, potentially reflecting the spatial smoothness of this response. However, we did find multivariate encoding of distractor congruence in lPFC, and multivariate encoding of target and distractor coherence in IPS. These multivariate profiles were consistent with our subspace encoding hypothesis. The reasons for a mix of segregated and subspace encoding across cortex is unclear, but this may speculatively reflect the segregation across functional networks. Like in dACC, distractor congruence had stronger encoding within the lPFC Control network, albeit without the

feature segregation (lPFC Control parcels also encoded target coherence in an orthogonal subspace). It is possible that these network segregations help bind related control processes [15,18,80], a hypothesis that future experiments should test with targeted paradigms (e.g., with subject-specific functional networks).

By comparing two different task goals (Attend-Color vs. Attend-Motion), our study was able to test whether coherence representations reflect control-dependent prioritization of information processing. Previous research has shown that these tasks differ dramatically in their control demands [10]. As in previous work, task performance was much better in Attend-Motion runs than Attend-Color runs, and participants were not sensitive to color distractors. Consistent with previous work on context-dependent decision-making, target evidence had similarly strong encoding across tasks, with generalizable encoding dimensions for choice and motion directions [36,41,45]. In contrast to these putative decision representations, we found that coherence representations disappeared in the easier Attend-Motion task. On its own, weaker encoding of color distractors in Attend-Motion could be explained by the weaker bottom-up salience of the color dimension. However, the stark drop in the encoding of target (motion) coherence in these blocks cannot be similarly accounted for – these differences in target coherence encoding showed the opposite relationship expected from salience: better encoding of low-salience color targets (hard Attend-Color task) and weaker encoding of high-salience motion targets (easy Attend-Motion task). Instead, this encoding profile is consistent with previous research finding that feature decoding is stronger for more difficult tasks [47,71,72,103] or when people are incentivized to use cognitive control [104,105].

Critically, stimuli and responses were matched across tasks, helping to rule out alternative accounts of coherence encoding based on 'bottom-up' stimulus salience, decision-making, or eye movements. Difficulty-dependent coherence encoding may instead reflect the involvement of an attention control system that can separately regulate target and distractor processing, speculatively indexing the top-down 'gain' or 'priority' on these features [17,58,102]. Supporting this account, coherence representations in cognitive control regions like IPS were aligned with performance representations, with target encoding strength aligned with better performance and distractor encoding strength aligned with poorer performance. Individual difference in feature-performance alignment was correlated across features, consistent with these representations reflecting the underlying processes (e.g., priority) that give rise to behavior, rather than performance monitoring or surprise (which would likely have the opposite relationship, e.g., high target coherence aligned with poorer performance).

Classic models of prefrontal involvement in cognitive control [77,82,106] propose that prefrontal cortex biases information processing in sensory regions. In line with this macro-scale organization, we found that coherence encoding in visual cortex was related to functional connectivity with the frontoparietal network. In particular, coherence encoding in visual cortex was aligned with patterns of functional connectivity to lateral prefrontal cortex, and this feature-seed relationship was mediated by IPS. The results of this novel multivariate connectivity analysis are consistent with previous research supporting a role for IPS in top-down control of visual encoding [83,84,107], as well as a granger-causal PFC-IPS-visual pathway during a similar decision-making task [42]. Here, we demonstrate stable 'communication subspaces' between visual cortex and PFC [108,109], which can

plausibly communicate feedback adjustments to feature gain. With that said, while our interpretation of the direction of communication is therefore supported by prior work, these connectivity methods are correlational [85], and cannot rule out the possibility that our mediation findings reflect a bottom-up pattern of communication (e.g., visual-IPS-PFC). The asymmetric mediation between regions (i.e., IPS mediates lPFC more than lPFC mediates IPS; Supplementary Figure 12) rules out a range of potential confounders, and these regions were selected based on the anatomical connectivity within the frontoparietal network, notably through the superior longitudinal fasciculus [110]. Future research should use temporally precise neuroimaging to account for directionality, causal manipulations to account for causality (e.g., [111]), and should explore the higher dimensional connectivity subspaces that link different regions [103,109]. These considerations notwithstanding, our findings are consistent with IPS, a critical site for orthogonal feature representations, playing a key role in linking prefrontal cortex with early perceptual processing.

Collectively, our findings provide new insights into how the brain may control multiple streams of information processing. While evidence for multivariate control has a long history in attentional tracking [28,112], including parametric relationships between attentional load and IPS activity [113–117], little is known about how the brain coordinates multiple control signals [2,5]. Future experiments should further elaborate on this frontoparietal control circuit, for instance by interrogating how incentives influence different task representations [104,105,118–120], or how neural and behavioral indices of control causally depend on perturbations of neural activity [111]. Future experiments should also use fast timescale neural recording technologies like (i)EEG or (OP-)MEG to better understand the within-trial dynamics of multivariate control [10,121]. In sum, this experiment provides new insights into the large-scale neural networks involved in multivariate cognitive control, and points towards new avenues for developing a richer understanding of goal-directed attention.

Methods

Participants

Twenty-nine individuals (17 females, Age: M = 21.2, SD = 3.4) participated in this experiment. All participants had self-reported normal color vision and no history of neurological disorders. Two participants missed one Attend-Color block (see below) due to a scanner removal, and one participant missed a motion localizer due to a technical failure, but all participants were retained for analysis. Participants provided informed consent, in accordance with Brown University's institutional review board.

Task

The main task closely followed our previously reported behavioral experiment [10]. On each trial, participants saw a random dot kinematogram (RDK) against a black background. This RDK consisted of colored dots that moved left or right, and participants responded to the stimulus with button presses using their left or right thumbs.

In Attend-Color blocks (six blocks of 150 trials), participants responded depending on which color was in the majority. Two colors were mapped to each response (four colors

total), and dots were a mixture of one color from each possible response. Dots colors were approximately isolument (uncalibrated; RGB: [239, 143, 143], [191, 239, 143], [143, 239, 239], [191, 143, 239]), and we counterbalanced their assignment to responses across participants.

In Attend-Motion blocks (six blocks of 45 trials), participants responded based on the dot motion instead of the dot color. Dot motion consisted of a mixture between dots moving coherently (either left or right) and dots moving in a random direction. Attend-Motion blocks were shorter because they acted to reinforce motion sensitivity and provide a test of stimulus-dependent effects.

Critically, dots always had color and motion, and we varied the strength of color coherence (percentage of dots in the majority) and motion coherence (percentage of dots moving coherently) across trials. Our previous experiments have found that in Attend-Color blocks, participants are still influenced by motion information, introducing a response conflict when color and motion are associated with different responses [10]. Target coherence (e.g., color coherence during Attend-Color) was linearly spaced between 65% and 95% with 5 levels, and distractor congruence (signed coherence relative to the target response) was linearly spaced between −95% and 95% with 5 levels. In order to increase the salience of the motion dimension relative to the color dimension, the display was large (~10 degrees of visual angle) and dots moved quickly (~10 degrees of visual angle per second).

Participants had 1.5 seconds from the onset of the stimulus to make their response, and the RDK stayed on the screen for this full duration to avoid confusing reaction time and visual stimulation (the fixation cross changed from white to gray to register the response). The inter-trial interval was uniformly sampled from 1.0, 1.5, or 2.0 seconds. This ITI was relatively short in order to maximize the behavioral effect, and because efficiency simulations showed that it increased power to detect parametric effects of target and distractor coherence (e.g., relative to a more standard 5 second ITI). The fixation cross changed from gray to white for the last 0.5 seconds before the stimulus to provide an alerting cue.

Procedure

Before the scanning session, participants provided consent and practiced the task in a mock MRI scanner. First, participants learned to associate four colors with two button presses (two colors for each response). After being instructed on the color-button mappings, participants practiced the task with feedback (correct, error, or 1.5 second time-out). Errors or time-out feedback were accompanied with a diagram of the color-button mappings. Participants performed 50 trials with full color coherence, and then 50 trials with variable color coherence, all with 0% motion coherence. Next, participants practiced the motion task. After being shown the motion mappings, participants performed 50 trials with full motion coherence, and then 50 trials with variable motion coherence, all with 0% color coherence. Finally, participants practiced 20 trials of the Attend-Color task and 20 trials of Attend-Motion tasks with variable color and motion coherence (same as scanner task).

Following the twelve blocks of the scanner task, participants underwent localizers for color and motion, based on the tasks used in our previous experiments [30]. Both localizers were block designs, alternating between 16 seconds of feature present and 16 seconds of feature absent for seven cycles. For the color localizer, participants saw an aperture the same size as the task, either filled with colored squares that were resampled every second during stimulus-on ('Mondrian stimulus'), or luminance-matched gray squares that were similarly resampled during stimulus-off. For the motion localizer, participants saw white dots that were moving with full coherence in a different direction every second during stimulus-on, or still dots for stimulus-off. No responses were required during the localizers.

