Behavioral/Cognitive

Connectivity-Defined Subdivisions of the Intraparietal Sulcus Respond Differentially to Abstraction during Decision-Making

Melissa Newton,^{1,2} Savannah L. Cookson,¹ Mark D'Esposito,^{1,3} and Andrew S. Kayser^{1,2,3}

¹Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, California 94720, ²Department of Neurology, University of California, San Francisco, San Francisco, California 94158, and ³Veterans Affairs Northern California Health Care System, Martinez, California 94553

The intraparietal sulcus (IPS) has been implicated in numerous functions that range from representation of visual stimuli to action planning, but its role in abstract decision-making has been unclear, in part because low-level functions often act as confounds. Here, we address this problem using a task that dissociates abstract decision-making from sensory salience, attentional control, motor planning, and motor output. Functional MRI data were collected from healthy female and male human subjects while they performed a policy abstraction task requiring use of a more abstract (second-order) rule to select between two less abstract (first-order) rules that informed the motor response. By identifying IPS subdivisions with preferential connectivity to prefrontal regions that are differentially responsive to task abstraction, we found that a caudal IPS (cIPS) subregion with strongest connectivity to the pre-premotor cortex was preferentially active for second-order cues, whereas a rostral IPS subregion (rIPS) with strongest connectivity to the dorsal premotor cortex was active during attentional control over first-order cues. These effects for abstraction were seen in addition to cIPS activity that was specific to sensory salience, and rIPS activity that was specific to motor output. Notably, topographic responses to the second-order cue were detected along the caudal-rostral axis of IPS, mirroring the broader organization seen in lateral prefrontal cortex. Together, these data demonstrate that subregions within IPS exhibit activity responsive to policy abstraction, and they suggest that IPS may be organized into frontoparietal subnetworks that support hierarchical cognitive control.

Key words: abstraction; connectivity; decision-making; hierarchy; intraparietal sulcus; prefrontal cortex

Significance Statement

Abstract decision-making allows us to flexibly adapt our behavior to new contexts. Although much previous work has focused on the role of lateral prefrontal cortex in such decisions, the contributions of parietal cortex have been relatively understudied. Here, we demonstrate that spatially segregated subregions of human IPS with strong functional connections to lateral prefrontal cortex demonstrate activity selective for abstract decisions. This activity can be distinguished from responses resulting from cognitive processes related to sensory salience, attentional control, motor planning, and movement. Together, these findings indicate that different task demands are reflected in the topography of IPS, and they explicitly reveal a role in abstract decision-making.

Received June 16, 2021; revised July 14, 2022; accepted July 18, 2022.

Author contributions: M.N. and A.S.K. designed research; M.N. performed research; M.N., S.L.C., M.D., and A.S.K. contributed unpublished reagents/analytic tools; M.N., S.L.C., and A.S.K. analyzed data; M.N. and A.S.K. wrote the paper.

This work was supported by National Eye Institute Grant EY024554 to A.S.K.; University of California, Berkeley Graduate Division Mentored Research Award to M.N.; the Milton I. and Florence Mack Neurology Research Fund through a partnership with the Helen Wills Neuroscience Institute at University of California, Berkeley to M.N.; and the University of California, Berkeley Henry H. Wheeler, Jr. Brain Imaging Center, where neuroimaging was performed, supported by the National Science Foundation through its Major Research Instrumentation Program Award BCS-0821855. We thank Karen Lee and Angie Wong for assistance with data collection and preliminary analysis and Swetha Shankar and Daniella Furman for theoretical and analytical assistance.

The authors declare no competing financial interests.

Correspondence should be addressed to Andrew Kayser at Andrew.Kayser@ucsf.edu. https://doi.org/10.1523/JNEUROSCI.1237-21.2022

Copyright © 2022 the authors

Introduction

The intraparietal sulcus (IPS) is critical for a remarkably diverse array of functions in humans, including but not limited to perception of visual motion (Konen and Kastner, 2008), representation of visual features (Xu, 2009), action planning (Grèzes and Decety, 2001), action execution (Culham and Valyear, 2006), and perceptual decision-making (Kayser et al., 2010a). The engagement of IPS in these varied processes reflects its key role in mediating stimulus-response relationships within the frontoparietal network, a set of brain regions responsible for perceptual and motor control. In support of these findings, lesions to IPS produce impairments in cognitive functions such as spatial attention (Gillebert et al., 2011), visual perception (Murphy et al., 2016), and reaching (Karnath and Perenin, 2005; Medina et al., 2020).

In line with this evidence that IPS plays a key role in linking stimuli with appropriate actions, it has been suggested that parietal cortex may contribute to decision-making through participation in networks distinguished by their level of policy abstraction (Choi et al., 2018). Policy abstraction (Badre, 2008; Botvinick, 2008) describes a hierarchical set of superordinate and subordinate action rules that permit behavior to adapt flexibly to novel environments. For example, based on the context of one's visit to a new restaurant, whether ordering takeout or eating on site (superordinate rule), one can then either identify the cashier counter for payment or the restaurant host for seating, respectively (subordinate rules).

Previous studies of policy abstraction have largely focused on the frontal lobe (Badre and D'Esposito, 2007; Badre et al., 2010; but see D'Mello et al., 2020). This literature has identified specific regions along the rostral/caudal axis of the lateral frontal cortex that respond hierarchically to increasingly abstract decisions (Nee and D'Esposito, 2016), including the dorsal premotor cortex (PMd), pre-premotor cortex (pre-PMd), inferior frontal sulcus, and rostral PFC (Badre et al., 2010; Kayser and D'Esposito, 2013). Yet, like lateral frontal cortex, regions within the IPS are polymodal areas (Mesulam, 1998) with complex cognitive functions (Xu, 2018). Because lateral frontal cortex and the IPS are not only highly connected to each other but also project to and receive projections from many of the same brain regions, understanding their respective contributions to complex cognition is critical to understanding circuit function (Goldman-Rakic, 1988). However, in part because of the multiple, potentially confounding perceptual and motor processes instantiated within IPS, it is not currently known whether IPS also contains localized higher-order representations responsive to policy abstraction.

To address this possibility, here, we capitalize on the fact that the brain is functionally segregated into large-scale intrinsic networks that are highly reproducible across subjects (Moussa et al., 2012) and parcellate the brain into discrete regions based on shared activity dynamics, even in the resting state (Schaefer et al., 2018). Moreover, within IPS, it has been demonstrated that multiple functional subregions constitute a complex, topographic organization (Silver and Kastner, 2009). If subregions in IPS participate in higher-order abstract decision-making, they are likely to be differentially connected to lateral frontal regions sensitive to distinct levels of policy abstraction. In addition, activity related to abstraction within IPS should be differentiable from other known functions of IPS. To this end, we created a task that systematically dissociates known sensory and motor control functions from goal-oriented, abstract functions. Specifically, we use an adaptation of the random-dot kinetogram task (Britten et al., 1992; Kayser et al., 2010a) to manipulate the salience of sensory stimuli, attentional control, the timing of motor responses, and the level of abstraction.

We hypothesized that distinct subregions within IPS would exhibit hierarchically ordered responses to policy abstraction that are consistent with intrinsically connected lateral frontal regions. Moreover, these responses should not be readily explained by other known functions of IPS. We find support for both hypotheses, and our results suggest that hierarchical cognitive control networks that include IPS may reflect an organizing principle that informs the multifaceted functions of parietal cortex.

Materials and Methods

Thirteen women and 10 men provided written informed consent to participate, in accordance with the Committee for the Protection of Human Subjects at the University of California, Berkeley. Eight women and three men (11 subjects, age 18-45 years) completed all study procedures and were included in analyses. Of the 12 excluded subjects, seven did not complete prescan training, one developed an MRI contraindication after consent, one showed a persistent response bias (i.e., indicating blue on >50% of color training runs), two demonstrated excessive head motion (>3 mm) during MRI scanning, and one subject fell asleep during scanning. Subjects had normal neural anatomy as assessed by a neurologist (A.K.), were right-handed, and had normal or corrected-to-normal vision. Before the first scan session, subjects were trained for a minimum of five 1.5 h training sessions to minimize learning effects during MRI scanning. Subjects then performed the task inside the fMRI scanner for five 1.5 h scan sessions, each consisting of six runs of 22 trials, for a total of 660 trials across the 5 d of testing. Each session occurred between 1 and 7 d after the previous session. The average number of days between the first and last scan session across all subjects was 13.4, with a median of 10 d.

