

Neural correlates of decision-making in human parietal cortex

By

Melissa Lynn Newton

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Andrew Kayser, Co-Chair
Professor Mark D'Esposito, Co-Chair
Professor Michael Silver

Summer 2021

Abstract

Neural correlates of decision-making in human parietal cortex

by

Melissa Lynn Newton

Doctor of Philosophy in Neuroscience

University of California, Berkeley

Professor Andrew Kayser, Co-Chair
Professor Mark D'Esposito, Co-Chair

Parietal cortex has long been considered essential for complex executive functions necessary during human decision-making (Mesulam, 1998). Following an introduction, Chapter 2 of this dissertation investigates one particularly pertinent region of parietal cortex called the intraparietal sulcus (IPS) in sensory and abstract decision-making processes. The 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 functions in sensorimotor processing often act as confounds. This thesis addresses this problem using a novel task developed to dissociate abstract decision-making from sensory salience, attentional control, and motor output. Functional MRI data were collected from healthy female and male 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 (pre-PMd) was preferentially active for second-order cues, while a rostral IPS subregion (rIPS) with strongest connectivity to the dorsal premotor cortex (PMd) was active for 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 fronto-parietal subnetworks that support hierarchical cognitive control.

Chapter 3 next assessed whether parietal areas involved in decision-making vary as a function of strategy. One can perform the same goal-oriented task using different strategies (Morrison, Rosenbaum, Fair, & Chein, 2016). Studies have shown dissimilarities in neural activity for strategies that depend on unique sensory or mnemonic systems (Iaria, Petrides, Dagher, Pike, & Bohbot, 2003), but few have investigated if different strategies in performing a cognitive control task produce activation differences within the fronto-parietal control network if those strategies differ

in the form of abstraction. Half of all trials in the same abstract decision-making task developed for this thesis required use of a hierarchically-defined “policy strategy”, and the other half permitted a heterarchical match/nonmatch “comparison strategy” (Badre, 2008). This unique design feature was next used to investigate whether frontal and parietal control regions implicated in policy abstraction differentially represent strategy when it employs different forms of abstraction. Previously collected fMRI data from healthy female and male subjects was re-analyzed. Accuracy was higher for trials allowing the comparison strategy, and a significant behavioral interaction was observed between sensory coherence of the second-order cue and strategy trial type. Across the brain, known decision-making regions for hierarchical policy abstraction tasks had increased activity for the policy strategy, including frontal regions pre-PMd and PMd. A region in the superior frontal sulcus covaried activity with subject strategy use. However, parietal activity did not strongly differentiate between the strategies, suggesting that IPS subregions specifically important for policy abstraction decision information may not preferentially represent control of hierarchical policy abstraction over other non-hierarchical types of task abstraction.

Finally, Chapter 4 explores if retained versus removed working memory items are reflected in parietal activity during decision-making. Every discrete decision requires selection from at least two unique options, and in the absence of consistent external representations, those options are held in working memory. As a result, relevant working memory decision items are retained for use in a future motor response while irrelevant items must be removed to reduce retrieval interference (Lewis-Peacock, Drysdale, Oberauer, & Postle, 2012). Parietal cortex and IPS specifically are associated with working memory item representation and control (Chatham, Frank, & Badre, 2014; Cowan et al., 2011; M. D’Esposito et al., 1998), implicating it in possible retention and removal processes during decision-making. To test whether an active process of removal and/or retention exists in parietal cortex, fMRI data collected from the same abstract decision-making task developed for this thesis work were re-analyzed. A subset of the data were selected where second-order cues were presented in position 2 within a trial, effectively acting as a retro-cue for the relevance or irrelevance of the previously presented cue in position 1. Investigating activity at this timepoint would reveal working memory item removal if the previous cue was irrelevant, and retention if the cue was relevant. Hypothesis 1 predicted that parietal cortex would reflect an active process of content unbinding during irrelevant working memory item removal, but results instead suggested that removal may be a latent process that is undetectable via univariate imaging. Hypothesis 2 investigated working memory item retention and predicted that rIPS would increase activity for retaining response-relevant first-order contents while cIPS would not change activity for working memory control of first-order information. Contrary to these hypotheses, results revealed that rIPS did not significantly change activity for working memory item retention of relevant first-order information, while cIPS increased activity at this timepoint. This finding supports an alternative hypothesis that cIPS and rIPS should instead be considered to be primarily engaged with decision-making aspects of the policy abstraction task, and not necessarily implicated in aspects of working memory item retention. Continued research is necessary to further expose the neural underpinnings of working memory item removal and retention.

Dedication

“This one is dedicated to all the ravers in the nation.” – Dune, *Hardcore vibes*

To Amy’s Kitchen pre-prepared frozen meals.