MRI sequence

We scanned participants with a Siemens Prisma 3T MR system. We used the following sequence parameters for our functional runs: field of view (FOV) = 211 mm × 211 mm (60 slices), voxel size = 2.4 mm^3, repetition time (TR) = 1.2 sec with interleaved multiband acquisitions (acceleration factor 4), echo time (TE) = 33 ms, and flip angle (FA) = 62°. Slices were acquired anterior to posterior, with an auto-aligned slice orientation tilted 15° relative to the AC/PC plane. At the start of the imaging session, we collected a high-resolution structural MPRAGE with the following sequence parameters: FOV = 205 mm × 205 mm (192 slices), voxel size = 0.8 mm^3, TR = 2.4 sec, TE1 = 1.86 ms, TE2 = 3.78 ms, TE3 = 5.7 ms, TE4 = 7.62, and FA = 7°. At the end of the scan, we collected a field map for susceptibility distortion correction (TR = 588ms, TE1 = 4.92 ms, TE2 = 7.38 ms, FA = 60°).

fMRI preprocessing

We preprocessed our structural and functional data using fMRIprep (v20.2.6; [122] based on the Nipype platform [123]. We used FreeSurfer and ANTs to nonlinearly register structural T1w images to the MNI152NLin6Asym template (resampling to 2mm). To preprocess functional T2w images, we applied susceptibility distortion correction using fMRIprep, co-registered our functional images to our T1w images using FreeSurfer, and slice-time corrected to the midpoint of the acquisition using AFNI. We then registered our images into MNI152NLin6Asym space using the transformation that ANTs computed for the T1w images, resampling our functional images in a single step. For univariate analyses, we smoothed our functional images using a Gaussian kernel (8mm FWHM, as dACC responses often have a large spatial extent). For multivariate analyses, we worked in the unsmoothed template space (see below).

fMRI univariate analyses

We used SPM12 (v7771) for our univariate general linear model (GLM) analyses. Due to high trial-to-trial collinearity from to our short ITIs, we performed all analyses across trials, rather than extracting single-trial estimates. Our regression models used whole trials as events (i.e., a 1.5 second boxcar aligned to the stimulus onset). We parametrically modulated these events with standardized trial-level predictors (e.g., linear-coded target coherence, or contrast-coded errors), and then convolved these predictors with SPM's canonical HRF, concatenating our voxel timeseries across runs. We included nuisance regressors to capture 1) run intercepts and 2) the average timeseries across white matter and CSF (as segmented

by fMIRPrep). To reduce the influence of motion artifacts, we used robust weighted least-squares [124,125], a procedure for optimally down-weighting noisy TRs.

We estimated contrast maps at the subject-level, which we then used for one-sample t-tests at the group-level. We controlled for family-wise error rate using threshold-free cluster enhancement [126], testing whether voxels have an unlikely degree of clustering under a randomized null distribution (Implemented in PALM [127]; 10,000 randomizations). To improve the specificity of our coverage (e.g., reducing white-matter contributions) and to facilitate our inference about functional networks (see below), we limited these analyses to voxels within the Kong2022 whole-brain parcellation [54,55]. This parcellation assigns regional labels to parcels (e.g., whether parcels are in 'SPL' or 'IPS'), which was used through-out to generate ROIs. Surface renders were generated using surfplot [128,129], projecting from MNI space to the Human Connectome Project's fsLR space (164,000 vertices).

dACC longitudinal axis analyses

To characterize the spatial organization of target difficulty and distractor congruence signals in dACC, we constructed an analysis mask that provided broad coverage across cingulate cortex and preSMA. This mask was constructed by 1) getting a meta-analytic mask of cingulate responses co-occurring with 'cognitive control' (Neurosynth uniformity test; [130], and taking any parcels from the whole-brain Schaefer parcellation (400 parcels; [54,55] that had a 50 voxel overlap with this meta-analytic mask. We used this parcellation because it provided more selective gray matter coverage than the Neurosynth mask alone and it allowed us to categorize voxels membership in putative functional networks.

To characterize the spatial organization within dACC, we first performed PCA on the masked voxel coordinates (y and z), getting a score for eac' voxel's position on the longitudinal axis of this ROI. We then regressed voxel's gradient scores against their regression weights from a model including linear target coherence and distractor congruence (both coded −1 to 1 across difficulty levels). We used linear mixed effects analysis to partially pool across subjects and accommodate within-subject correlations between voxels. Our model predicted gradient score from the linear and quadratic expansions of the target and distractor betas (gradientScore ~ 1 + target + target^2 + distractor + distractor^2 + (1 + target + target^2 + distractor + distractor^2 | subject)). To characterize the network-dependent organization of target and distractor encoding, we complexity-penalized fits between models that either 1) predicted target or distractor betas from linear and quadratic expansions of gradient scores, or 2) predicted target/distractor betas from dummy-coded network assignment from the Schaefer parcellation, comparing these models against a model that used both network and gradient information.

Encoding Geometry Analysis (EGA)

We adapted functions from the pcm-toolbox and rsatoolbox packages for our multivariate analyses [65,131]. We first fit whole-brain GLMs without spatial smoothing, separately for each scanner run. These GLMs estimated the parametric relationship between task variables and BOLD response (e.g., linearly coded target coherence), with a pattern

of these parametric betas across voxels reflecting linear encoding subspace [59]. Within each Schaefer parcel (n=400), we spatially pre-whitened these beta maps, reducing noise correlations between voxels that can inflate pattern similarity and reduce reliability [63]. We then computed the cross-validated Pearson correlation, estimating the similarity of whitened patterns across scanner runs. We used a correlation metric to estimate the alignment between encoding subspaces, rather than distances between condition patterns, to normalize biases and scaling across stimuli (e.g., greater sensitivity to targets vs distractors) and across time (e.g., representational drift). Note that this analysis approach is related to 'Parallelism Scores' [43], but here we use parametric encoding models and emphasize not only deviations from parallel/orthogonal, but also the direction of alignment between features (e.g., Figures 5 and 7).

We computed subspace alignment between contrasts of interest within each participant, and then tested these against zero at the group level. Since our correlations were less than $r = |0.5|$, we did not transform correlations before analysis. We used a Bayesian t-test to test for orthogonality (bayesFactor toolbox in MATLAB, based on [132]). The Bayes factor from this t-test gives evidence for either non-orthogonality (BF_{10} further from zero) or orthogonality (BF_{10} closer to zero, often defined as the reciprocal BF_{01}). Using a standard prior (Cauchy, width = 0.707), our strongest possible evidence for the orthogonality is $BF_{01} = 5.07$ or equivalently $logBF = -0.705$ (i.e., the Bayes factor when $t_{(28)} = 0$).

Our measure of encoding strength was whether encoding subspaces were reliable across blocks (i.e., whether within-feature encoding pattern correlations across runs were significantly above zero at the group level). We used pattern reliability as a geometric proxy for how well a linear encoder would predict held-out brain data, as reliability provides the similarity between the cross-validated model and the best linear unbiased estimator of the within-sample data. We confirmed through simulations that pattern reliability is a good proxy for the traditional encoding metric of predicting held-out timeseries [59]. However, we found that pattern reliability is more powerful, due to it being much less sensitive to the magnitude of residual variance (these two methods are identical in the noise-free case; see Supplementary Figure 3).

When looking at alignment between two subspaces across parcels, we first selected parcels that significantly encoded both factors ('jointly reliable parcels', both $p < .001$ uncorrected). This selection process acts as a thresholded version of classical correlation disattenuation [66,67], and we confirmed through simulations this selection procedure does not increase type 1 error rate. We corrected for multiple comparisons using non-parametric max-statistic randomization tests across parcels [133]. These randomization tests determine the likelihood of an observed effects under a null distribution generated by randomizing the sign of alignment correlations across participants and parcels 10,000 times. Within each randomization, we saved the max and min group-level effect sizes across all parcels, estimating the strongest parcel-wise effect we'd expect if there wasn't a systematic group-level effect.

Some of our first-level models had non-zero levels of multicollinearity, due to conditioning on trials without omission errors or when including feature coherence and performance in the same model. Multicollinearity was far below standard thresholds for concern, generally

(much) less than 5 for a standard threshold is 30 (ratio between largest and smallest singular values in the design matrix, using MATLAB colintest; [134]). However, we wanted to confirm that predictor correlations wouldn't bias our estimates of encoding alignment. We simulated data from a pattern component model [131] in which two variables were orthogonal (generated by separate variance components with no covariance), but were generated from a design matrix with correlated predictors. These simulations confirmed that cross-validated similarity measures were not biased by predictor correlations (Supplementary Figure 16).

To provide further validation for our parametric analyses, we estimated encoding profiles using an analysis with fewer parametric assumptions. First, we fit a GLM with separate predictors for levels of target and distractor evidence ('Evidence Levels' GLM in Table 2). Next, we estimated a traditional (cross-validated) representational dissimilarity matrix across all feature levels. Finally, we visualized these encoding profiles using classical multidimensional scaling (eigenvalue decomposition; see Figure 4B and Supplementary Figure 8).

Multivariate Connectivity Analysis

To estimate what information is plausibly communicated between cortical areas, we measured the alignment between multivariate connectivity patterns (i.e., the 'communication subspace' with a seed region, [108]) and local feature encoding patterns. First, we residualized our Performance GLM (see Table 2) from a seed region's timeseries, and then extracted the variance-weighted average timecourse (i.e., the leading eigenvariate from SPM12's volume of interest function). We then re-estimated our Performance GLM, including the block-specific seed timeseries as a covariate, and performed EGA between seed and coherence patterns (see Equations 1–3). We found convergent results when we residualized a quadratic expansion of our Performance GLM from our seed region, helping to confirm that connectivity alignment wasn't due to underfitting. Note that our cross-validated EGA helps avoid false positives due to any correlations in the design matrix (see above). We localized this connectivity analysis to color- and motion-sensitive cortex by finding the bilateral Kong22 parcels that roughly covered the area of strongest block-level contrast during our localizer runs. Note that these analyses reflect 'functional connectivity', which is susceptible to unmodelled confounders [85].