Experimental design and statistical analyses

Task design. Subjects performed an abstract decision-making task (Fig. 1A) in which each trial consisted of three sequentially presented visual stimulus displays (color, motion, and shape stimuli, respectively) shown in randomized, counterbalanced order (Fig. 1B), followed by an explicit response period. The sensory salience of each visual stimulus was varied by manipulating its coherence level (Fig. 1C). To address the role of policy abstraction (Botvinick, 2008; Badre and D'Esposito, 2009), the task was structured so that the coherence discrimination for one of the three stimuli determined the relevance of the other two stimuli, which in turn were associated with a pair of stimulus-response mappings (Fig. 1D). Specifically, one of the three stimuli in each trial served as the more abstract contextual cue (second-order rule) necessary for selecting which of the other two stimuli determined the motor response (firstorder rule). The third remaining stimulus in each trial was therefore unnecessary for the response and could be ignored (irrelevant cue). To dissociate specific visual aspects of the sensory stimulus from its role in the task, the second-order, first-order, and irrelevant cues were equally represented by the color, motion, and shape stimuli across all runs. Moreover, as noted above, the position of each of the three cues (second order, first order, and irrelevant) was randomized and counterbalanced across trials to avoid confounds that might potentially be associated with specific positions. For example, the uncertainty of the motor response and novelty responses related to the identity of the specific stimulus (color, motion, or shape) would be maximal for the first position. Thus, all runs were pseudorandomized to ensure that position, coherence, and cue type were equally distributed across trials (Fig. 1D), under the assumption that doing so reduces the likelihood that features particular to any one trial sequence could drive behavioral and neural effects averaged across all sequences.

The same second-order rule was active for an entire run but varied across runs. For example, a full 8.5 min color run would always use the color stimulus as the second-order cue, but the next run might use the shape stimulus as the second-order cue. Moreover, specific second-order rule mappings, that is, from specific second-order feature to first-order rule (Fig. 1*C*), were consistent within a subject but counterbalanced across subjects. For example, one subject's second-order rule for a color run might map blue to the shape stimulus, whereas a different subject's color rule mappings from stimulus feature to response were consistent across all trials and runs within a subject but counterbalanced across subjects. For example, the first-order rule mappings to the right button for one subject to the left button and triangles to the right button for one subject but the opposite for another (Fig. 1*D*).

At the beginning of each run, an instruction screen indicated which of the second-order rules was active for the run and therefore which of the stimuli was serving as the second-order cue (Fig. 1A). On each trial, subjects were then required to perform two perceptual discriminations,



Figure 1. *A*, Task design. At the start of each run, an instruction screen notified subjects of the active second-order rule. Each trial then consisted of three sequentially presented stimuli accompanied by first-order rule reminders that were displayed as symbols on the left and right of the screen (to indicate mappings to the left and right buttons, respectively). Following the presentation of the third stimulus, subjects made a button press response, after which a reminder of the active second-order rule was displayed during the intertrial interval. On this example trial, the second-order cue, the color stimulus, indicates that the relevant first-order cue is the shape stimulus, and the irrelevant cue is the motion stimulus. The correct response is a left button press, corresponding to the circle feature. *B*, Stimuli were presented equally often in positions 1, 2, and 3 across all trials. Each position permutation was pseudorandomly presented. Three of the possible six permutations are shown here. *C*, Stimuli were presented at either high coherence (high sensory salience) or low coherence (low sensory salience) equally often across all trials. *D*, Hierarchical structure of the policy abstraction task for the example trial. Across all runs, color, motion, and shape stimuli served equally often as the second-order cue.

a discrimination for the second-order cue and a discrimination for the relevant first-order cue, while the irrelevant cue could be ignored. For example, during a color run (Fig. 1A), the color stimulus might be displayed first in position 1. After judging whether there were more gray dots or blue dots present, the subject would know whether to use the shape rule or the motion rule to determine the button press response. In this case, if the color stimulus in this trial contained more blue dots, the shape stimulus would serve as the first-order cue. Thus, if position 2 contained the motion stimulus, the subject would disregard it. The shape stimulus in position 3 would then be the relevant first-order cue, during which the subject would make the appropriate perceptual discrimination to determine whether there were more circles or more triangles present. Each of the sequentially displayed cues within a trial was separated by an interstimulus interval of 1, 3, or 5 s, pseudorandomly presented from a uniform distribution. To reduce working memory load, small symbols on the left and right sides of the screen outside the circular stimulus display aperture reminded subjects of the first-order rule mapping from stimulus to button press for each of the stimuli (Fig. 1A).

To dissociate motor execution from the decision itself, subjects were required to withhold the appropriate button press until prompted to respond at the end of each trial. Once the response cue appeared, subjects were given 2 s to press the appropriate button as quickly and as accurately as possible. The left button was always pressed using the right index finger, and the right button was always pressed with the right middle finger. During the intertrial interval (3, 5, or 7 s duration, pseudorandomly drawn from a uniform distribution), subjects viewed a reminder screen about the currently active abstract rule to further reduce working memory demand (Fig. 1*A*). At the end of each 8.5 min run, subjects were shown their overall accuracy for the preceding run; no other feedback was provided. The task was programmed in MATLAB using Psychophysics Toolbox version 3 (http://psychtoolbox.org/), adapted from previous code (Kayser et al., 2010b; Shankar and Kayser, 2017).

Behavioral training. In training sessions 1–3, each subject completed a minimum of six behavioral runs per session to learn to perceptually discriminate between competing features of a single stimulus type (i.e., two runs limited to discrimination of blue or gray colored dots, two runs limited to discrimination of 2-D circles or triangles), counterbalanced across sessions. Coherence for the motion stimulus was defined as the percentage of the total number of dots moving coherently (Kim and Shadlen, 1999; Heekeren et al., 2006; Kayser et al., 2010a). For shape and color stimuli, coherence for the target feature was defined as the difference between the number of target features and the number of dots and 40 gray dots were shown as the color stimulus, the color coherence would equal (60-40)/(60 + 40) = 20%, and the correct response would be blue. Lower coherence produced lower sensory salience and therefore a perceptually more ambiguous display for all three stimulus types.

Perceptual discriminations during training were preset to a range of seven coherences, that is, 0, 2, 4, 8, 16, 32, and 64% for the motion stimulus (Kayser et al., 2010b); 0, 4, 7, 11, 18, 36, and 68% for the color stimulus; and 0, 7, 13% 20, 33, 67, and 87% for shape. The training coherence values captured the full range of behavioral performance from chance (50%) performance to 100% accuracy, as fit by a psychometric curve derived from the logistic function $[f(x) = \frac{a}{1+e^{-bx}}]$, where *a* defines the maximal accuracy at a given coherence level, and *b* defines the slope of the curve.

In training sessions 1 and 2, subjects were provided with auditory feedback at the end of each trial. A correct trial was communicated with a low-frequency tone of 0.15 s duration, whereas an incorrect trial was communicated with a corresponding high-frequency tone. When subjects made the transition to training on the full task in session 3, trial-by-trial feedback was discontinued. At the end of training session 3, individual high and low coherence values (one high and one low value) were selected for each subject based on their interpolated behavioral performance at 95% and 75% accuracy, respectively, to ensure comparable performance for every subject. These

individualized high and low coherence values were used in training sessions 4-5, in which subjects performed the full task.

To acclimate subjects to assessing the second-order cue, the first half of training session four used a version of the full task in which the second-order cue was always presented in position 1. In the second half of training session 4, the second-order cue appeared equally as likely in position 1, 2, or 3, consistent with the task that subjects performed in the MRI scanner. Auditory feedback on each trial was again provided for full task training during session 4 before being discontinued for session 5. Training session 5 used the identical task completed during fMRI scanning sessions. All subjects reached stable performance of at least 75% accuracy before continuing with scanning sessions.