Table of Contents

Abstract.....	1
Dedication.....	i
Table of Contents.....	ii
List of Figures.....	iii
Acknowledgements.....	iv
Publication Related to This Work.....	v
Chapter 1: Introduction.....	1
Chapter 2: What areas of IPS are involved in sensory and abstract decision-making processes, and are they separable?	7
Chapter 3: Does activity in parietal areas involved in decision-making vary as a function of strategy?	25
Chapter 4: Are retained versus removed working memory items reflected in parietal activity during decision-making?	38
Chapter 5: Conclusions.....	48
References.....	50

List of Figures

Figure 1: Chapter 2, Task design.....	9
Figure 2: Chapter 2, Behavioral effect of coherence on accuracy.....	15
Figure 3: Chapter 2, Univariate whole brain task activation with ROIs.....	16
Figure 4: Chapter 2, ROI results.....	17
Figure 5: Chapter 2, Abstraction gradient.....	19
Figure 6: Chapter 2, Extended data for all ROIs.....	24
Figure 7: Chapter 3, Strategy trial types.....	27
Figure 8: Chapter 3, Behavioral effect of strategy.....	29
Figure 9: Chapter 3, Additional behavioral measures of strategy.....	30
Figure 10: Chapter 3, Univariate results for strategy trial types.....	31
Figure 11: Chapter 3, ROI results for strategy trial types.....	32
Figure 12: Chapter 3, Region that covaries with abstraction strategy.....	33
Figure 13: Chapter 4, Working memory removal/retention trial subtypes.....	41
Figure 14: Chapter 4, Behavioral effect of cue position.....	43
Figure 15: Chapter 4, Univariate results for working memory item removal/retention....	44

Acknowledgements

Warm thank you to my thesis advisor, Dr. Andrew Kayser, for the continuous support. I was profoundly lucky to have an advisor that I truly enjoyed working with.

Thank you to committee members Dr. Mark D'Esposito and Dr. Michael Silver for contributions to my scientific development. And huge thank you Dr. William Prinzmetal for serving as an invaluable committee member despite administrative shenanigans from the online filing system. Although he is not listed officially as a thesis committee member, he contributed substantially to my advancement in the PhD program over the years and provided much needed support.

Thank you to the following organizations, departments and funding sources for assistance across both personal and academic domains: UC Berkeley Graduate Division Mentored Research Award, Helen Wills Neuroscience Institute, UC Berkeley Office for Graduate Diversity, UC Berkeley Disabled Students Program, The University Section Club at Berkeley, UC Berkeley Graduate Assembly, UC San Francisco Neurology Department, University Health Services at UC Berkeley

Gratitude to the following scientist friends and colleagues who aided immensely in numerous formal and informal endeavors over the course of my graduate journey: Swetha Shankar, Savannah Cookson, Jocelyn Breton, Liz Lawler, Elizabeth Lorenc, Dan Lurie, Daniella Furman, Katarina Slama, Vanessa Carels, Chris Rodgers, Liberty Hamilton, Kelly Clancy, Katelyn Begany, Amy LeMessurier, Kim Long, Wren Thomas, Yvonne Fonken, Alex Naka, Marina Garrett, Charles Frye, Jon Jui, Eric Weiss, Cameron Baker, and all the sweet Knight Lab postdocs who chatted with me about research over tea in the Barker kitchen

Thank you to my fantastic undergraduate research assistants: Joshua Price, Angie Wong, and Karen Lee

Heartfelt thanks to these departmental administrative heroes and heroines: Candace Groskreutz, Kati Markowitz, John Schindel

Extra special acknowledgements to non-science friends and family who carried me throughout these long years in my lowest points: Barron the Cocker Spaniel and best dog-child, Alice the probable Himalayan Spaniel and tied for best dog-child, Jeff Bliss the best husband, Andrea Newton the best sister, and the following best friends across categories: Christina Patel, Nathalia Torres, Emily Partridge, my Southwind Circle family

<3

And finally, thank you to all the haters, ya'll know who you are, for giving me the opportunity to prove to myself that I can achieve my highest aspirations in spite of significant barriers and without your approval.

Publication Related To This Work

"Connectivity-defined subdivisions of the intraparietal sulcus respond differentially to abstraction during decision making", manuscript submitted for publication (2021-06-16)

Chapter 1: Introduction

Decision-making is a “deliberative process that results in the commitment to a categorical proposition” (Gold & Shadlen, 2007). Each day, we make a vast multitude of decisions. Some are concrete and immediately enacted, such as seeing a road hazard while driving and turning the steering wheel left or right to avoid it. Other decisions are relatively abstract, such as turning your steering wheel either left or right while driving to optimize your route based on the context of the current time in order to avoid rush hour traffic.