To estimate the mediation of lPFC connectivity by IPS, we compared models in which just lPFC or just IPS were used for EGA against a model where both seeds were included as covariates in the same model ([86]; see Equations 4–5). Our test of mediation was the group-level difference in lPFC seed-coherence alignment before and after including IPS. While these analyses are inherently cross-sectional (i.e., lPFC and IPS are measured at the same time), we supplemented these analyses by showing that the mediating effect of IPS on lPFC was much larger than the mediating effect of lPFC on IPS (see Figure 7B; Supplementary Figure 12). Unlike traditional mediation analyses looking at the first-order change in regression estimates, our analysis looks at the second-order change in the multivariate alignment between regression estimates, using the same core rationale.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Data Availability:

Data will be made available upon publication.

References

1. Musslick S, Shenhav A, Botvinick M & Cohen J A Computational Model of Control Allocation based on the Expected Value of Control. in 2nd Multidisciplinary Conference on Reinforcement Learning and Decision Making (2015).

2. Badre D, Bhandari A, Keglovits H & Kikumoto A The dimensionality of neural representations for control. Curr Opin Behav Sci 38, 20–28 (2021). [PubMed: 32864401]

3. Danielmeier C, Eichele T, Forstmann BU, Tittgemeyer M & Ullsperger M Posterior medial frontal cortex activity predicts post-error adaptations in task-related visual and motor areas. J. Neurosci. 31, 1780–1789 (2011). [PubMed: 21289188]

4. Egner T Multiple conflict-driven control mechanisms in the human brain. Trends Cogn. Sci. 12, 374–380 (2008). [PubMed: 18760657]

5. Ritz H, Leng X & Shenhav A Cognitive Control as a Multivariate Optimization Problem. J. Cogn. Neurosci. 34, 569–591 (2022). [PubMed: 35061027]

6. Friedman NP & Miyake A Unity and diversity of executive functions: Individual differences as a window on cognitive structure. Cortex 86, 186–204 (2017). [PubMed: 27251123]

7. Danielmeier C & Ullsperger M Post-error adjustments. Front. Psychol. 2, 233 (2011). [PubMed: 21954390]

8. Fischer AG, Nigbur R, Klein TA, Danielmeier C & Ullsperger M Cortical beta power reflects decision dynamics and uncovers multiple facets of post-error adaptation. Nat. Commun. 9, 5038 (2018). [PubMed: 30487572]

9. Leng X, Yee D, Ritz H & Shenhav A Dissociable influences of reward and punishment on adaptive cognitive control. PLoS Comput. Biol. 17, e1009737 (2021). [PubMed: 34962931]

10. Ritz H & Shenhav A Humans reconfigure target and distractor processing to address distinct task demands. bioRxiv 2021.09.08.459546 (2021) doi:10.1101/2021.09.08.459546.

11. Kalman RE On the general theory of control systems. IFAC Proceedings Volumes 1, 491–502 (1960).

12. Shenhav A, Botvinick MM & Cohen JD The expected value of control: an integrative theory of anterior cingulate cortex function. Neuron 79, 217–240 (2013). [PubMed: 23889930]

13. MacDonald AW 3rd, Cohen JD, Stenger VA & Carter CS Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. Science 288, 1835–1838 (2000). [PubMed: 10846167]

14. Smith EH et al. Widespread temporal coding of cognitive control in the human prefrontal cortex. Nat. Neurosci. 66, 83 (2019).

15. Menon V & D'Esposito M The role of PFC networks in cognitive control and executive function. Neuropsychopharmacology 1–14 (2021) doi:10.1038/s41386-021-01152-w.

16. Kerns JG et al. Anterior cingulate conflict monitoring and adjustments in control. Science 303, 1023–1026 (2004). [PubMed: 14963333]

17. Gottlieb J, Cohanpour M, Li Y, Singletary N & Zabeh E Curiosity, information demand and attentional priority. Current Opinion in Behavioral Sciences 35, 83–91 (2020).

18. Gordon EM et al. Precision Functional Mapping of Individual Human Brains. Neuron 95, 791–807.e7 (2017). [PubMed: 28757305]

19. Gratton C, Laumann TO, Gordon EM, Adeyemo B & Petersen SE Evidence for Two Independent Factors that Modify Brain Networks to Meet Task Goals. Cell Rep. 17, 1276–1288 (2016). [PubMed: 27783943]

20. Kragel PA et al. Generalizable representations of pain, cognitive control, and negative emotion in medial frontal cortex. Nat. Neurosci. 21, 283–289 (2018). [PubMed: 29292378]

21. Fu Z et al. The geometry of domain-general performance monitoring in the human medial frontal cortex. Science 376, eabm9922 (2022). [PubMed: 35511978]

22. Vermeylen L et al. Shared Neural Representations of Cognitive Conflict and Negative Affect in the Medial Frontal Cortex. J. Neurosci. 40, 8715–8725 (2020). [PubMed: 33051353]

23. Brown JW & Braver TS Learned predictions of error likelihood in the anterior cingulate cortex. Science 307, 1118–1121 (2005). [PubMed: 15718473]

24. Rushworth MF & Behrens TE Choice, uncertainty and value in prefrontal and cingulate cortex. Nat. Neurosci. 11, 389–397 (2008). [PubMed: 18368045]

25. Grinband J et al. The dorsal medial frontal cortex is sensitive to time on task, not response conflict or error likelihood. Neuroimage 57, 303–311 (2011). [PubMed: 21168515]

26. Yarkoni T, Barch DM, Gray JR, Conturo TE & Braver TS BOLD correlates of trial-by-trial reaction time variability in gray and white matter: a multi-study fMRI analysis. PLoS One 4, e4257 (2009). [PubMed: 19165335]

27. Mumford JA et al. The response time paradox in functional magnetic resonance imaging analyses. bioRxiv 2023.02.15.528677 (2023) doi:10.1101/2023.02.15.528677.

28. Pylyshyn ZW & Storm RW Tracking multiple independent targets: evidence for a parallel tracking mechanism. Spat. Vis. 3, 179–197 (1988). [PubMed: 3153671]

29. Beldzik E & Ullsperger M A thin line between conflict and reaction time effects on EEG and fMRI brain signals. bioRxiv 2023.02.14.528515 (2023) doi:10.1101/2023.02.14.528515.

30. Shenhav A, Straccia MA, Musslick S, Cohen JD & Botvinick MM Dissociable neural mechanisms track evidence accumulation for selection of attention versus action. Nat. Commun. 9, 2485 (2018). [PubMed: 29950596]

31. Taren AA, Venkatraman V & Huettel SA A parallel functional topography between medial and lateral prefrontal cortex: evidence and implications for cognitive control. J. Neurosci. 31, 5026–5031 (2011). [PubMed: 21451040]

32. Venkatraman V, Rosati AG, Taren AA & Huettel SA Resolving response, decision, and strategic control: evidence for a functional topography in dorsomedial prefrontal cortex. J. Neurosci. 29, 13158–13164 (2009). [PubMed: 19846703]

33. Fu Z et al. Single-Neuron Correlates of Error Monitoring and Post-Error Adjustments in Human Medial Frontal Cortex. Neuron 101, 165–177.e5 (2019). [PubMed: 30528064]

34. Zarr N & Brown JW Hierarchical error representation in medial prefrontal cortex. Neuroimage 124, 238–247 (2016). [PubMed: 26343320]

35. Ebitz BR et al. Human dorsal anterior cingulate neurons signal conflict by amplifying task-relevant information. bioRxiv 2020.03.14.991745 (2020) doi:10.1101/2020.03.14.991745.

36. Flesch T, Juechems K, Dumbalska T, Saxe A & Summerfield C Orthogonal representations for robust context-dependent task performance in brains and neural networks. Neuron 0, (2022).

37. Minxha J, Adolphs R, Fusi S, Mamelak AN & Rutishauser U Flexible recruitment of memory-based choice representations by the human medial frontal cortex. Science 368, (2020).

38. Ebitz RB & Hayden BY The population doctrine in cognitive neuroscience. Neuron (2021) doi:10.1016/j.neuron.2021.07.011.

39. Cunningham JP & Yu BM Dimensionality reduction for large-scale neural recordings. Nat. Neurosci. 17, 1500–1509 (2014). [PubMed: 25151264]

40. Rigotti M et al. The importance of mixed selectivity in complex cognitive tasks. Nature 497, 585–590 (2013). [PubMed: 23685452]

41. Mante V, Sussillo D, Shenoy KV & Newsome WT Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503, 78–84 (2013). [PubMed: 24201281]

42. Kayser AS, Erickson DT, Buchsbaum BR & D'Esposito M Neural representations of relevant and irrelevant features in perceptual decision making. J. Neurosci. 30, 15778–15789 (2010). [PubMed: 21106817]

43. Bernardi S et al. The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. Cell 0, (2020).

44. Panichello MF & Buschman TJ Shared mechanisms underlie the control of working memory and attention. Nature 1–5 (2021) doi:10.1038/s41586-021-03390-w.

45. Takagi Y, Hunt LT, Woolrich MW, Behrens TE & Klein-Flügge MC Adapting non-invasive human recordings along multiple task-axes shows unfolding of spontaneous and over-trained choice. Elife 10, (2021).