MRI parameters. MRI scanning took place at the Henry H. Wheeler Jr. Brain Imaging Center at the University of California, Berkeley using a Siemens TIM/Trio 3 T MRI and a 12-channel receive-only head coil. Functional images were acquired with a single-shot gradient echoplanar imaging protocol in contiguous descending order (32 slices, TR = 1.8 s, TE = 23 ms, FOV 225 mm, matrix size $70 \times 70 \times 32$, voxel size $3.2 \times$ 3.2×3 mm). Anatomical images were acquired with a T1-weighted MP-RAGE imaging protocol (160 slices, TR = 2.3 s, TE = 2.98 ms, FOV 256 mm, matrix size 256 \times 256, voxel size 1 \times 1 \times 1 mm). Head movement was restricted using foam padding. A projector (Avotec SV-6011) displayed the task on a screen inside the scanner bore. Once a mirror was placed over the head coil, the distance from the subject's eye to the screen was 29 cm, and the presented images subtended a visual angle of 7.5°. An MRI-compatible fiber-optic four-button response device (inline model HH-1 \times 4-L, Cambridge Research Systems) was used for subject responses.

fMRI preprocessing. fMRI data were preprocessed using a version of the open-source Analysis of Functional NeuroImages (AFNI) pipeline written by S.L.C. (https://github.com/savannahcookson/AFNI-Pipeline), which was customized to manage multisession data. This pipeline supports preprocessing and first-level analysis of data based on the afni_proc.py function available in the AFNI software package (Cox, 1996). DICOM (Digital Imaging and Communications in Medicine) images were converted to NIFTI format using MRIcron DCM2NII (https:// people.cas.sc.edu/rorden/mricron/dcm2nii.html); they were then aligned to a T1 anatomic image acquired on the same day as the functional images. Data were despiked, corrected for slice timing, registered to the second functional image of the first run acquired that day, resampled to $3 \times 3 \times 3$ mm space, and then aligned to the skull-stripped anatomic image warped to the standard Montreal Neuroscience Institute (MNI) atlas (ICBM152; https://www.mcgill.ca/bic/). Data were smoothed with a 6 mm FWHM Gaussian kernel and scaled to have a voxelwise mean of 100 by dividing the time series of each voxel by its mean signal and multiplying by 100. Individual volumes that contained >10% outlier voxels were censored from further analysis. Data for each subject were then concatenated across all scanning days for first-level analysis.

Univariate analysis

All univariate analyses were completed with AFNI software through a combination of the above preprocessing pipeline and additional custom scripts. The AFNI function 3dDeconvolve was used to create a general linear model that evaluated nine different conditions as separate regressors. Six regressors evaluated correct trials only during the following stimulus presentation conditions: (1) the second-order cue at high coherence, (2) the second-order cue at low coherence, (3) the first-order cue at high coherence, (4) the first-order cue at low coherence, (5) the irrelevant cue at high coherence, and (6) the irrelevant cue at low coherence. Three additional regressors captured trials and/or events of lesser interest: (7) all three stimulus displays on error trials, (8) the instruction screen displaying the active second-order rule, and (9) the button press response. Because regressors 1-6 were limited to correct trials only, the number of trials contributing to these regressors was less than the 660 total trials each subject completed. For each subject, 542 \pm 19 correct trials (mean \pm SD), on average, were included within these six regressors.

Each regressor was created from the convolution of a gamma probability density function that peaked at 6 s with a condition-specific vector of stimulus onsets and durations (2 s for stimulus presentation, 6 s for

the instruction screen, and a stick function for the button press response). Resultant estimated β coefficients for each subject were mapped to the MNI template for use in group-level analyses. Wholebrain voxelwise contrasts were next created for the second-order cue versus first-order cue (Second-order vs First-order), first-order cue versus irrelevant cue (First-order vs Irrelevant), all high coherence cues versus all low coherence cues (High vs Low), and the time point of button press versus time points during task stimulus presentation (Button press vs Visual stimuli). Stimulus type (color, motion, or shape) for each cue was intentionally collapsed to capture parietal activity that is not specific to visual features. To remove areas that deactivated during task performance, each contrast was masked by the positive main effect of task, created by contrasting all conditions versus baseline at a loose threshold (p < 0.05, uncorrected).

Regions of interest selection. All regions of interest (ROIs) were selected from a 1000-area parcellation (Schaefer et al., 2018) derived from a well-established seven-network resting-state fMRI (rs-fMRI) analysis (Yeo et al., 2011), which was obtained at 1 mm resolution and resampled to $3 \times 3 \times 3$ mm task space (https:// github.com/ThomasYeoLab/CBIG/tree/master/stable_projects/brain_ parcellation/Schaefer2018_LocalGlobal). The Schaefer et al. (2018) parcel that included MNI coordinates (-30, -10, 68) in dorsal premotor cortex identified by Badre and D'Esposito (2007) was selected as the PMd ROI [parcel #144, MNI centroid (-28, -12, 60)]. Likewise, the parcel that included MNI coordinates [-38, 10, 34] in pre-premotor cortex from Badre and D'Esposito (2007) was selected as the pre-PMd ROI [parcel #355, centroid (-40, 10, 32)].

Resting-state functional connectivity analysis. Independent rs-fMRI data from 100 human subjects studied in the Human Connectome Project (HCP; healthy young adult Fix_extended rs-fMRI) was then used to identify parietal parcels with maximal functional connectivity to pre-PMd and PMd parcels, respectively. These data from 100 HCP subjects were selected for minimal motion and were further processed using the standard ICA-FIX pipeline (https://www.humanconnectome.org/ software/hcp-mr-pipelines), followed by regression of the global average brain signal and a bandpass filter (0.009-0.08 Hz) to remove physiological artifacts. Parcel-level time series averages were obtained by fitting HCP resting-state data to the Schaefer et al. (2018) 1000-area parcellation using MATLAB and AFNI. Pearson's correlations between all parcels were z-scored using Fisher's R-to-Z transform via the AFNI 3dNetCorr command.

The resultant z-scored correlation matrices were averaged across all 100 HCP subjects and used to identify the parietal parcel with the highest positive correlation to the pre-PMd parcel, where parietal cortex was defined using cytoarchitectonic macro labels from the standard AFNI atlas CA_ML_18_MNI. The same method was repeated for the PMd parcel. This analysis yielded a caudal subregion of the left IPS with greatest connectivity to pre-PMd [MNI centroid coordinates (-32, -70, 50), parcel #321], hereafter referred to as caudal IPS (cIPS), and a more rostral IPS subregion with greatest connectivity to PMd [MNI centroid coordinates (-38, -42, 56), parcel #211] hereafter referred to as rostral IPS (rIPS).

To determine whether these parietal ROIs had differential connectivity to pre-PMd compared with PMd, vectors for each HCP subject's z-scored correlations of cIPS/rIPS with pre-PMd and cIPS/rIPS with PMd were first standardized. For each parietal ROI, we then implemented a paired t test comparing the correlations to pre-PMd and PMd using a Bonferroni-corrected significance threshold of p = 0.025. Effect size for this test and all paired two-sample *t* tests in this study was determined using the standard equation for Cohen's d, $(d = \frac{d'}{\sqrt{1-r}})$, where $d' = \frac{\overline{x_1 - \overline{x_2}}}{s}$, *r* is the correla-

tion between samples, and the pooled SD *s* is defined as follows:

$$s = \sqrt{\frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2}}.$$

We next evaluated whether the choice of regions whose connectivity differentiated pre-PMd and PMd was critical to our results by further identifying parietal regions maximally connected to both pre-PMd and

PMd. This approach allowed us to compare task activity in parietal regions that have selective connectivity to one particular region in the frontal hierarchy with activity in parietal regions that have more generalized connectivity to multiple frontal regions. To account for overall differences in connection strength for the two regions, we separately standardized the vectors of z-scored correlations between pre-PMd/PMd and all other brain parcellations by subtracting the mean and dividing by the SD for each. We next took the item-by-item product of those two correlation vectors. The products were ranked, and the five parietal parcels with the highest positive products were selected [IPSa with MNI centroid coordinates (-42, -40, 42), IPSb (-52, -34, 44), IPSc (-32, -48, 46), IPSd (-48,-42, 50), and IPSe (-38, -52, 58); for further ROI analysis, see Extended data Table 4-1]. We evaluated whether these IPS subregions with more nonspecific connectivity had frontal connectivity profiles distinct from the selectively connected rIPS and cIPS using a one-way ANOVA and a post hoc Tukey HSD multiple comparisons test. A vector of the correlation of each parietal ROI with pre-PMd was subtracted from the correlation with PMd for all 100 HCP subjects, and these absolute differences were treated as seven groups for the statistical tests implemented in MATLAB.

For multiple reasons, we focused on areas within the left hemisphere. As is true for many cognitive neuroscience studies, only right-handed individuals were included, and our task explicitly involves hierarchical processing of action rules that result in a motor output from the dominant (right) hand. As a consequence, responses are likely to rely on frontoparietal action representations in the left hemisphere. In addition, because subjects might use language rules to facilitate hierarchical responses, there may be differences between hemispheres linked to language-based strategies, or even to subvocal rehearsal. Consistent with these ideas, previous reports investigating similar abstract tasks have found predominantly left-sided activations (Badre and D'Esposito, 2007; Nee and Brown, 2012, 2013; Jeon and Friederici, 2013; Jeon et al., 2014; Nee and D'Esposito, 2016).

fMRI-based ROI analysis. After selecting the parcel used for each ROI, an ROI mask was created by resampling from the Schaefer et al. (2018) 1000-area map. The mask was then applied to each subject's univariate whole-brain contrast dataset (Second-order vs First-order, First-order vs Irrelevant, High vs Low, and Button press vs Visual stimuli). We then took the mean β value from all voxels within the selected parcel for each subject and performed ANOVAs and *t* tests at the group level using custom MATLAB scripts. All ANOVA tests assessed effect size with partial η -squared, $\eta^2 p = \frac{SS_{effect}}{SS_{effect} + SS_{error}}$. Effect sizes for all one-sample β contrast ROI results were found using the standard Cohen's d equation, $d = \frac{\overline{x_1} - \mu}{\sigma}$. Graded activity for second-order abstraction in IPS. To investigate

Graded activity for second-order abstraction in IPS. To investigate whether selectivity for second-order abstraction varied systematically across the multiple subregions of IPS defined above, we conducted an exploratory analysis. To do so, we initially computed the Pearson correlation coefficient between the mean β value obtained from the Second-order versus First-order contrast and the *y* coordinate (MNI) of the centroid for the seven IPS ROIs previously identified by connectivity methods, IPSa, IPSb, IPSc, IPSd, IPSe, rIPS, and cIPS. We repeated this analysis for the First-order versus Irrelevant contrast, then calculated the significance of the difference between the Second-order versus First-order correlation and the First-order versus Irrelevant correlation.