The decision-making process is incredibly rich in naturalistic environments, even for seemingly straightforward choices. Multiple neural systems are necessary depending on the nature of the decision; for example, visual processing streams necessary in visual tasks and motor systems for indicating decision responses. But many aspects of decision-making that occur between sensory processing and motor production are still not understood in terms of which brain areas are needed and how those regions support the decision process. Real-world decisions necessitate neural regions that successfully integrate multiple information streams.

Consequently, this thesis will investigate various aspects of decision-making and their associated neural correlates, focusing on contributions of the parietal cortex. As an established component of a frontoparietal network (FPN) control system, the multidimensional functions of parietal cortex are investigated here to further delineate the extent, spatial localization, and functional specificity of cognitive control processes. Cognitive control is the capacity to organize goal-directed behaviors that take place in novel surroundings, that act against habits, or that adjudicate multiple competing actions (Badre & Nee, 2018; Duncan, Emslie, Williams, Johnson, & Freer, 1996). Cognitive control and decision-making are highly related research subfields in the sense that cognitive control tasks require decisions to be made by virtue of their design, and complex decision-making tasks rely on the appropriate amount of cognitive control to be applied in order to be successfully executed.

Cognitive control orchestrates an assortment of interrelated executive functions required during decision-making. Among these functions include the assessment of current sensory information, synthesis of that information with goals and subgoals, implementation of a strategy for reaching the decision, subsequent selection of a relevant decision option and removal of irrelevant options, and then performance of motor responses in support of that decision. To address this fact, I developed a novel, high-dimensional, abstract decision-making task that allows for controlled investigation of each of those processes within parietal cortex. By collecting fMRI data from healthy adult human subjects while they perform the task across 660 trials, I am able to isolate neural activity at particular conditions when distinct executive functions are performed during decision-making.

Abstract decision-making

Abstract decisions are those that require internally-held conceptual information such as a goals that are not necessarily represented in external sensory stimuli. Abstract information can take the form of contextual rules, motivations, or schema, and

decisions utilizing this information often require sequences of behavior that are enacted over time. As such, abstract decisions require increased cognitive control. A broad range of tasks are used to study abstract decision-making. Variations of the classic Stroop task have been used, which requires subjects to report either the color of text of a displayed word, or the word itself. Increased cognitive control is required to override interference when a word is inconsistent with the color that it is displayed in (e.g. when the word “blue” is displayed in red font) (Stroop, 1935). Other tasks have used variations of the AX-continuous performance test used to clinically assess control (Cohen, Barch, Carter, & Servan-Schreiber, 1999), where serially presented letters must be monitored in order to respond uniquely to an “X” but only if it was preceded by an “A”. Control is assessed with consideration to statistically-manipulated probe combinations such as B-X presentations, where subjects must inhibit responses by maintaining the appropriate goal state and processing the current context (Braver, 2012).

This dissertation will focus explicitly on policy abstraction (Badre & D’Esposito, 2009; Badre & Nee, 2018). A policy is another term for an abstract rule that asserts the contextual relationship of a stimulus with an action or another rule. For example, a common policy in everyday life states that when a traffic light is red, that means to press the brake pedal to stop the car, but if the light is green, then we can continue through the intersection by pressing the acceleration pedal. Policy abstraction tasks train subjects on novel rules that they use to select a single response out of multiple competing potential responses. The lowest level of policy abstraction, or first-order decisions, are concrete rules that link just one external stimulus to one motor response (Badre & D’Esposito, 2007). For example, if there are more blue dots being displayed on a screen, then press button 1 on a button box, but if there are actually more gray dots on the screen, then press button 2. Higher-order policy abstraction decisions (e.g., second-order, third-order, and so on) require multiple layers of hierarchical rules to set the context for lower-order decisions and are increasingly further away conceptually from initiating an appropriate motor response. These rules are hierarchical because they dictate the behavioral relevance of lower order rules. For example, a second-order policy abstraction rule may govern if a first-order cue displaying blue and gray dots is relevant for response, or if a different first-order cue of displayed shapes is the relevant rule for response, rendering the stimulus of colored dots useless for action production under some circumstances (See Chapter 2).

Abstract decisions are not all hierarchical. A higher-order decision could be heterarchical, such as in the case of match/nonmatch discriminations in relational complexity tasks (Badre, 2008; Badre & D’Esposito, 2009; Badre & Nee, 2018). Heterarchical organization requires “neither modulation nor inheritance of information asymmetrically from higher to lower levels” (Badre, 2008). In these tasks, a heterarchical higher-order rule may assert that if two lower-order cues match, then to press button 1, but if the lower-order cues do not match, then to press button 2. A variant of heterarchical abstraction decisions using match/nonmatch rules is described as a strategy of task performance in Chapter 3.