46. Cohen JD, Dunbar K & McClelland JL On the control of automatic processes: a parallel distributed processing account of the Stroop effect. Psychol. Rev. 97, 332–361 (1990). [PubMed: 2200075]

47. Woolgar A, Hampshire A, Thompson R & Duncan J Adaptive coding of task-relevant information in human frontoparietal cortex. J. Neurosci. 31, 14592–14599 (2011). [PubMed: 21994375]

48. Nee DE, Wager TD & Jonides J Interference resolution: insights from a meta-analysis of neuroimaging tasks. Cogn. Affect. Behav. Neurosci. 7, 1–17 (2007). [PubMed: 17598730]

49. Shenhav A, Straccia MA, Botvinick MM & Cohen JD Dorsal anterior cingulate and ventromedial prefrontal cortex have inverse roles in both foraging and economic choice. Cogn. Affect. Behav. Neurosci. 1127–1139 (2016) doi:10.3758/s13415-016-0458-8. [PubMed: 27580609]

50. Fleming SM, van der Putten EJ & Daw ND Neural mediators of changes of mind about perceptual decisions. Nat. Neurosci. 21, 617–624 (2018). [PubMed: 29531361]

51. Shenhav A & Karmarkar UR Dissociable components of the reward circuit are involved in appraisal versus choice. Sci. Rep. 9, 1958 (2019). [PubMed: 30760824]

52. Clairis N & Pessiglione M Value, confidence, deliberation: a functional partition of the medial prefrontal cortex demonstrated across rating and choice tasks. J. Neurosci. 42, 5580–5592 (2022). [PubMed: 35654606]

53. Yeo BTT et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J. Neurophysiol. 106, 1125–1165 (2011). [PubMed: 21653723]

54. Kong R et al. Individual-Specific Areal-Level Parcellations Improve Functional Connectivity Prediction of Behavior. Cereb. Cortex 31, 4477–4500 (2021). [PubMed: 33942058]

55. Schaefer A et al. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. Cereb. Cortex 28, 3095–3114 (2018). [PubMed: 28981612]

56. Kong R et al. Spatial Topography of Individual-Specific Cortical Networks Predicts Human Cognition, Personality, and Emotion. Cereb. Cortex 29, 2533–2551 (2019). [PubMed: 29878084]

57. Bisley JW & Mirpour K The neural instantiation of a priority map. Curr. Opin. Psychol. 29, 108–112 (2019). [PubMed: 30731260]

58. Yantis S & Serences JT Cortical mechanisms of space-based and object-based attentional control. Curr. Opin. Neurobiol. 13, 187–193 (2003). [PubMed: 12744972]

59. Kriegeskorte N & Diedrichsen J Peeling the Onion of Brain Representations. Annu. Rev. Neurosci. 42, 407–432 (2019). [PubMed: 31283895]

60. Cohen MR & Maunsell JHR A neuronal population measure of attention predicts behavioral performance on individual trials. J. Neurosci. 30, 15241–15253 (2010). [PubMed: 21068329]

61. Libby A & Buschman TJ Rotational dynamics reduce interference between sensory and memory representations. Nat. Neurosci. 1–12 (2021) doi:10.1038/s41593-021-00821-9. [PubMed: 33303973]

62. Kimmel DL, Elsayed GF, Cunningham JP & Newsome WT Value and choice as separable and stable representations in orbitofrontal cortex. Nat. Commun. 11, 3466 (2020). [PubMed: 32651373]

63. Walther A et al. Reliability of dissimilarity measures for multi-voxel pattern analysis. Neuroimage 137, 188–200 (2016). [PubMed: 26707889]

64. Diedrichsen J & Kriegeskorte N Representational models: A common framework for understanding encoding, pattern-component, and representational-similarity analysis. PLoS Comput. Biol. 13, e1005508 (2017). [PubMed: 28437426]

65. Nili H et al. A toolbox for representational similarity analysis. PLoS Comput. Biol. 10, e1003553 (2014). [PubMed: 24743308]

66. Spearman C The Proof and Measurement of Association between Two Things. Am. J. Psychol. 100, 441–471 (1987). [PubMed: 3322052]

67. Thornton MA & Mitchell JP Consistent Neural Activity Patterns Represent Personally Familiar People. J. Cogn. Neurosci. 29, 1583–1594 (2017). [PubMed: 28557690]

68. Hunt LT et al. Mechanisms underlying cortical activity during value-guided choice. Nat. Neurosci. 15, 470–6, S1–3 (2012). [PubMed: 22231429]

69. Kayser AS, Buchsbaum BR, Erickson DT & D'Esposito M The functional anatomy of a perceptual decision in the human brain. J. Neurophysiol. 103, 1179–1194 (2010). [PubMed: 20032247]

70. Kriegeskorte N, Goebel R & Bandettini P Information-based functional brain mapping. Proc. Natl. Acad. Sci. U. S. A. 103, 3863–3868 (2006). [PubMed: 16537458]

71. Woolgar A, Williams MA & Rich AN Attention enhances multi-voxel representation of novel objects in frontal, parietal and visual cortices. Neuroimage 109, 429–437 (2015). [PubMed: 25583612]

72. Woolgar A, Afshar S, Williams MA & Rich AN Flexible Coding of Task Rules in Frontoparietal Cortex: An Adaptive System for Flexible Cognitive Control. J. Cogn. Neurosci. 27, 1895–1911 (2015). [PubMed: 26058604]

73. Jackson J, Rich AN, Williams MA & Woolgar A Feature-selective Attention in Frontoparietal Cortex: Multivoxel Codes Adjust to Prioritize Task-relevant Information. J. Cogn. Neurosci. 29, 310–321 (2017). [PubMed: 27626230]

74. Woolgar A, Thompson R, Bor D & Duncan J Multi-voxel coding of stimuli, rules, and responses in human frontoparietal cortex. Neuroimage 56, 744–752 (2011). [PubMed: 20406690]

75. Aoi MC, Mante V & Pillow JW Prefrontal cortex exhibits multidimensional dynamic encoding during decision-making. Nat. Neurosci. (2020) doi:10.1038/s41593-020-0696-5.

76. Pagan M et al. A new theoretical framework jointly explains behavioral and neural variability across subjects performing flexible decision-making. bioRxiv 2022.11.28.518207 (2022) doi:10.1101/2022.11.28.518207.

77. Miller EK & Cohen JD An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. 24, 167–202 (2001). [PubMed: 11283309]

78. Stringer C et al. Spontaneous behaviors drive multidimensional, brainwide activity. Science 364, 255 (2019). [PubMed: 31000656]

79. Goldman-Rakic PS Topography of cognition: parallel distributed networks in primate association cortex. Annu. Rev. Neurosci. 11, 137–156 (1988). [PubMed: 3284439]

80. Corbetta M & Shulman GL Control of goal-directed and stimulus-driven attention in the brain. Nat. Rev. Neurosci. 3, 201–215 (2002). [PubMed: 11994752]

81. Suzuki M & Gottlieb J Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. Nat. Neurosci. 16, 98–104 (2013). [PubMed: 23242309]

82. Kastner S & Ungerleider LG Mechanisms of visual attention in the human cortex. Annu. Rev. Neurosci. 23, 315–341 (2000). [PubMed: 10845067]

83. Kay KN & Yeatman JD Bottom-up and top-down computations in word- and face-selective cortex. Elife 6, (2017).

84. Saalmann YB, Pigarev IN & Vidyasagar TR Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. Science 316, 1612–1615 (2007). [PubMed: 17569863]

85. Reid AT et al. Advancing functional connectivity research from association to causation. Nat. Neurosci. 22, 1751–1760 (2019). [PubMed: 31611705]

86. MacKinnon DP, Fairchild AJ & Fritz MS Mediation analysis. Annu. Rev. Psychol. 58, 593–614 (2007). [PubMed: 16968208]

87. Friedman NP & Robbins TW The role of prefrontal cortex in cognitive control and executive function. Neuropsychopharmacology 1–18 (2021) doi:10.1038/s41386-021-01132-0.

88. Holroyd CB & McClure SM Hierarchical control over effortful behavior by rodent medial frontal cortex: A computational model. Psychol. Rev. 122, 54–83 (2015). [PubMed: 25437491]

89. Holroyd CB & Yeung N An Integrative Theory of Anterior Cingulate Cortex Function: Option Selection in Hierarchical Reinforcement Learning. Neural Basis of Motivational and Cognitive Control 332–349 Preprint at 10.7551/mitpress/9780262016438.003.0018 (2011).

90. Vassena E, Deraeve J & Alexander WH Predicting Motivation: Computational Models of PFC Can Explain Neural Coding of Motivation and Effort-based Decision-making in Health and Disease. J. Cogn. Neurosci. 29, 1633–1645 (2017). [PubMed: 28654358]

91. Koechlin E & Summerfield C An information theoretical approach to prefrontal executive function. Trends Cogn. Sci. 11, 229–235 (2007). [PubMed: 17475536]

92. Badre D & D'Esposito M Is the rostro-caudal axis of the frontal lobe hierarchical? Nat. Rev. Neurosci. 10, 659–669 (2009). [PubMed: 19672274]

93. Badre D & Nee DE Frontal Cortex and the Hierarchical Control of Behavior. Trends Cogn. Sci. 22, 170–188 (2018). [PubMed: 29229206]

94. Serences JT & Yantis S Spatially selective representations of voluntary and stimulus-driven attentional priority in human occipital, parietal, and frontal cortex. Cereb. Cortex 17, 284–293 (2007). [PubMed: 16514108]

95. Yantis S et al. Transient neural activity in human parietal cortex during spatial attention shifts. Nat. Neurosci. 5, 995–1002 (2002). [PubMed: 12219097]

96. Greenberg AS, Esterman M, Wilson D, Serences JT & Yantis S Control of spatial and feature-based attention in frontoparietal cortex. J. Neurosci. 30, 14330–14339 (2010). [PubMed: 20980588]

97. Esterman M, Chiu Y-C, Tamber-Rosenau BJ & Yantis S Decoding cognitive control in human parietal cortex. Proc. Natl. Acad. Sci. U. S. A. 106, 17974–17979 (2009). [PubMed: 19805050]

98. Serences JT, Schwarzbach J, Courtney SM, Golay X & Yantis S Control of object-based attention in human cortex. Cereb. Cortex 14, 1346–1357 (2004). [PubMed: 15166105]

99. Molenberghs P, Mesulam MM, Peeters R & Vandenberghe RRC Remapping attentional priorities: differential contribution of superior parietal lobule and intraparietal sulcus. Cereb. Cortex 17, 2703–2712 (2007). [PubMed: 17264251]

100. Adam KCS & Serences JT History modulates early sensory processing of salient distractors. J. Neurosci. (2021) doi:10.1523/JNEUROSCI.3099-20.2021.