Conjunction maps. To determine whether our ROI-based results were specific to the connectivity-defined regions evaluated above, we also performed an exploratory conjunction analysis. As for the ROI analyses, whole-brain maps of interest (see below) were first loosely masked by the positive main effect of task (p < 0.05, uncorrected) to remove any voxels that explicitly deactivated in response to task performance. An anatomic mask of parietal cortex, again defined by the cytoarchitectonic macro labels from the standard AFNI atlas, CA_ML_18_MNI, was then applied to limit the search to parietal voxels.

To assess the same conjunction of properties seen in the ROI analyses, we evaluated contrasts that identified voxels responding differentially to



Figure 2. Group-level behavioral effect of coherence on trial accuracy, collapsed across cue types, for the second-order, first-order, and irrelevant cues. Error bars indicate between-subject SEM. Asterisks indicate p < 0.05; n.s. indicates non-significant.

Second-order versus First-order cues, Relevant versus Irrelevant stimuli, High versus Low coherence, and/or Button press versus Other task stimuli, as indicated. Pertinent maps were defined by loose thresholds (p < 0.05 uncorrected), after which conjunction analyses were performed across properties matched to the ROI-based findings. For example, to identify cIPS-like regions, we evaluated BOLD activity that, like cIPS, differentiates second-order from first-order cues and high from low coherence but does not distinguish relevant from irrelevant stimuli or the button press from other task stimuli (see below, Results). Similarly, we evaluated BOLD activity that, like rIPS, differentiates relevant from irrelevant stimuli and button press from other task stimuli but does not distinguish second-order from first-order cues or high from low coherence. Lastly, the resulting maps were overlapped with the resting-state connectivity-defined parcels (i.e., cIPS and rIPS) to determine conjoint responses.

Sample size. This study was designed to maximize within-subject power so that multiple cognitive processes could be differentiated in within-individual contrasts rather than between-groups comparisons. More broadly, this approach borrows from more recently developed techniques that acquire large numbers of scans in a small number of subjects (Gordon et al., 2017). Based on our estimates, recruitment of at least 10 subjects provides 80% power to identify large effect sizes (Cohen's d = 1.0) at an α value of 0.05. Including subject attrition noted above, the actual number of subjects that ultimately survived all quality checks was 11.

Results

Behavior

Subjects completed a task in which factors including the order (level) of abstraction, type of visual stimulus, sensory salience, and cue position were manipulated (Fig. 1). As expected, subject performance improved with increased sensory salience; accuracy was higher for pertinent high coherence compared with low coherence cues except when the cue was irrelevant (second-order cue at high coherence vs second-order cue at low coherence, $t_{(10)} = 7.95$, $p = 1.24 \times 10^{-5}$, d = 3.45; first-order cue at high coherence vs first-order cue at low coherence, $t_{(10)} = 13.69$, $p = 8.36 \times 10^{-8}$, d = 5.90; irrelevant cue at high coherence vs irrelevant cue at low coherence, $t_{(10)} = 0.23$, p = 0.82, n.s.; d = 0.10; Fig. 2). This result confirmed that subjects attended to appropriate task-relevant stimuli and used coherence in their decisions. To ensure that subjects evaluated the abstract relationship of the second-order cue rather than merely the visual stimulus feature that composed it, we compared accuracy for trials in which motion, color, and shape stimuli, respectively, served as the second-order cue. No differences were seen $(F_{(2,30)} = 0.16)$, p = 0.85, n.s., $\eta^2 p = 0.01$), suggesting that abstract decisions were not dependent on the low-level visual feature. As all cue types were presented equally often in positions 1, 2, and



Figure 3. Univariate whole-brain analysis depicting the positive main effect of task relative to baseline across the group, FDR corrected, q < 0.05. ROIs are overlaid from the Schaefer et al. (2018) parcellation map. Dashed white lines indicate preferential connectivity defined by analysis of HCP data.

3 across all trials to minimize potential expectation effects, we next confirmed that the position of the second-order cue did not affect accuracy ($F_{(2,30)} = 1.53$, p = 0.23, n.s., $\eta^2 p = 0.09$). Furthermore, accuracy did not differ significantly among trials for the six corresponding permutations of position for the second-order, first-order, and irrelevant cues ($F_{(5,60)} = 0.81$, p = 0.55, n.s., $\eta^2 p = 0.06$). Consistent with data for accuracy, response time did not differ across visual stimulus type or position of the second-order cue, and there were no differences in response time between high and low coherence cues, suggesting that decisions were made in advance of the response cue.

Univariate functional imaging maps

A whole-brain, voxelwise contrast of all task conditions versus baseline yielded a characteristic frontoparietal network consistent with prior studies of hierarchical cognitive control in policy abstraction tasks (Fig. 3). To more directly analyze target regions implicated in policy abstraction, we next directed our attention to lateral prefrontal regions previously implicated in abstract decision-making, including regions sensitive to second-order (pre-PMd) and first-order (PMd) rules, to identify their functionally connected regions in parietal cortex.

Connectivity

Using resting-state connectivity derived from Human Connectome Project data, we identified parietal regions that were maximally functionally connected to lateral frontal regions implicated in second-order (pre-PMd) and first-order (PMd) abstract decisions. Using the two lateral frontal parcels containing pre-PMd and PMd, respectively, as seed regions (Schaefer et al., 2018; see above, Materials and Methods), we identified the parietal parcels with the highest positive raw *z*-scored correlation for pre-PMd (parcel 321 in cIPS) and PMd (parcel 211 in rIPS; Fig. 3). These two IPS subregions showed selective connectivity to their corresponding regions in frontal cortex, with strongly differential connectivity to either pre-PMd or PMd (rIPS, $t_{(99)} = -20.57$, $p = 1.61 \times 10^{-37}$, d = -2.06; cIPS, $t_{(99)} = 14.80$, $p = 7.79 \times 10^{-27}$, d = 1.48). These results support hypotheses that regions within parietal and frontal cortex may participate in distinct functional networks for different orders of abstract tasks (Choi et al., 2018).

Functional imaging ROI analysis

Using these four ROIs—pre-PMd, PMd, cIPS and rIPS—we performed a three-way ANOVA that included factors of ROI, order of abstraction (second-order or first-order), and stimulus coherence (high or low), with subjects treated as a random effect. There was a strongly significant interaction between ROI and abstraction $(F_{(1.78,17.81)} = 32.27, p = 1.99 \times 10^{-6}, \eta^2 p = 0.76)$ and a significant interaction between ROI and coherence ($F_{(1.48,14.76)} = 7.21, p =$ 0.01, $\eta^2 p = 0.42$), though no three-way interaction between ROI, abstraction, and coherence ($\eta^2 p = 0.04$).

The interaction between ROI and abstraction was driven by differences in the responses of pre-PMd and cIPS to secondorder stimuli. Pre-PMd had significantly greater activity for the second-order versus the first-order cue (Second-order vs Firstorder, $t_{(10)} = 5.45$, $p = 2.81 \times 10^{-4}$, d = 1.64), as did cIPS ($t_{(10)} = 6.29$, $p = 9.07 \times 10^{-5}$, d = 1.90), but neither PMd ($t_{(10)} = -1.00$, p = 0.339, d = -0.30 nor rIPS ($t_{(10)} = 0.71, p = 0.493, d = 0.21$) showed such changes (Fig. 4A, top). In addition, these differences could not be explained by general attentional mechanisms. When responses to relevant (i.e., putatively attended) first-order cues and irrelevant (putatively unattended) cues were compared in these regions, we observed a dissociation with respect to the results for order (Fig. 4A, bottom). Specifically, PMd and rIPS strongly differentiated relevant from irrelevant first-order cues (First-order vs Irrelevant, PMd, $t_{(10)} = 3.26$, p = 0.009, d = 0.98; rIPS, $t_{(10)} = 5.77$, $p = 1.80 \times 10^{-4}$, d = 1.74), but pre-PMd ($t_{(10)} =$ -0.15, p = 0.880, n.s., d = -0.05) and cIPS ($t_{(10)} = -0.48$, p =0.641, n.s., d = -0.14) did not. These results suggest that cIPS and rIPS might serve complementary roles in policy abstraction. Moreover, along with previous reports that the bottom of the hierarchy does not respond differentially to secondorder policy abstraction (Badre and D'Esposito, 2007; Badre and Nee, 2018), these data demonstrate that PMd and rIPS are dissociably responsive to attentional manipulations for firstorder stimuli.