Perceptual decision-making

Abstract decision-making differs from other types of decisions in the components of the decision space as well as in the neural representations themselves. In contrast to abstract decisions, perceptual decisions have been extensively studied in non-human primates and rodents (Brody & Hanks, 2016; Schall, 2003), revealing separate but partially overlapping neural mechanisms. Perceptual decision-making (PDM) is a fundamental cognitive process in which sensory information provides the evidence necessary to select an action among multiple options. PDM depends on the resolution of sensory uncertainty, when there is some amount of sensory noise that impedes a discrimination judgement. Common visual PDM tasks use a random dot kinetogram stimulus consisting of a field of moving dots, with a proportion that are moving coherently in the same direction amidst the remainder moving in random directions (Gold & Shadlen, 2003). This task requires subjects to decide the direction of coherent motion and communicate that direction with an associated motor response. Other PDM tasks have required subjects to discern the category of a static image (e.g. either a human face or a house) that has been degraded by statistically-introduced visual noise (H. R. Heekeren, Marrett, Bandettini, & Ungerleider, 2004), or in the case of rodents, to discern if either a left or right speaker had a greater number of clicks when simultaneous streams of auditory tones were played (Brunton, Botvinick, & Brody, 2013).

The developed task used in this dissertation combines elements of the random dot kinetogram to assess the effect of sensory salience during decision-making (see Chapter 2). When there is increased noise in a low coherence stimulus, then sensory salience is low, and when a stimulus has high coherence then sensory salience is high. Among the functions of parietal cortex, sensory salience has been demonstrated as a prominently represented factor (Buschman & Miller, 2007; Chen et al., 2020). It was thus used to dissociate aspects of abstract decision-making within the experimental design. Since PDM was not the central focus of this thesis, the term “perceptual decision” is used interchangeably with the more general term “sensory decision”.

Frontal cortex and abstract/sensory decisions

Lateral frontal cortex and prefrontal cortex (PFC) are known to be fundamental for both abstract and perceptual decisions (Hanks et al., 2015; Hauke R. Heekeren, Marrett, & Ungerleider, 2008; Rahnev, Nee, Riddle, Larson, & D’Esposito, 2016) but are particularly implicated in abstract tasks. As a human subject performs decisions that are increasingly abstract and hierarchical in organization, then progressively more rostral portions of lateral prefrontal cortex (LPFC) become engaged (Badre & D’Esposito, 2007; Koechlin & Jubault, 2006; Koechlin, Ody, & Kouneiher, 2003). First-order abstract decisions have been shown to activate dorsal premotor cortex (PMd) within the frontal lobe, while second-order decisions activate a more rostral region called dorsal pre-premotor cortex (pre-PMd) in addition to PMd. Third- and fourth-order decisions additionally activate anterior frontal regions including mid-dorsolateral prefrontal cortex (mid-DLPFC) and rostrolateral prefrontal cortex (RLPFC) (Badre, 2008; Badre & D’Esposito, 2007; Badre & Nee, 2018), although the details and extent of hierarchical organization within frontal cortex is still debated (Fedorenko, Duncan, & Kanwisher, 2013; Pischedda, Go, Haynes, & Reverberi, 2017).

For performance of perceptual decisions, additional areas of frontal cortex have been implicated for different task demands. PDM tasks have notably been shown to activate the frontal eye fields (FEF) in visuo-motor attention (Vernet, Quentin, Chanes, Mitsumasu, & Valero-Cabré, 2014), superior frontal gyrus for strength of evidence in the decision (H. R. Heekeren et al., 2004), and middle frontal gyrus for additional decision-related information (Kayser, Buchsbaum, Erickson, & D'Esposito, 2010). Despite some differences in regional frontal activation, perceptual and abstract decisions have both been strongly associated with activity across the DLPFC (Kim & Shadlen, 1999).

Parietal cortex and abstract/sensory decisions

The direct role of parietal cortex in abstract decisions, particularly policy abstraction, has been comparatively less studied than frontal cortex. Despite this, rule-based tasks have repeatedly implicated parietal cortex in decisions that are detached from sensorimotor functions (S. A. Bunge, 2004). It was recently shown that the inferior parietal lobules bilaterally represent transferable latent states of an abstract task, or states that cannot be resolved using sensory information alone (Vaidya, Jones, Castillo, & Badre, 2021). Similarly, bilateral inferior parietal lobules and the left intraparietal sulcus (IPS) were active for discovery of hierarchical structure in abstraction tasks (Badre, Kayser, & D'Esposito, 2010; Eichenbaum, Scimeca, & D'Esposito, 2020).

As for sensory decisions, parietal participation in PDM is accepted to be crucial. Activity specifically in IPS is notably involved. Initially documented in the rhesus macaque monkey lateral intraparietal area (LIP) (Shadlen & Newsome, 2001), subregions of human IPS have been noted to change activity with respects to the amount of available sensory information accumulated over time leading up to a threshold when the decision is made (H. R. Heekeren et al., 2004; Hauke R. Heekeren et al., 2008; Kayser, Buchsbaum, et al., 2010).