101. Soutschek A, Stelzel C, Paschke L, Walter H & Schubert T Dissociable effects of motivation and expectancy on conflict processing: an fMRI study. J. Cogn. Neurosci. 27, 409–423 (2015). [PubMed: 25203271]

102. Bisley JW & Goldberg ME Attention, intention, and priority in the parietal lobe. Annu. Rev. Neurosci. 33, 1–21 (2010). [PubMed: 20192813]

103. Rust NC & Cohen MR Priority coding in the visual system. Nat. Rev. Neurosci. 1–13 (2022) doi:10.1038/s41583-022-00582-9. [PubMed: 34754097]

104. Etzel JA, Cole MW, Zacks JM, Kay KN & Braver TS Reward Motivation Enhances Task Coding in Frontoparietal Cortex. Cereb. Cortex 26, 1647–1659 (2016). [PubMed: 25601237]

105. Hall-McMaster S, Muhle-Karbe PS, Myers NE & Stokes MG Reward Boosts Neural Coding of Task Rules to Optimize Cognitive Flexibility. J. Neurosci. 39, 8549–8561 (2019). [PubMed: 31519820]

106. Desimone R & Duncan J Neural mechanisms of selective visual attention. Annu. Rev. Neurosci. 18, 193–222 (1995). [PubMed: 7605061]

107. Lauritzen TZ, D'Esposito M, Heeger DJ & Silver MA Top–down flow of visual spatial attention signals from parietal to occipital cortex. J. Vis. 9, 18–18 (2009). [PubMed: 19761333]

108. Semedo JD, Zandvakili A, Machens CK, Yu BM & Kohn A Cortical Areas Interact through a Communication Subspace. Neuron 102, 249–259.e4 (2019). [PubMed: 30770252]

109. Srinath R, Ruff DA & Cohen MR Attention improves information flow between neuronal populations without changing the communication subspace. bioRxiv 2021.03.31.437940 (2021) doi:10.1101/2021.03.31.437940.

110. Petrides M & Pandya DN Efferent association pathways originating in the caudal prefrontal cortex in the macaque monkey. J. Comp. Neurol. 498, 227–251 (2006). [PubMed: 16856142]

111. Jackson JB, Feredoes E, Rich AN, Lindner M & Woolgar A Concurrent neuroimaging and neurostimulation reveals a causal role for dlPFC in coding of task-relevant information. Commun Biol 4, 588 (2021). [PubMed: 34002006]

112. Vul E, Alvarez G, Tenenbaum J & Black M Explaining human multiple object tracking as resource-constrained approximate inference in a dynamic probabilistic model. in Advances in Neural Information Processing Systems (eds. Bengio Y, Schuurmans D, Lafferty J, Williams C & Culotta A) vol. 22 (Curran Associates, Inc., 2009).

113. Ritz H, Wild CJ & Johnsrude IS Parametric Cognitive Load Reveals Hidden Costs in the Neural Processing of Perfectly Intelligible Degraded Speech. J. Neurosci. 42, 4619–4628 (2022). [PubMed: 35508382]

114. Culham JC, Cavanagh P & Kanwisher NG Attention response functions: characterizing brain areas using fMRI activation during parametric variations of attentional load. Neuron 32, 737–745 (2001). [PubMed: 11719212]

115. Culham JC et al. Cortical fMRI activation produced by attentive tracking of moving targets. J. Neurophysiol. 80, 2657–2670 (1998). [PubMed: 9819271]

116. Howe PD, Horowitz TS, Morocz IA, Wolfe J & Livingstone MS Using fMRI to distinguish components of the multiple object tracking task. J. Vis. 9, 10.1–11 (2009). [PubMed: 19761325]

117. Jovicich J et al. Brain areas specific for attentional load in a motion-tracking task. J. Cogn. Neurosci. 13, 1048–1058 (2001). [PubMed: 11784443]

118. Peck CJ, Jangraw DC, Suzuki M, Efem R & Gottlieb J Reward modulates attention independently of action value in posterior parietal cortex. J. Neurosci. 29, 11182–11191 (2009). [PubMed: 19741125]

119. Wisniewski D, Reverberi C, Momennejad I, Kahnt T & Haynes J-D The Role of the Parietal Cortex in the Representation of Task–Reward Associations. J. Neurosci. 35, 12355–12365 (2015). [PubMed: 26354905]

120. Parro C, Dixon ML & Christoff K The neural basis of motivational influences on cognitive control. Hum. Brain Mapp. 39, 5097–5111 (2018). [PubMed: 30120846]

121. Weichart ER, Turner BM & Sederberg PB A model of dynamic, within-trial conflict resolution for decision making. Psychol. Rev. (2020) doi:10.1037/rev0000191.

122. Esteban O et al. fMRIPrep: a robust preprocessing pipeline for functional MRI. Nat. Methods 16, 111–116 (2019). [PubMed: 30532080]

123. Gorgolewski K et al. Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. Front. Neuroinform. 5, 13 (2011). [PubMed: 21897815]

124. Diedrichsen J & Shadmehr R Detecting and adjusting for artifacts in fMRI time series data. Neuroimage 27, 624–634 (2005). [PubMed: 15975828]

125. Jones MS, Zhu Z, Bajracharya A, Luor A & Peelle JE A multi-dataset evaluation of frame censoring for task-based fMRI. bioRxiv 2021.10.12.464075 (2021) doi:10.1101/2021.10.12.464075.

126. Smith SM & Nichols TE Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. Neuroimage 44, 83–98 (2009). [PubMed: 18501637]

127. Winkler AM, Ridgway GR, Webster MA, Smith SM & Nichols TE Permutation inference for the general linear model. Neuroimage 92, 381–397 (2014). [PubMed: 24530839]

128. Vos de Wael R et al. BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. Commun Biol 3, 103 (2020). [PubMed: 32139786]

129. Gale DJ, Vos de Wael R, Benkarim O & Bernhardt. Surfplot: Publication-ready brain surface figures. (2021). doi:10.5281/zenodo.5567926.

130. Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC & Wager TD Large-scale automated synthesis of human functional neuroimaging data. Nat. Methods 8, 665–670 (2011). [PubMed: 21706013]

131. Diedrichsen J, Yokoi A & Arbuckle SA Pattern component modeling: A flexible approach for understanding the representational structure of brain activity patterns. Neuroimage 180, 119–133 (2018). [PubMed: 28843540]

132. Rouder JN, Morey RD, Speckman PL & Province JM Default Bayes factors for ANOVA designs. J. Math. Psychol. 56, 356–374 (2012).

133. Nichols TE & Holmes AP Nonparametric permutation tests for functional neuroimaging: a primer with examples. Hum. Brain Mapp. 15, 1–25 (2002). [PubMed: 11747097]

134. Belsley DA, Kuh E & Welsch RE Wiley Series in Probability and Statistics. Regression Diagnostics: Identifying Influential Data and Sources of Collinearity 293–300 (1980).
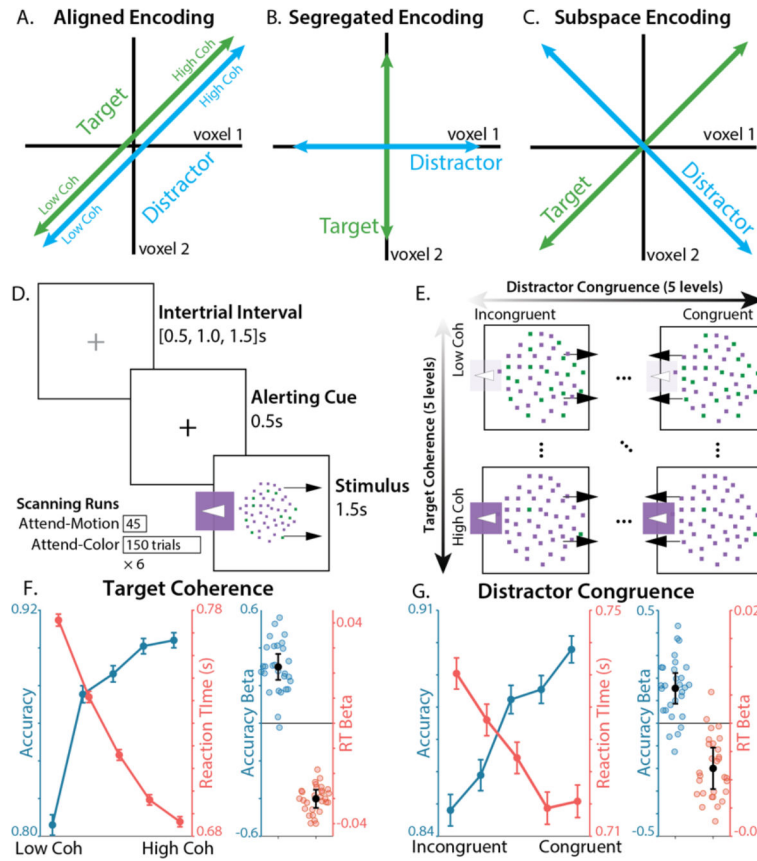
Figure 1. Task and Behavior.