In the above comparisons, functionally connected frontal and parietal regions responded equivalently to abstraction-related task features. Conversely, when we examined the significant interaction between coherence and ROI, we found that the effect of coherence was more strongly represented in parietal than in frontal regions, specifically in cIPS (Fig. 4B, top). When high and low coherence stimuli were compared, collapsed across order of abstraction, cIPS responded more strongly to high coherence (High vs Low, $t_{(10)} = 2.93$, p = 0.015, d = 0.88). Notably, as indicated by the absence of a three-way interaction between ROI, coherence, and order, cIPS did not differentiate between coherence for second-order and first-order cue stimuli (cIPS interaction of coherence and order, $F_{(1,10)} = 1.08$, p = 0.324, n.s., $\eta^2 p = 0.10$). To determine whether cIPS might therefore have a specific role in sensorimotor transformations (Erickson and Kayser, 2013), we evaluated whether any of the four frontal and parietal regions were differentially active when movement plans were implemented (one-way ANOVA Button press vs Visual stimuli, $F_{(3,40)} = 4.54, p = 0.008, \eta^2 p = 0.25$). cIPS did not distinguish between the motor output and visual stimuli ($t_{(10)} = -1.26$, p =0.237, n.s., d = -0.38), but rIPS was significantly more active for the button press ($t_{(10)} = 3.18$, p = 0.009, d = 0.96). Neither of the two prefrontal ROIs showed significantly different activity for this contrast (pre-PMd, $t_{(10)} = -1.96$, p = 0.079, n.s., d = -0.59; PMd $t_{(10)} = 0.64$, p = 0.536, n.s., d = 0.19). Together, these results demonstrate that IPS subregions show selectivity for abstraction consistent with their respective functionally connected frontal regions. In addition, these IPS regions demonstrate sensitivity to task factors not seen in the pre-premotor and dorsal premotor areas previously shown to be sensitive to second- and first-order rules, respectively. Specifically, cIPS processed information about



Figure 4. R0I-based results. *A*, Within-R0I activity for the univariate contrast of all second-order cues versus first-order cues (top) and all first-order cues versus irrelevant cues (bottom). *B*, Activity for the same R0Is in the contrast of all high coherence versus low coherence cues (top) and the time point of button press versus task stimulus displays (bottom). The β parameter is equivalent to percentage signal change (see above, Materials and Methods). Error bars indicate between-subject SEM. Asterisks indicate p < 0.05; n.s. indicates non-significant (Extended Data Table 4-1).

sensory salience, whereas first-order rIPS was differentially active during motor output.

To address the possibility that other functions of parietal cortex, such as motor planning, might explain the findings in cIPS, we next evaluated distinct orderings of the task stimuli on individual trials, specifically, the position of the second-order cue relative to the relevant first-order cue. If cIPS is responsible for motor planning, its BOLD activity in response to the first-order cue should differ (1) when the first-order cue is presented after the second-order cue, in which case motor planning is possible, compared with (2) its presentation before the second-order cue, in which case no motor planning is possible. As expected, responses in the first-order regions, PMd and rIPS, to the relevant first-order cue were significantly different ($F_{(3,40)} = 3.38, p = 0.027$, $\eta^2 p = 0.20$) from responses to the relevant first-order cue in the second-order regions, pre-PMd and cIPS; and activity in PMd and rIPS was significantly larger when the relevant first-order cue was presented after the second-order cue (PMd, $t_{(10)} = 3.86$, p = 0.003, d = 1.16; rIPS, $t_{(10)} = 2.64$, p = 0.025, d = 0.80). However, activity in pre-PMd and cIPS was unchanged (for both ROIs, p > 0.50, n.s.). In keeping with this last finding, when BOLD activity in response to the second-order cue was compared between (1) when the second-order cue occurred after the first-order cue, in which case motor planning is possible, and (2) when the secondorder cue occurred before the first-order cue, in which case no motor planning is possible, there was no difference seen between conditions ($F_{(3,40)} = 0.675$, p = 0.573, n.s., $\eta^2 p = 0.05$, and none of the individual areas showed responses that differed from zero (all p values > 0.36). Together, these results indicate that motor planning demands may be represented in rIPS but not in cIPS.

By the same token, the cue positions within a trial might affect results related to relevant versus irrelevant information. For first-order cues in position 1, subjects do not yet know which cue is relevant and which is irrelevant as they have not yet seen the second-order cue. Thus, in this position, there is as yet no relevant/irrelevant distinction. Consistent with this prediction, there were no significant differences in the BOLD response between the (ultimately) relevant and irrelevant cues for any of the ROIs, including PMd and rIPS, for position 1 (all p values >0.21). However, in position 3, all the information for the trial is available, and the differences between relevant and irrelevant first-order cues should be maximal. As expected, PMd ($t_{(10)}$ = 3.31, p = 0.008, d = 1.0) and rIPS ($t_{(10)} = 4.08$, p = 0.002, d =1.23), but not pre-PMd ($t_{(10)} = 0.24$, p = 0.812, n.s., d = 0.07) and cIPS ($t_{(10)} = 0.02$, p = 0.982, d = 0.01), strongly and significantly distinguished relevant from irrelevant cues. As demonstrated in Figure 4, these differences are strong enough to remain significant when collapsed across positions.

An additional potential concern is that the results in Figure 4 hold only for a specific level of coherence. To address this possibility, we analyzed the response within cIPS and rIPS separately for high and low coherence stimuli. For the contrasts of second-order cue versus the relevant first-order cue for high coherence conditions only and for low coherence conditions only, results were similar. Specifically, both pre-PMd and cIPS responded significantly differently than did PMd and rIPS for both high ($F_{(3,40)} = 21.2$, $p = 2.19 \times 10^{-8}$, $\eta^2 p = 0.61$) and low coherence stimuli ($F_{(3,40)} = 14.7$, $p = 1.3 \times 10^{-6}$, $\eta^2 p = 0.52$), and these two second-order regions showed consistently greater activity for the second-order stimulus in both cases (cIPS high, $t_{(10)} = 7.0$, $p = 3.6 \times 10^{-5}$, d = 2.12; cIPS low, $t_{(10)} = 4.8$, $p = 7.3 \times 10^{-4}$, d = 1.44; pre-PMd high, $t_{(10)} = 5.7$, $p = 2.0 \times 10^{-4}$, d = 1.72; pre-PMd low, $t_{(10)} = 4.7$, $p = 8.0 \times 10^{-4}$, d = 1.43). In contrast, PMd and rIPS did not show differential responses to these cues (all p values > 0.087).



Figure 5. Variation in responsiveness to abstraction along the rostral-caudal axis of IPS. The Pearson's correlation for the β parameter estimate for each contrast (see below) is plotted versus the MNI *y* coordinate of seven IPS subregions defined by independent connectivity measures (see above, Materials and Methods). Information related to the level of abstraction (second-order cue vs first-order cue contrast) is indicated by the circular symbols and black solid line, whereas information related to the relevance of the first-order cue (first-order cue vs irrelevant cue contrast) is denoted by the square symbols and black dashed line. Left, A significant difference between the two correlations is represented by the solid gray bracket. The color of each data point corresponds to the color-coded ROI shown on the surface at right (lime green, clPS; pink, IPSe; purple, IPSc; orange, IPSd; bright blue, rIPS; red, IPSa; brown, IPSb). Asterisks indicate *p* < 0.05, n.s. indicates non-significant. Extended Data Table 5-1 shows the within-subject correlations of abstraction versus rostral-caudal position.