Network relationships of the IPS and frontal cortex in cognitive control

Regional differences in frontal and parietal cortex across abstract and PDM tasks can be explored using network analyses, potentially shedding light on their differential contributions to cognitive control. It is well established that frontal and parietal cortices are structurally connected (Petrides, Tomaiuolo, Yeterian, & Pandya, 2012; Thiebaut de Schotten, Dell'Acqua, Valabregue, & Catani, 2012), and resting state network functional MRI (rsfMRI) connectivity analyses confirm the high correlation of activity between frontal and parietal cortices (Yeo et al., 2011), with existent evidence of agreement between the strength of anatomical and resting-state connectivity within a rsfMRI network (Greicius, Supekar, Menon, & Dougherty, 2009). The frontoparietal control network, also referred to as the FPN, includes portions of RLPFC, middle frontal gyrus, superior frontal gyrus, superior frontal sulcus, inferior parietal lobule, and IPS (Dixon et al., 2018; Yeo et al., 2011). These regions exhibit flexible connectivity across networks, particularly with the dorsal attention network (DAN) in attentional control (Yeo et al., 2011). The DAN, in turn, heavily influences activity in the somato-motor network engaged for perceptual and motor processes (Buschman & Kastner, 2015; Corbetta & Shulman, 2002).

When regions involved with hierarchical cognitive control for policy abstraction are categorized by their large-scale intrinsic networks established by Yeo et al (2011),

frontal region pre-PMd implicated in second-order decisions belongs to the FPN (control network), while region PMd implicated in first-order decisions belongs to the somato-motor network (Schaefer et al., 2018). Recent evidence suggests that pre-PMd and PMd may belong to distinct task-active subnetworks with parietal cortex that support decision-making at different levels of abstraction (Choi, Drayna, & Badre, 2018). It has not yet been determined whether the contribution of parietal cortex in policy abstraction is to represent attentional and sensorimotor components of the decision akin to functions during PDM, or if parietal cortex differentially represents decision information based on its level of abstraction.

Chapter 2 of this thesis identifies subnetworks of pre-PMd and PMd with parietal cortex to investigate multiple aspects of parietal functions including representations of sensory salience necessary during PDM, attentional control, motor output, and level of policy abstraction. If IPS differentiates activity based on level of task abstraction, then it may serve as a link between contextual goals at the second-order of abstraction and stimulus-motor relationships at the first-order of abstraction by interfacing with frontal nodes in the FPN and the somato-motor network (Cole et al., 2013). Using a fine-grained parcellation that divides the cerebral cortex into 1000 distinct areas (Schaefer et al., 2018), Chapter 2 additionally sheds light on the role of IPS by recognizing functional subdivisions as a reconciling property to explain multifaceted activity.

Representations of strategy in the IPS

The same task can be completed using different cognitive strategies and subsequently engage differential brain regions. This is seen for example in findings that reading aloud can recruit PMd for a non-semantic strategy reliant on phonological word properties, or can recruit ventral inferior frontal cortex for a semantic strategy reliant on mnemonic factors (Binder, Medler, Desai, Conant, & Liebenthal, 2005; Kherif, Josse, Seghier, & Price, 2009). In the context of abstract decision-making tasks, strategies that differ based on the type of abstract rules being applied but that produce the same action outcomes potentially engage differential brain regions (Nee, Jahn, & Brown, 2014). Parietal cortex has been implicated in a number of related processes for performance of rule-based tasks. Formative studies have shown parietal activation during rule-switching (Silvia A Bunge et al., 2003; Rushworth, Hadland, Paus, & Sipila, 2002), long-term storage of action rules (Silvia A. Bunge, Hazeltine, Scanlon, Rosen, & Gabrieli, 2002), and rule-specific information (Sakai & Passingham, 2006) that is important for different strategies. It has been suggested that activity specifically in IPS initiates and adapts control between task trials (Dosenbach et al., 2007), which is necessary when different rule-based strategies are employed across trial types. The method by which a rule is enacted for decision-making has further been shown to activate bilateral IPS, with increased activity for self-chosen rather than an experimenter-chosen rule set (Zhang, Kriegeskorte, Carlin, & Rowe, 2013).

Chapter 3 explores the finding that the same abstract decision-making task developed for this thesis could be performed using either a hierarchically-organized strategy dependent on use of a second-order policy (“policy strategy”), or by using a heterarchical match/nonmatch strategy (Badre, 2008) that compares response outcomes of first-order rules (“comparison strategy”). In this case, the policy strategy employs a rule to assess the second-order cue in order to next assess the relevant first-

order cue for a response, as instructed during training. Alternatively, the subject-initiated comparison strategy utilizes a rule in which they may assess only the two first-order cues to determine if they are a match or nonmatch in motor output for response.