A-C) Three hypothesized encoding schemes. **A)** In *aligned encoding* features are represented similarly, e.g., encode performance variables like error likelihood or time-on-task. **B)** In *segregated encoding* features are encoded independently, in distinct voxel populations (i.e., voxel-level pure selectivity [40]). **C)** In *subspace encoding,* features are encoding independently, in overlapping voxel populations (i.e., voxel-level mixed selectivity). **D)** Participants responded to a color-motion random dot kinematogram (RDK) with a button press. Participants either responded to the left/right motion direction of the RDK (Attend-Motion runs) or based on the majority color (Attend-Color runs; critical condition). **E)** We parametrically and independently manipulated target coherence (% of dots in the majority color) and distractor congruence (motion coherence signed relative to the target response). **F)** Participants were faster and more accurate when the target was more coherent. **G)** Participants were faster and more accurate when the distractor was more congruent with the target. Error bars on line plots reflect within-participant SEM, error bars for regression fixed-effect betas reflect 95% CI.
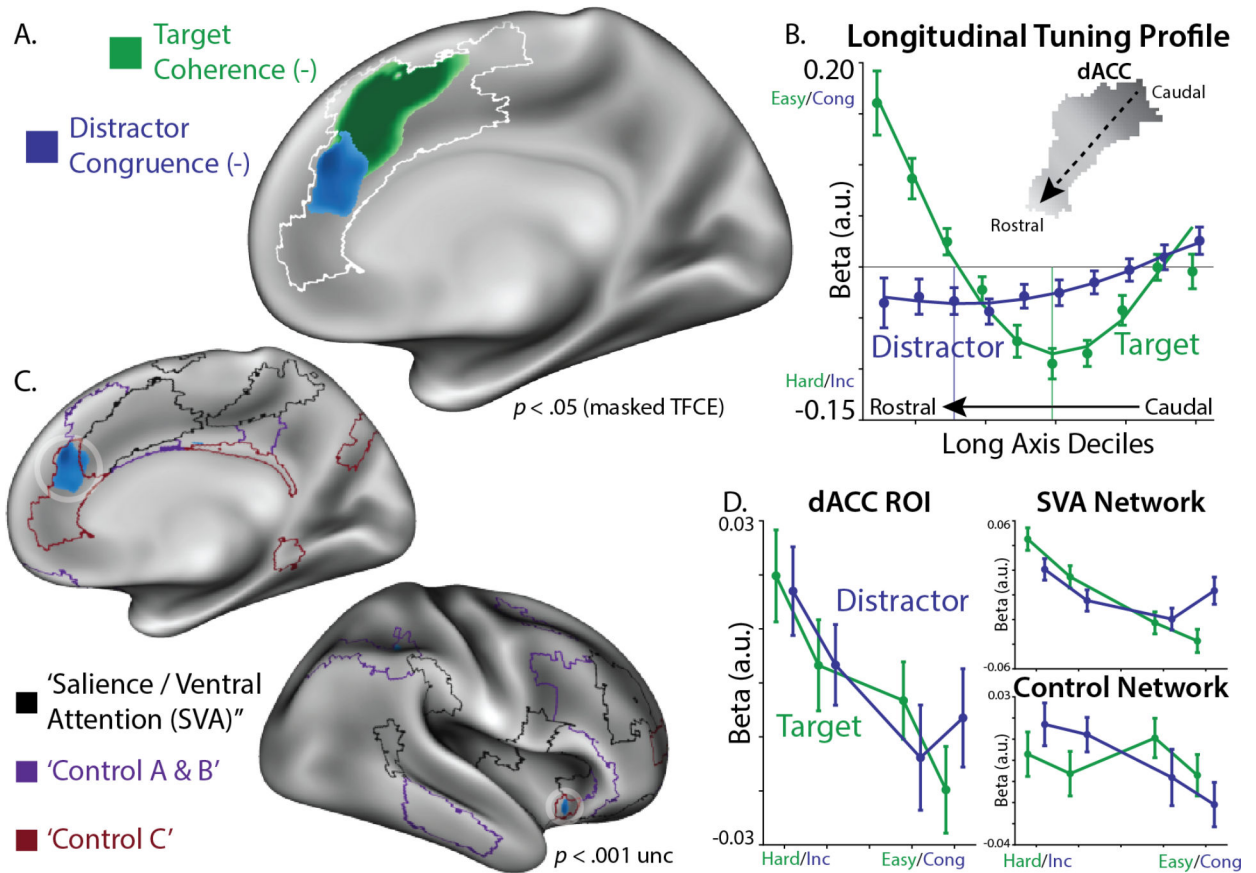
Figure 2. Distinct coding of target and distractor difficulty in dACC.
A) We looked for linear target coherence and distractor congruence signals within an a priori dACC mask (white outline; overlapping Kong22 parcels and medial 'cognitive control' Neurosynth mask). We found that voxels in the most caudal dACC reflected target difficulty (green), more rostral voxels reflected distractor incongruence (blue). Statistical tests are corrected using non-parametric threshold-free cluster enhancement. **B)** We extracted the long axis of the dACC using a PCA of the voxel coordinates. We plotted the target coherence (green) and distractor congruence (blue) along the deciles of this long axis. Fit lines are the quantized predictions from a second-order polynomial regression. We used these regression betas to estimate the minima for target and distractor tuning (i.e., location of strongest difficulty effects), finding that the target difficulty peak (vertical green line) was more caudal than the distractor incongruence peak (vertical blue line). **C)** Plotting the uncorrected whole-brain response, distractor incongruence responses (blue) were strongest within the 'Control C' sub-network (red), both in dACC and anterior insula. **D)** BOLD responses across levels of target coherence and distractor congruence, plotted within the whole dACC ROI (left), or the 'Salience/Ventral Attention (SVA)' network and 'Control' network parcels within the dACC ROI (right). GLMs: A-C: Feature UV, D: Difficulty Levels, see Table 2.
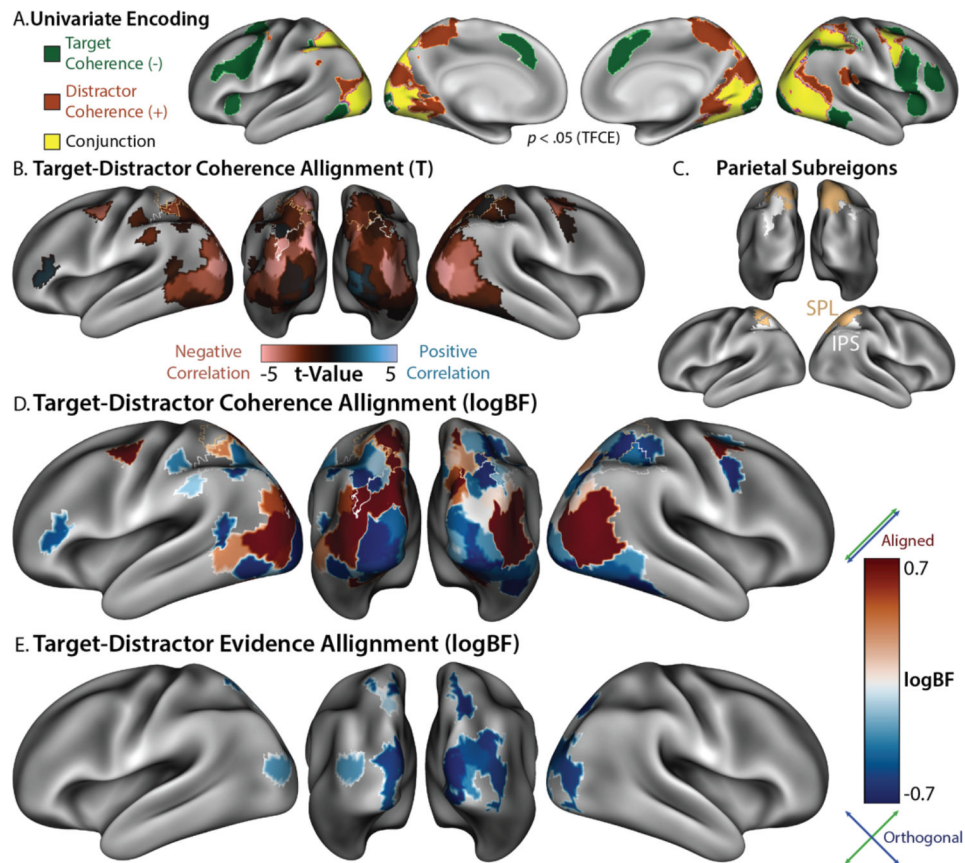
Figure 3. Encoding Geometry Analysis (EGA) dissociates target and distractor encoding.
A) Parametric univariate responses to weak target coherence (green; percentage of dots in majority color), strong distractor coherence (orange; percentage of dots with coherent motion), and their conjunction (yellow). Statistical tests are corrected for multiple comparisons using non-parametric threshold-free cluster enhancement (TFCE). **B)** Encoding alignment within parcels in which target and distractor encoding was jointly reliable (both $p < .001$ uncorrected). Representations were negatively correlated within Superior Parietal Lobule (SPL in gold; Kong22 labels), and uncorrelated within Intraparietal Sulcus (IPS in white; Kong22 labels). **C)** Anatomical labels for parietal regions, based on the labels in the Kong22 parcellation. **D)** Bayesian analyses provide explicit evidence for orthogonality within IPS (i.e., negative BF; theoretical minima: −0.71). **E)** Coherence coded in terms of evidence (i.e., supporting a left vs right choice). Target and distractor evidence encoding overlapped in visual cortex and SPL and was represented orthogonally. GLMs: A: Feature UV, B-E: Feature MV, see Table 2.
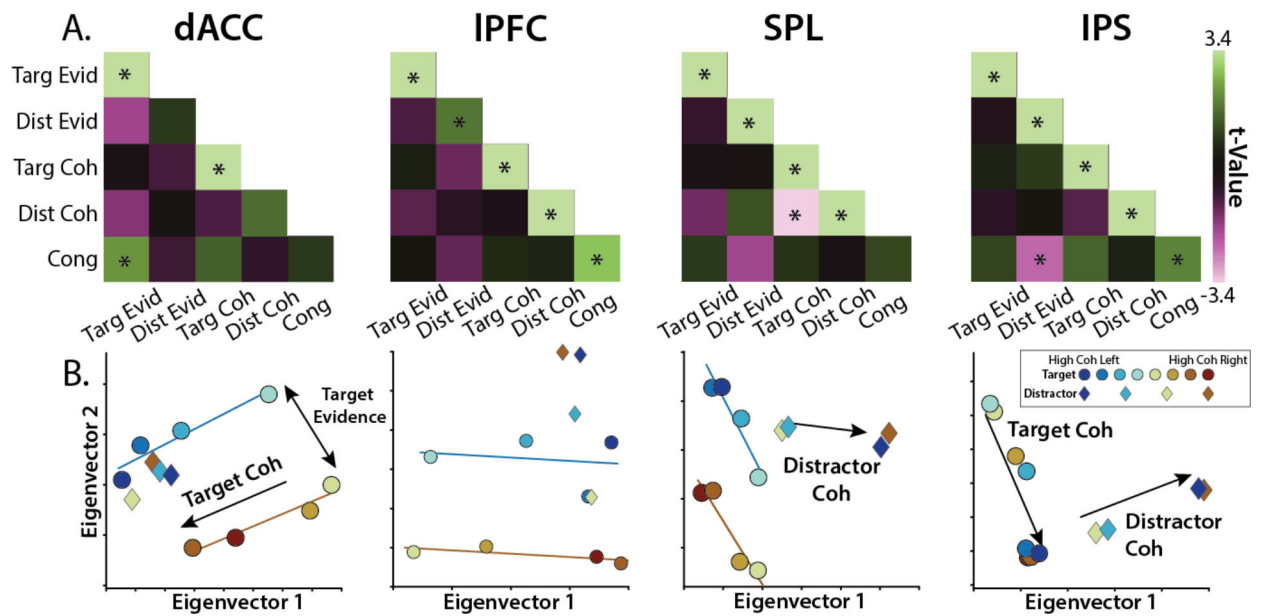
Figure 4. Region-specific feature encoding.
A) Similarity matrices for dACC, lPFC, SPL, and IPS, correlating feature evidence ('Evid'), feature coherence ('Coh'), and feature congruence ('Cong'). Encoding strength on diagonal (right-tailed *p*-value), encoding alignment on off-diagonal (two-tailed *p*-value). **B)** Classical MDS embedding of target (circle) and distractor (diamond) representations at different levels of evidence. Colors denote responses, hues denote coherence. GLMs: A: Feature MV, B: Evidence Levels, see Table 2.
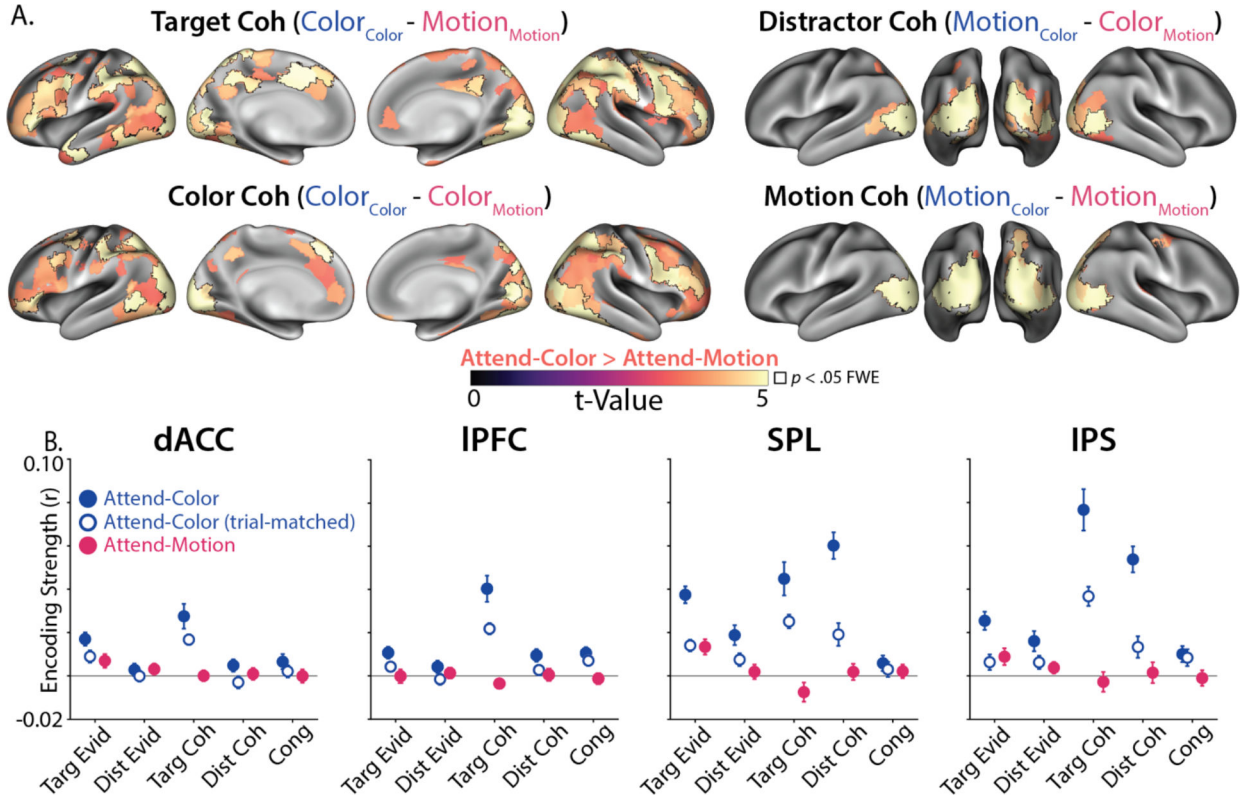
Figure 5. Task-dependent encoding strength.