For the contrast of relevant versus irrelevant cues, the results for high coherence demonstrated significant differences between firstorder and second-order regions ($F_{(3,40)} = 3.61$, p = 0.02, $\eta^2 p =$ 0.21). However, although both PMd and rIPS moved in the expected direction numerically, the contrast was not significantly different from zero in either of these regions [rIPS, $t_{(10)} = 2.03$, p =0.07 (trend), d = 0.61; PMd, $t_{(10)} = 0.96$, p = 0.36, n.s., d = 0.29]. For low coherence, although high variance in the pre-PMd and cIPS responses rendered the differences between regions nonsignificant in a one-way ANOVA ($F_{(3,40)} = 0.44$, p = 0.73, n.s., $\eta^2 p = 0.03$), responses within PMd and rIPS, but not in pre-PMd and cIPS, remained significantly different from zero after Bonferroni correction (rIPS, $t_{(10)} = 3.5$, p = 0.0056, d = 1.06; PMd, $t_{(10)} = 3.56$, p =0.0052, d = 1.07; pre-PMd and cIPS, both p values > 0.49, n.s.). Thus, in keeping with the lack of a three-way interaction between ROI, abstraction, and coherence, these data suggest that coherence does not directly modulate the difference between second-order and first-order responses in these areas.

IPS subregions with connectivity to both pre-PMd and PMd

Given that cIPS and rIPS were selected because of their strongly differentiated connectivity with pre-PMd/PMd, and they subsequently exhibited complementary differences in task activity, we predicted that parietal regions with less specific connectivity to both pre-PMd and PMd would demonstrate less differentiated task responses. Thus, we next searched parietal cortex for regions with strong connectivity to both pre-PMd and PMd in HCP rsfMRI data (see above, Materials and Methods). We identified the following five parcels: parcel 197 [IPSa, MNI centroid (-42, -40, 42)], parcel 196 [IPSb (-52, -34, 44)], parcel 204 [IPSc (-32,-48, 46)], parcel 326 [IPSd (-48, -42, 50)], and parcel 212 [IPSe (-38, -52, 58); Extended Data Table 4-1] All parietal parcels with nonspecific connectivity were located in IPS and situated along the rostral-caudal axis of the sulcus. Frontal connectivity discrepancies in IPSa-e were different from cIPS and rIPS but not significantly different from each other $(F_{(6,693)} = 16.63,$ $p = 5.712 \times 10^{-18}$, $\eta^2 p = 0.13$; see above, Materials and Methods), allowing us to draw comparisons between subregions with less

specific connectivity to pre-PMd/PMd and those with more selective connectivity.

Confirming our predictions, these IPSa-e subregions showed task activations that were both largely similar to each other and also dissimilar to those of rIPS and cIPS (Extended Data Table 4-1). A three-way ANOVA with factors of ROI, coherence, and order produced an interaction of coherence and ROI ($F_{(6,60)}$ = 6.92, $p = 1.28 \times 10^{-5}$, $\eta^2 p = 0.41$) that was driven entirely by coherence effects in cIPS. In post hoc analyses, the only subregion with a significant difference in activity for sensory salience was cIPS, and the only subregion with a significant effect of motor output was rIPS (see nonsignificant results for IPSa-e in Extended Data Table 4-1). With respect to motor planning and attentional control over first-order cues, only rIPS and IPSc (IPSc, $t_{(10)} = 2.42$, p = 0.036, d = 0.73) had significantly greater activity for first-order relevant versus irrelevant cues. However, post hoc analyses of the interaction of order and ROI ($F_{(2.54,25,36)} = 9.34$, p = 0.004, $\eta^2 p =$ 0.48) revealed that all five of the nonspecific connectivity subregions had significantly greater activity for the second-order cue compared with the first-order cue (Extended Data Table 4-1; see also below). To further investigate this latter result, we evaluated whether this responsiveness to second-order cues demonstrated systematic variation.

Graded activity for second-order decisions along the rostrocaudal axis of IPS

After discovering that the five previously identified IPS subregions along the rostrocaudal axis (see above, Materials and Methods) had varying sensitivity to the second-order cue, we conducted an exploratory analysis to determine whether a relationship existed between spatial location and higher-order abstraction. We hypothesized that the response to secondorder versus first-order stimuli might vary with rostrocaudal position (Choi et al., 2018). We therefore correlated the anteroposterior position of all seven IPS subregions (IPSa-e, cIPS, rIPS) with their respective β values for the Second-order v First-order contrast (Fig. 5). We found a significant correlation between the anteroposterior position, as indexed by the



Figure 6. *A*, cIPS-like regions: four clusters are identified. Voxels only defined in the exploratory analysis are shown in red, voxels only present in the cIPS parcel are displayed in blue, and voxels that overlap with the connectivity-defined cIPS region are shown in purple. *B*, rIPS-like regions: two clusters are identified. Left, voxels only defined in the exploratory analysis are shown in red, voxels only present in the rIPS parcel are displayed in blue, and voxels that overlap with the connectivity-defined rIPS region are shown in purple. However, because of the transformation from a volume to a surface rendering for display purposes, purple voxels are not visible in *B*.

MNI *y* coordinate of the centroid, and β value [r(5) = -0.77, p = 0.044; Fig. 5] This result remained strongly significant when we assessed the correlation within individual subjects, for which the slopes were all negative and different from zero $(r = -0.64 \pm 0.24, t_{(10)} = -8.67, p = 5.8 \times 10^{-6}, d = -2.62;$ Extended Data Table 5-1). Interestingly, the increase in selectivity for the second-order cue in more caudal IPS subregions was topographically opposite to that of frontal cortex, where more rostral regions show greater responsiveness to more abstract stimuli. Notably, the correlations were not significant if we instead correlated anteroposterior position with strength of resting-state functional connectivity to pre-PMd, [r(5) = -0.47, p = 0.29, n.s.] or to PMd [r(5) = 0.67, p = 0.10, n.s.].

To test whether motor planning and attentional control over the relevant first-order cue was similarly represented in spatially graded activity, we next correlated β values of the First-order versus Irrelevant dataset with the MNI *y* coordinate of the seven IPS subregions (IPSa–e, rIPS, cIPS). We found no correlation between these variables [r(5) = 0.55, *p* = 0.202, n.s.]. However, a comparison of the Second-order versus First-order spatial relationship and the First-order versus Irrelevant spatial relationship revealed that these two correlations were significantly different from each other (*z* = -2.31, *p* = 0.010; Fig. 5). These results suggest that a representation of higher-order abstract information within a spatial gradient may be a topographic feature of IPS.

Conjunction maps for parietal cortex

To further determine whether the ROI-based results shown in Figure 4 were specific for connectivity-defined parcels, we expanded our search to the rest of parietal cortex. To this end, we computed conjunction maps within parietal cortex between pertinent contrasts for both cIPS and rIPS (see above, Materials and Methods). To identify additional cIPS-like regions, we evaluated BOLD activity, which, like cIPS, differentiates second-order from first-order cues and high from low coherence but does not distinguish relevant from irrelevant cues or the button press from other task stimuli. As shown in Figure 6A, the largest

cluster, 85 voxels centered at MNI coordinates [-34, -64, 49] includes substantial overlap with the connectivity-defined cIPS region (purple voxels). Additional areas only defined in the exploratory analysis extend forward along the sulcus (red voxels), consistent with the gradient of IPS regions found in Figure 5. The smallest cluster, 17 voxels centered at [38, -59, 47], identifies a homologous area in the right hemisphere. The final two clusters, 47 voxels at [6, -63, 51] and 42 voxels at [-7, -67, 52], encompass medial parietal regions on the borders of the task active mask that are typically grouped with medial parietal default mode areas.

To identify rIPS-like regions, we evaluated BOLD activity, which, like rIPS, differentiates relevant from irrelevant cues and button presses from task stimuli but does not distinguish second-order from first-order cues or high from low coherence. Two clusters of voxels were identified (Fig. 6*B*), one 29-voxel cluster with center of mass at [-39, -41, 58] in which rIPS and the exploratory contrast overlap (left) and one 21-voxel cluster with center of mass at [-52, -27, 43] in which only the exploratory contrast is identified (right). This latter region is within the postcentral gyrus, the presumptive location of primary somatosensory cortex.

Discussion

IPS is implicated in a number of cognitive functions, from the mapping of attention within visual space to action execution. Here, we demonstrate that IPS regions distinguished by their intrinsic functional connectivity with frontal cortex respond differentially to distinct levels of policy abstraction in a manner separable from attentional control, motor planning, sensory features, and motor responding.