Previous abstract decision-making research has failed to show consistent differences in brain activity between tasks that are hierarchical or heterarchical in design (Badre & Nee, 2018; Nee et al., 2014; Reynolds, Reilly, Cohen, & Braver, 2012), but those studies were exploratory in nature. Building on findings from Chapter 2, the explicit representation of hierarchical policy abstraction as a strategy is investigated within ROIs shown to be important for policy-specific decisions, providing defined hypotheses for expected activity. The specificity of IPS subregions that represent higher order policy information is explored, potentially delineating distinct roles of parietal and frontal cortices in abstraction strategy.

Representations of working memory item retention/removal in the IPS

The posterior parietal cortex (PPC) has long been considered a potential target for working memory associated with the control of abstract decision-making (Mesulam, 1998). Early fMRI studies have asserted that PPC interacts with the PFC to retrieve and maintain relevant rule-informed response contingencies across delays (Silvia A Bunge et al., 2003; M. D'Esposito et al., 1998), with sustained delay period activation for both lower-order and higher-order rule-based tasks (Toni, 1999) that is not affected by sensory modality of the cue type associated with the rule (Silvia A Bunge et al., 2003). More recent research using a hierarchical second-order abstraction task found that when working memory contents were selected for behavior in a process termed “output gating” (Badre, 2012; Frank & Badre, 2012), the PPC increased activity compared to when the relevant response information did not need to be maintained across a trial (Chatham et al., 2014). This implicates parietal regions involved in task abstraction for working memory item retention in the ultimate service of upcoming action output.

But what happens to abstract decision information held in working memory that is determined to be irrelevant for behavior? This information must be efficiently removed, or it can cause retrieval interference with relevant information (Bjork, 1970; Lechner, Squire, & Byrne, 1999) and takes up valuable space within limited working memory capacity (Buschman, Siegel, Roy, & Miller, 2011). Behavioral and modeling studies suggest that working memory item removal may be an active process that unbinds sensory details about the item from other details regarding the context in which it was encoded (N. Burgess & Hitch, 2006; Lewis-Peacock, Kessler, & Oberauer, 2018).

It has not been determined how activity in specific IPS subregions important for second-order policy abstraction might reflect when relevant decision information is retained to generate a future motor response or when irrelevant information is removed. Using the same policy abstraction task developed for this thesis, Chapter 4 re-analyzes a subset of trials at timepoints where first-order information held in working memory is retro-cued for relevance or irrelevance, allowing contrasts of neural activity for directed retention or removal, respectively. The process of active unbinding is investigated for working memory item removal, with results contributing to a growing body of research suggesting latent neural mechanisms.

Chapter 2:

What areas of IPS are involved in sensory and abstract decision-making processes, and are they separable?

Abstract

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, and motor output. Functional MRI data were collected from healthy female and male 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 (pre-PMd) was preferentially active for second-order cues, while a rostral IPS subregion (rIPS) with strongest connectivity to the dorsal premotor cortex (PMd) was active for attentional control over first-order cues. These effects of abstraction were seen in addition to cIPS activity that was specific for 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 in IPS, mirroring the 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.

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 & Kastner, 2008), representation of visual features (Xu, 2009), action planning (Grèzes & Decety, 2000), action execution (Culham & Valyear, 2006), and perceptual decision making (Kayser, Buchsbaum, et al., 2010). The engagement of IPS in these varied processes reflects its key role in mediating stimulus-response relationships within the frontoparietal network (FPN), 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, Leopold, Humphreys, & Welchman, 2016), and reaching (Karnath & Perenin, 2005; Medina, Jax, & Coslett, 2020).

In line with this evidence that IPS plays a key role in linking stimuli with appropriate actions, it has recently been suggested that parietal cortex may contribute to decision making through participation in networks distinguished by their level of policy abstraction (Choi, Drayna, Badre 2018). Policy abstraction (Badre et al., 2009; 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 food for 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 & D’Esposito, 2007; Badre et al., 2010). This literature has identified specific regions along the rostral/caudal axis of the lateral frontal cortex that respond hierarchically to increasingly abstract decisions (Nee & 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 & D’Esposito, 2013). Like lateral frontal cortex, regions within the IPS are also higher-order (“polymodal”) areas (Mesulam, 1998) that have 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 of policy abstraction.

To address this possibility, here we take advantage of the fact that the brain is functionally segregated into large-scale intrinsic networks that are highly reproducible across subjects (Moussa, Steen, Laurienti, & Hayasaka, 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 & 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 utilize an adaptation of the random-dot-kinetogram task (Britten, Shadlen, Newsome, & Movshon, 1992; Kayser, Erickson, Buchsbaum, & D’Esposito, 2010) 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 be distinguishable from 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 helps to explain the multifaceted functions of parietal cortex.