A) Across cortex, feature coherence encoding was stronger during Attend-Color than Attend-Motion, matched for the same number of trials. Attend-Color had stronger encoding when comparing target coherence (top left), distractor coherence (top right), color coherence (bottom left) and motion coherence (bottom right). Parcels are thresholded at $p < .001$ (uncorrected); outlined parcels are significant at $p < .05$ I (max-statistic randomization test across all parcels). Condition labels in title parentheses are coded 'Feature$_{Task}$'. **B)** Target and distractor coherence information was encoded more strongly during Attend-Color than Attend-Motion in dACC, lPFC, SPL and IPS. Attend-Color encoding plotted from the whole sample (blue fill) and a trial-matched sample (first 45 trials of each run; white fill) In Attend-Motion runs, only target evidence was significantly encoded (magenta). **C)** Target and distractor coherence was not reliably encoded during the Attend-Motion task (liberally thresholded at $p < .01$ uncorrected). GLM: Between-Task, see Table 2.
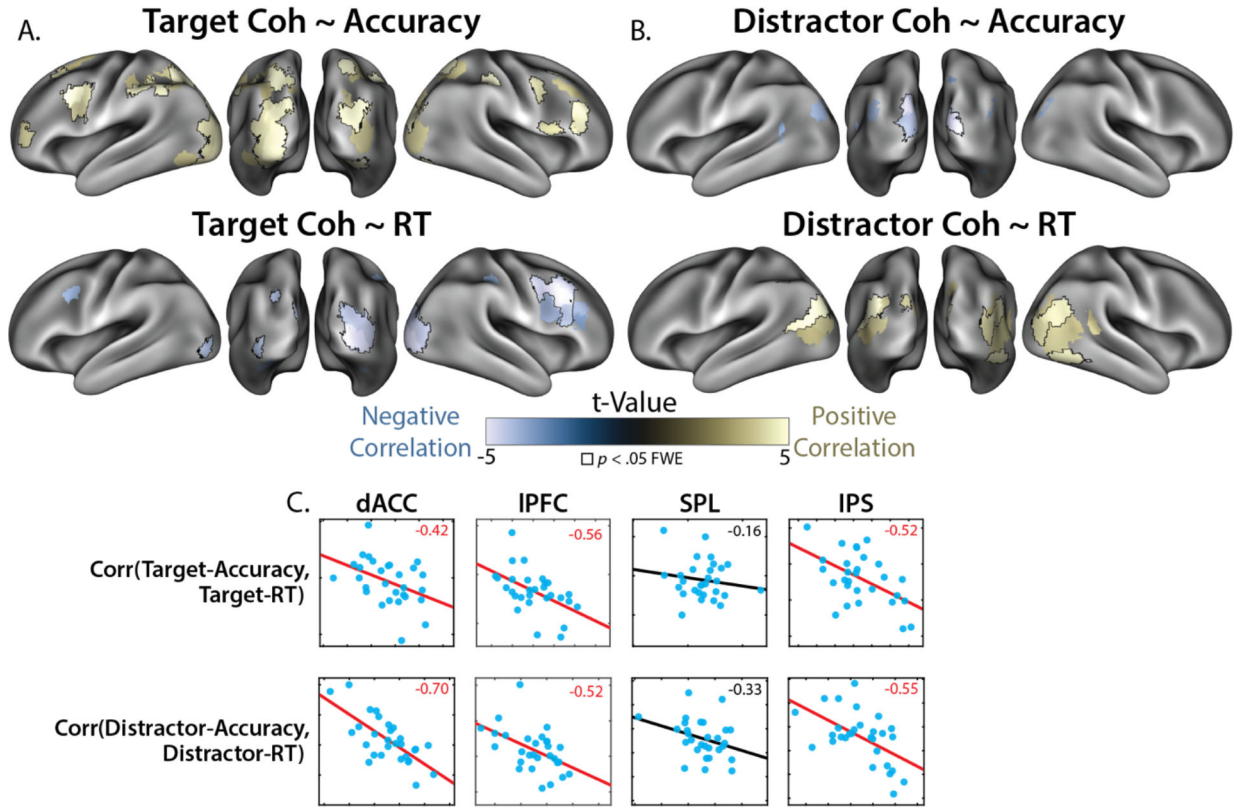
Figure 6. Alignment between feature and performance encoding.