These results follow previous work on policy abstraction establishing that parietal cortical activity may respond to increasing abstraction. Nee and D'Esposito (2016) demonstrated that parietal cortex is strongly active across different levels of control demands, although they did not evaluate whether such activity varied parametrically with abstraction. Similarly, in a pioneering rs-fMRI study (Choi et al., 2018) that analyzed previously published task data (Badre and D'Esposito, 2007), lateral frontal regions whose activity correlated with progressively more abstract decisions in caudal-torostral fashion were shown to participate in intrinsic networks whose parietal components were consistent with a rostral-tocaudal gradient. However, their network analysis did not investigate more granular parietal activations in areas implicated in sensorimotor transformations (Erickson and Kayser, 2013). Here, we used a multimodal parcellation to identify specific subdivisions of IPS. Moreover, in a post hoc analysis, we found that the cIPS parcel (parcel 321) identified in our connectivity-based analysis was also the parcel with peak activity for the second-order cue when restricted only by atlas boundaries of parietal cortex (see above, Materials and Methods), strengthening the link between resting-state connectivity and task-based results. Another important difference is neuroanatomical focus; the task in Choi et al. (2018) was designed to understand policy abstraction while controlling for confounds critical in frontal cortex. Here, we build on this work by controlling for additional processes specific to parietal cortex, for example, by equating the featural complexity of visual stimuli.

In keeping with the potential importance of parietal-specific confounds, these data demonstrate that the functions of frontal and parietal areas within intrinsic networks are not homogeneous



Figure 7. *A*, The locations of cIPS and rIPS (dark blue) relative to the most overlapping regions (left, IPS 2; right, IPS 5) from the functional atlas of Wang et al. (2015), where overlap is defined by the number of shared voxels. cIPS and rIPS are largely outside the majority of the probability mass that defines IPS 2 and IPS 5, suggesting that they may be somewhat distinct from these previously defined areas. *B*, The locations of cIPS and rIPS relative to regions defined in a series of other studies. cIPS is closest to the short-term visual working memory area identified by Todd and Marois (2004), although it is somewhat lateral to this region and posterior to the motor planning (action) regions evaluated in Frey et al. (2005) and Culham et al. (2003). Consistent with results that rIPS is more directly involved in defining motor responses than is cIPS, rIPS lies between these previously identified motor planning areas.

(Fig. 4) but instead make distinct contributions to decision-making. Here, we find that cIPS (and not pre-PMd) responds strongly to coherence differences, whereas rIPS (and not PMd) is sensitive to the motor response. The former result is consistent with findings that a region within right inferior parietal lobule correlates with perceptual capacity, whereas an area within left middle frontal gyrus correlates with cognitive control capacity (Eayrs and Lavie, 2019). These findings are also compatible with other significant frontal/parietal dissociations. For example, parietal cortex, which is well situated to serve as an episodic buffer in working memory (Gelastopoulos et al., 2019), might participate in the manipulation but not the monitoring of items within working memory, unlike lateral frontal cortex (Champod and Petrides, 2010). Likewise, in a spatial working memory task performed in patients with lesions, frontal lesions reduced resistance to distraction, whereas parietal lesions impaired the ability to benefit from predictable spatial positions of targets (Saj et al., 2018).

Less consistent with specializations within parietal cortex is the distinction between cIPS and rIPS in the response to coherence. Behaviorally, subjects performed better when both secondand first-order cues were high rather than low coherence (Fig. 2), but only activity in cIPS significantly distinguished the two coherence levels without regard to hierarchical order. This result may reflect the multimodal nature of responses in parietal cortex (Xu, 2018), and IPS in particular. The location of cIPS in this study (MNI centroid -32, -70, 50) approximates the location of IPS2 (Fig. 7), a region previously shown to contain a retinotopic map of spatial attention (Silver and Kastner, 2009; Wang et al., 2015), among other functions (Kastner et al., 2017; Xu, 2018). That cIPS activation varies with coherence across orders may therefore reflect its general responsiveness to differences in stimulus salience. The spatial conjunction of this response to both perceptual and higher-order features might enable cIPS to access both cues as needed, although we did not detect a corresponding change in univariate activity reflecting this conjunction.

More generally, these findings build on work demonstrating that IPS responds to stimuli that are more abstract by definitions other than policy abstraction. Left IPS has been shown to respond in load-dependent fashion for both visual and verbal stimuli, generalizing across sensory modality (Cowan et al., 2011). Moreover, IPS connectivity with object-sensitive regions in temporal cortex distinguishes between conditions in which objects must be either categorized as part of a group or identified as a specific exemplar (McMenamin et al., 2016). In keeping with the known importance of IPS for motor function (Medina et al., 2020), left IPS may also hold effector-independent motor plans (Swinnen et al., 2010). Of course, this diversity of findings emphasizes the importance of defining abstraction, so that common elements that link definitions of abstraction across studies can be discerned (Badre, 2008).

There are also limits to what IPS may encode. IPS may not encode abstract signals related to overall task set or performance monitoring, for example. In support of the primacy of IPS for shaping the response but not the task set, Palenciano et al. (2019) found that IPS represented response complexity but not instructions. Similarly, during a switch task, transcranial magnetic stimulation over left IPS early in the task disrupted updating of response sets but not task goals, whereas the same stimulation later in the preparatory phase also disrupted goal updating, suggesting that IPS might inherit goal updates from linked frontal regions (Muhle-Karbe et al., 2014). These different findings are broadly consistent with a view in which posterior parietal cortex, including the IPS, might represent action-independent, nonspatial visual information (Xu, 2018).

At the same time, cIPS may certainly have capacities for abstract cognitive processes beyond specific second-order policy abstraction decisions. Parietal cortex includes multiple polymodal

areas, and a single macroscopic region can have multiple functions. Moreover, it is not yet known how unitary the cognitive processes that underlie abstract rule processing might be (Badre and Nee, 2018) nor their relationship to other known functions of parietal cortex. To answer such questions, additional task manipulations and converging evidence are likely necessary. For example, to assess whether information manipulation (as opposed to policy abstraction per se) explains the function of this area, one could vary the size of the response set and the level of abstraction independently (Badre and D'Esposito, 2007). In comparison with a two-alternative second-order cue condition, a four-alternative first-order condition would have increased selection demands but reduced abstraction. If cIPS were only responsible for information manipulation independent of abstraction, it might not differentiate these conditions.

Notably, the current data do not directly indicate the mechanisms by which IPS might support policy abstraction. Because the behavioral results indicate that second- and first-order rules have been encoded and retained within a trial, working memory is undoubtedly required. However, for trials in which the second-order rule was encoded in the first position, and subjects therefore retained it until the first-order stimulus was viewed in either position 2 or 3, cIPS activity did not differ between any position (repeated-measures ANOVA, $F_{(5,50)} = 2.19$, p = 0.07, $\eta^2 p = 0.18$). Moreover, cIPS activity was numerically lower for the second-order rule in position 1 (for presumptive encoding of the rule into working memory) than it was in response to the irrelevant cue in position 3 (when the response would already have been decided, and maintaining the second-order rule was no longer necessary). Thus, these data do not strongly support a purely working memory-based account of cIPS activity. However, work that more directly manipulates working memory demands would be important to fully understand mechanisms underlying cIPS activity.

Related to these questions of what IPS encodes and how IPS encodes it is the problem of where IPS encodes it (Fig. 7). Previous work has identified topographically organized IPS subregions, including attentional maps with systematic representations of retinotopic space (Silver et al., 2005) extending along the sulcus (Silver and Kastner, 2009). This organization supports theories that IPS includes priority maps corresponding to the attentional weighting of a spatial location (Jerde and Curtis, 2013), so that lesions in these areas disrupt saccades to remembered locations (Mackey et al., 2016). In contrast, more anteriorsuperior regions of IPS may respond more to features than to locations (Bettencourt and Xu, 2016), with the most anterior regions integrating feature and action to enable tool use (Kastner et al., 2017). How the abstract representations evaluated here are organized with respect to other feature responses in previously defined IPS subdivisions (e.g., retinotopic IPS0-5) remains an important topic for future studies.

Altogether, these findings demonstrate that topographically organized regions within IPS respond differentially to policy abstraction, both distinct from representations of attentional control, sensory features, and motor responses, and within specific networks identifiable by their intrinsic frontal connectivity. This work builds on previous studies that emphasize the polymodal nature of parietal cortex (Goldman-Rakic, 1988; Mesulam, 1998) and the importance of IPS in other complex cognitive processes such as sensorimotor transformations (Erickson and Kayser, 2013; Kastner et al., 2017; Xu, 2018). Future work should continue to define how the representations and connectivity of IPS contribute to human performance in abstract, hierarchically organized tasks.