Materials and Methods

Thirteen women and ten 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 (eleven completed subjects, ages 18-45 yrs) completed all study procedures and were included in analyses. Of the twelve excluded subjects, seven did not complete pre-scan training, one subject developed an MRI contraindication after consent, one showed a persistent response bias (i.e., indicating “blue” on more than 50% of color training runs), two subjects demonstrated

excessive head motion (>3mm) during MRI scanning, and one subject fell asleep during scanning. Subjects had normal neural anatomy as assessed by a neurologist (A.S.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-hr training sessions to minimize learning effects during MRI scanning. Subjects then performed the task inside the fMRI scanner for five 1.5-hr scan sessions, each consisting of six runs of 22 trials, for a total of 660 trials across the five days of testing. Each session occurred between 1 and 7 days after the previous session. The average number of days between first and last scan session across all subjects was 13.4, with a median of 10 days.

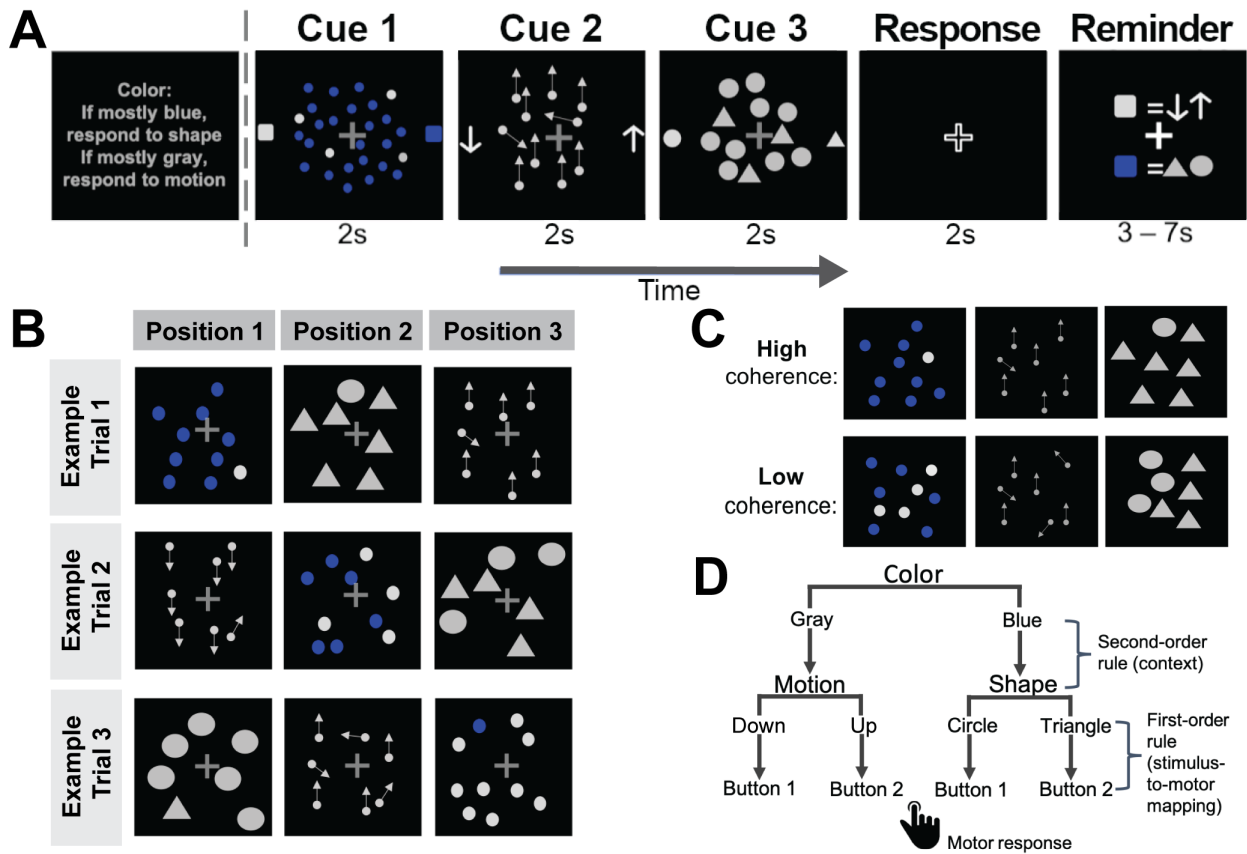


Figure 1. A. Task Design. At the start of each run, an instruction screen notified subjects of the active second-order rule. Each trial then consisted of three sequentially-presented stimuli accompanied by first-order rule reminders that were displayed as symbols on the left and right of the screen (to 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 inter-trial interval. On this example trial, the second-order cue, the color stimulus, indicates that the shape stimulus is the relevant first-order cue, and the motion stimulus is the irrelevant cue. 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 pseudo-randomly presented. **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.