A) Alignment between encoding of target coherence and performance (top row: Accuracy, bottom row: RT). **B)** Alignment between encoding of distractor coherence and performance (top row: Accuracy, bottom row: RT). Across A and B, parcels are thresholded at $p < .001$ (uncorrected, in jointly reliable parcels), and outlined parcels are significant at $p < .05$ I (max-statistic randomization test across jointly reliable parcels). **C)** Individual differences in feature-RT alignment correlated with feature-accuracy alignment across regions (correlation values in top right; $p < .05$ in red). See Supplementary Table 1 for partial correlations controlling for reliability. GLM: Performance, see Table 2.
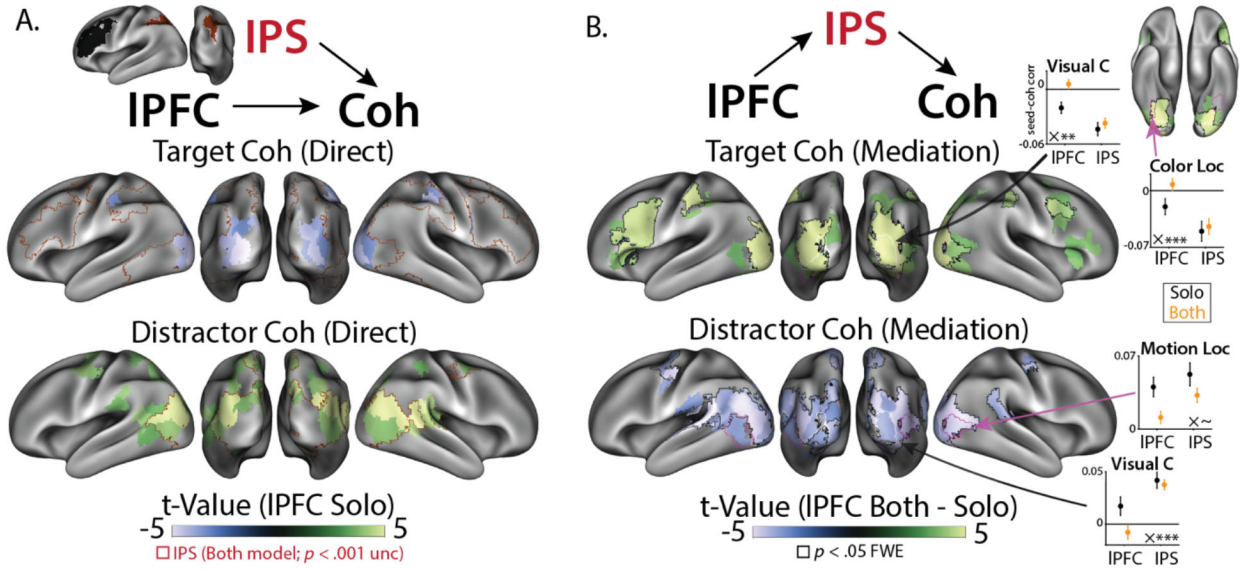
Figure 7. IPS mediates alignment between lPFC and feature encoding.
A) Connectivity patterns from lPFC (color) and IPS (red outline) were aligned with target and distractor coherence patterns ($p < .001$ uncorrected, in jointly reliable parcels). IPS effects are outlined to show overlap, with all effects in a consistent direction to lPFC.
B) lPFC-feature alignment contrasted between lPFC-only model ('Solo') and lPFC + IPS model ('Both'). Including IPS reduced the alignment between lPFC and feature encoding (compare the sign of the main effect in A to the contrast in B). Parcels are thresholded at $p < .001$ (uncorrected, jointly reliable parcels), and outlined parcels are significant at $p < .05$ I (max-statistic randomization test across jointly reliable parcels). Insets graphs: seed-coherence alignment in Solo models (black) and Both model (orange) across visual regions. 'Visual C' is defined by our parcellation [54]; Color and Motion localizers are parcels near the peak response identified during feature localizer runs (see Methods). In general, lPFC alignment was more affected by IPS than IPS alignment was affected by lPFC (inset 'X': difference of differences; $\sim p < .10$, $* p < .01$, $*** p < .001$). GLM: Performance CX, see Table 2.

Table 1.

Feature encoding contrasted across parietal cortex.

Task Feature	SPL	IPS	SPL - IPS
Target Evidence	$t(28) = 10.39$, $p = 4.07 \times 10^{-11}$, logBF $= 8.37$	$t(28) = 5.82$, $p = 3.00 \times 10^{-6}$, logBF $= 3.81$	$t(28) = 3.89$, $p = .000562$, logBF $= 1.75$
Distractor Evidence	$t(28) = 4.42$, $p = 1.34 \times 10^{-4}$, logBF $= 2.30$	$t(28) = 3.62$, $p = 0.0012$, logBF $= 1.47$	$t(28) = 0.896$, $p = .378$, logBF $= -0.545$
Target-Distractor Evidence Alignment	$t(28) = -0.703$, $p = .488$, logBF $= -0.606$	$t(28) = -0.436$, $p = 0.666$, logBF $= -0.667$	$t(28) = -0.145$, $p = .886$, logBF $= -0.701$
Target Coherence	$t(28) = 5.82$, $p = 2.94 \times 10^{-6}$, logBF $= 3.82$	$t(28) = 7.73$, $p = 2.00 \times 10^{-8}$, logBF $= 5.83$	$t(28) = -3.89$, $p = 9.36 \times 10^{-9}$, logBF $= 6.14$
Distractor Coherence	$t(28) = 8.88$, $p = 1.25 \times 10^{-9}$, logBF $= 6.97$	$t(28) = 8.53$, $p = 2.80 \times 10^{-9}$, logBF $= 6.63$	$t(28) = 1.40$, $p = .170$, logBF $= -0.320$
Target-Distractor Coherence Alignment	$t(28) = -4.75$, $p = 5.50 \times 10^{-5}$, logBF $= 2.65$	$t(28) = -1.06$, $p = 0.294$, logBF $= -0.479$	$t(28) = -2.99$, $p = .00580$, logBF $= 0.861$

Encoding of feature evidence and coherence within SPL, within IPS, and contrasted between SPL and IPS. Note the stronger target evidence encoding in SPL, stronger target coherence encoding in IPS, and stronger target-distractor coherence alignment in SPL.

Table 2.

fMRI models.

Model Name	Trial selection	Predictors (z-scored)
Feature UV	No omission errors; run-concatenated	target coherence, distractor coherence, target evidence, distractor evidence, distractor congruence; omission errors (run-concatenated)
Difficulty Levels	No omission errors; run-concatenated	Separate levels (1,2,4,5) of target coherence, separate levels (1,2,4,5) of distractor congruence; omission errors (run-concatenated)
Feature MV	No errors; run-separated	target coherence, distractor coherence, target evidence, distractor evidence, distractor congruence; errors (run-concatenated)
Evidence Levels	No errors; run-separated	Levels (1–5, 6–10) of target evidence, Levels (1,2,4,5) of distractor evidence; errors (run-concatenated)
Between-Task	No errors; run-separated	target coherence, distractor coherence, target evidence, distractor evidence, distractor congruence; errors (run-concatenated); reaction time (run-concatenated)
Performance	No omission errors; run-separated	target coherence, distractor coherence, target evidence, distractor evidence, distractor congruence, reaction time, accuracy; omission errors (run-concatenated)
Performance CX	No omission errors; run-separated	target coherence, distractor coherence, target evidence, distractor evidence, distractor congruence, reaction time, accuracy; omission errors (run-concatenated); seed timeseries (run-separated)

First-level general linear models used for univariate and multivariate fMRI analyses. Coherence: percentage of dots supporting the same response ('unsigned coherence'). Evidence: % dots supporting a rightwards vs leftwards response ('signed coherence'). Distractor Congruence: % dots supporting the same response as the target dimension. All predictors were z-scored within their run. For difficulty and feature levels, we included each level as a separate predictor, with collinearity with the block intercept preventing all levels from being included. For Evidence Levels, targets had greater granularity due to distractors being coded relative to targets (5 levels of congruence led to 5 levels of coherence). For Performance CX, seed timeseries were included as run-separated regressors (see Multivariate Connectivity Analysis in Methods).