References

- Badre D (2008) Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. Trends Cogn Sci 12:193–200.
- Badre D, D'Esposito M (2007) Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. J Cogn Neurosci 19:2082–2099.
- Badre D, D'Esposito M (2009) Is the rostro-caudal axis of the frontal lobe hierarchical? Nat Rev Neurosci 10:659–669.
- Badre D, Nee DE (2018) Frontal cortex and the hierarchical control of behavior. Trends Cogn Sci 22:170–188.
- Badre D, Kayser AS, D'Esposito M (2010) Frontal cortex and the discovery of abstract action rules. Neuron 66:315–326.
- Bettencourt KC, Xu Y (2016) Decoding the content of visual short-term memory under distraction in occipital and parietal areas Nat Neurosci 19:150–157.
- Botvinick MM (2008) Hierarchical models of behavior and prefrontal function. Trends Cogn Sci 12:201–208.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci 12:4745–4765.
- Champod AS, Petrides M (2010) Dissociation within the frontoparietal network in verbal working memory: a parametric functional magnetic resonance imaging study. J Neurosci 30:3849–3856.
- Choi EY, Drayna GK, Badre D (2018) Evidence for a functional hierarchy of association networks. J Cogn Neurosci 30:722–736.
- Cowan N, Li D, Moffitt A, Becker TM, Martin EA, Saults JS, Christ SE (2011) A neural region of abstract working memory. J Cogn Neurosci 23:2852– 2863.
- Cox RW (1996) AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput Biomed Res 29:162–173.
- Culham JC, Valyear KF (2006) Human parietal cortex in action. Curr Opin Neurobiol 16:205–212.
- Culham JC, Danckert SL, DeSouza JFX, Gati JS, Menon RS, Goodale MA (2003) Visually guided grasping produces fMRI activation in dorsal but not ventral stream brain areas. Exp Brain Res 153:180–189.
- D'Mello AM, Gabrieli JDE, Nee DE (2020) Evidence for hierarchical cognitive control in the human cerebellum. Curr Biol 30:1881–1892.e3.
- Eavrs JO, Lavie N (2019) Individual differences in parietal and frontal cortex structure predict dissociable capacities for perception and cognitive control. Neuroimage 202:116148.
- Erickson DT, Kayser AS (2013) The neural representation of sensorimotor transformations in a human perceptual decision making network. Neuroimage 79:340–350.
- Frey SH, Vinton D, Norlund R, Grafton ST (2005) Cortical topography of human anterior intraparietal cortex active during visually guided grasping. Brain Res Cogn Brain Res 23:397–405.
- Gelastopoulos A., Whittington MA, Kopell NJ (2019). Parietal low beta rhythm provides a dynamical substrate for a working memory buffer. Proc Natl Acad Sci U S A 116:16613–16620.
- Gillebert CR, Mantini D, Thijs V, Sunaert S, Dupont P, Vandenberghe R (2011) Lesion evidence for the critical role of the intraparietal sulcus in spatial attention. Brain 134:1694–1709.
- Goldman-Rakic PS (1988) Topography of cognition: parallel distributed networks in primate association cortex. Annu Rev Neurosci 11:137–156.
- Gordon EM, Laumann TO, Gilmore AW, Newbold DJ, Greene DJ, Berg JJ, Ortega M, Hoyt-Drazen C, Gratton C, Sun H, Hampton JM, Coalson RS, Nguyen AL, McDermott KB, Shimony JS, Snyder AZ, Schlaggar BL, Petersen SE, Nelson SM, Dosenbach NUF (2017) Precision functional mapping of individual human brains. Neuron 95:791–807.e7.
- Grèzes J, Decety J (2001) Functional anatomy of execution, mental simulation, observation, and verb generation of actions: a meta-analysis. Hum Brain Mapp 12:1–19.
- Heekeren HR, Marrett S, Ruff DA, Bandettini PA, Ungerleider LG (2006). Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. Proc Natl Acad Sci U S A 103:10023–10028.

- Jeon H-A, Friederici AD (2013) Two principles of organization in the prefrontal cortex are cognitive hierarchy and degree of automaticity. Nat Commun 4:2041.
- Jeon HA, Anwander A, Friederici AD (2014) Functional network mirrored in the prefrontal cortex, caudate nucleus, and thalamus: high-resolution functional imaging and structural connectivity. J Neurosci 34:9202–9212.
- Jerde TA, Curtis CE (2013) Maps of space in human frontoparietal cortex. J Physiol Paris 107:510–516.
- Karnath HO, Perenin MT (2005) Cortical control of visually guided reaching: evidence from patients with optic ataxia. Cereb Cortex 15:1561–1569.
- Kastner S, Chen Q, Jeong SK, Mruczek REB (2017) A brief comparative review of primate posterior parietal cortex: a novel hypothesis on the human toolmaker. Neuropsychologia 105:123–134.
- Kayser AS, D'Esposito M (2013) Abstract rule learning: the differential effects of lesions in frontal cortex. Cereb Cortex 23:230–240.
- Kayser AS, Buchsbaum BR, Erickson DT, D'Esposito M (2010a) The functional anatomy of a perceptual decision in the human brain. J Neurophysiol 103:1179–1194.
- Kayser AS, Erickson DT, Buchsbaum BR, D'Esposito M (2010b) Neural representations of relevant and irrelevant features in perceptual decision making, J Neurosci 30:15778–15789.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat Neurosci 2:176–185.
- Konen CS, Kastner S (2008) Representation of eye movements and stimulus motion in topographically organized areas of human posterior parietal cortex. J Neurosci 28:8361–8375.
- Mackey WE, Devinsky O, Doyle WK, Golfinos JG, Curtis CE (2016) Human parietal cortex lesions impact the precision of spatial working memory. J Neurophysiol 116:1049–1054.
- McMenamin BW, Marsolek CJ, Morseth BK, Speer MF, Burton PC, Burgund ED (2016) Conflicting demands of abstract and specific visual object processing resolved by frontoparietal networks. Cogn Affect Behav Neurosci 16:502–515.
- Medina J, Jax SA, Coslett HB (2020) Impairments in action and perception after right intraparietal damage. Cortex 122:288–299.
- Mesulam MM (1998) From sensation to cognition. Brain 121:1013-1052.
- Moussa MN, Steen MR, Laurienti PJ, Hayasaka S (2012) Consistency of network modules in resting-state fMRI connectome data. PLoS One 7: e44428.
- Muhle-Karbe PS, Andres M, Brass M, Andres M (2014) Transcranial magnetic stimulation dissociates prefrontal and parietal contributions to task preparation. J Neurosci 34:12481–12489.

- Murphy AP, Leopold DA, Humphreys GW, Welchman AE (2016) Lesions to right posterior parietal cortex impair visual depth perception from disparity but not motion cues. Phil Trans R Soc B 371:20150263.
- Nee DE, Brown JW (2012) Rostral-caudal gradients of abstraction revealed by multi-variate pattern analysis of working memory. Neuroimage 63:1285–1294.
- Nee DE, Brown JW (2013) Dissociable frontal-striatal and frontal-parietal networks involved in updating hierarchical contexts in working memory. Cereb Cortex 23:2146–2158.
- Nee DE, D'Esposito M (2016) The hierarchical organization of the lateral prefrontal cortex. Elife 5:e12112.
- Palenciano AF, González-García C, Arco JE, Pessoa L, Ruz M (2019) Representational organization of novel task sets during proactive encoding. J Neurosci 39:8386–8397.
- Saj A, Verdon V, Hauert CA, Vuilleumier P (2018) Dissociable components of spatial neglect associated with frontal and parietal lesions. Neuropsychologia 115:60–69.
- Schaefer A, Kong R, Gordon EM, Laumann TO, Zuo X-N, Holmes AJ, Eickhoff SB, Yeo BTT (2018) Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. Cereb Cortex 28:3095–3114.
- Shankar S, Kayser AS (2017) Perceptual and categorical decision making: goal-relevant representation of two domains at different levels of abstraction. J Neurophysiol 117:2088–2103.
- Silver MA, Kastner S (2009) Topographic maps in human frontal and parietal cortex. Trends Cogn Sci 13:488–495.
- Silver MA, Ress D, Heeger DJ (2005) Topographic maps of visual spatial attention in human parietal cortex. J Neurophysiol 94:1358–1371.
- Swinnen SP, Vangheluwe S, Wagemans J, Coxon JP, Goble DJ, Van Impe A, Sunaert S, Peeters R, Wenderoth N (2010) Shared neural resources between left and right interlimb coordination skills: the neural substrate of abstract motor representations. Neuroimage 49:2570–2580.
- Todd JJ, Marois R (2004) Capacity limit of visual short-term memory in human posterior parietal cortex. Nature 428:751–754.
- Wang L, Mruczek REB, Arcaro MJ, Kastner S (2015) Probabilistic maps of visual topography in human cortex. Cereb Cortex 25:3911–3931.
- Xu Y (2009) Distinctive neural mechanisms supporting visual object individuation and identification. J Cogn Neurosci 21:511–518.
- Xu Y (2018) The posterior parietal cortex in adaptive visual processing. Trends Neurosci 41:806–822.
- Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR, Fischl B, Liu H, Buckner RL (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J Neurophysiol 106:1125–1165.