Experimental design and statistical analyses

Task Design.

Subjects performed an abstract decision-making task (Figure 1A) in which each trial consisted of three sequentially presented visual stimulus displays (color, motion, and shape stimuli, respectively) shown in randomized, counterbalanced order (Figure 1B), followed by an explicit response period. The sensory salience of each visual stimulus was varied by manipulating its coherence level (Figure 1C). To address the role of policy abstraction (Badre & D'Esposito, 2009; M. M. Botvinick, 2008), the task was structured such that outcome of the perceptual 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 (Figure 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 (first-order 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. All runs were pseudo-randomized to ensure that position, coherence, and cue type were equally distributed across trials (Figure 1D).

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 – i.e., from specific second-order feature to first-order rule (Figure 1C) – 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, while a different subject's color rule might map blue to the motion stimulus. Likewise, the first-order 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 might map circles to the left button and triangles to the right button for one subject, but the opposite for another (Figure 1D).

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 (Figure 1A). On each trial, subjects were then required to perform two perceptual discriminations: 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 (Figure 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 seconds, pseudo-randomly presented

from a uniform distribution. To reduce working memory load, small symbols on the left and right sides of the screen outside of the circular stimulus display aperture reminded subjects of the first-order rule mapping from stimulus to button press for each of the stimuli (Figure 1A).

In order to dissociate motor execution from decision making, 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 seconds 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-sec duration, pseudo-randomly presented), subjects viewed a reminder screen about the currently active abstract rule to further reduce any working memory demand (Figure 1A). 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, Erickson, et al., 2010; Shankar & Kayser, 2017).

Behavioral training.

In training sessions 1 – 3, each subject completed a minimum of six behavioral runs per session in order 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 upward or downward moving dots, and 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 (H R Heekeren, Marrett, Ruff, Bandettini, & Ungerleider, 2006; Kayser, Buchsbaum, et al., 2010; Kim & Shadlen, 1999). For shape and color stimuli, coherence for the target feature was defined as the difference between the number of target features present and the number of distractor features, divided by their sum. For example, if 60 blue 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 (0%, 2%, 4%, 8%, 16%, 32%, 64% for the motion stimulus; 0%, 4%, 7%, 11%, 18%, 36%, 68% for the color stimulus; and 0%, 7%, 13%, 20%, 33%, 67%, 87% for shape (Kayser, Erickson, et al., 2010)). 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 $a = 1$ for 100% maximal accuracy at a given coherence level and $b = 0.5$ for the approximated 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 0.15s duration, while an incorrect trial was communicated with a high frequency tone. When subjects transitioned 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, in order 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 4 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 UC 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 = 1.8s, TE = 23ms, FOV 225mm, matrix size 70 x 70 x 32, voxel size 3.2 x 3.2 x 3mm). Anatomical images were acquired with a T1-weighted MP-RAGE imaging protocol (160 slices, TR=2.3s, TE = 2.98ms, FOV 256mm, matrix size 256 x 256, voxel size 1 x 1 x 1mm). 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 degrees. An MRI-compatible fiberoptic 4-button response device (Cambridge Research Systems, Inline model HH-1x4-L) was used for subject responses.

fMRI preprocessing.

fMRI data were pre-processed using a version of the open-source pipeline written by one of the authors (<https://github.com/savannahcookson/AFNI-Pipeline>) that was customized to support multi-session data. This pipeline supports preprocessing and first-level analysis of data based on the `afni_proc.py` function available through the Analysis of Functional NeuroImages (AFNI) software package (Cox J.S., 1996). DICOM 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 anatomical 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 x 3 x 3mm space, and then aligned to the skull-stripped anatomical image warped to the standard MNI atlas (Montreal Neuroscience Institute ICBM152, <https://www.mcgill.ca/bic/>). Data were smoothed with a 6-mm FWHM Gaussian kernel and scaled to have a voxel-wise mean of 100 by dividing the timeseries of each voxel by its mean signal and multiplying by 100. Individual volumes that contained more than 10% outlier voxels were censored from further analysis. Data

for each subject were then concatenated across all scanning days for subsequent first level analysis.

Univariate analysis.

All univariate analyses were completed with AFNI software through a combination of the same preprocessing pipeline above 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. Each regressor was created from the convolution of a gamma probability density function that peaked at 6-sec with a condition-specific vector of stimulus onsets and durations (2-sec for stimulus presentation, 6-sec for the instruction screen, and a stick function for the button press response). Resultant estimated beta coefficients for each subject were mapped to the MNI template for use in group-level analyses. Whole-brain voxelwise contrasts were next created for the second-order cue versus first-order cue (Second-order v First-order), first-order cue versus irrelevant cue (First-order v Irrelevant), all high coherence cues versus all low coherence cues (High v Low), and the timepoint of button press versus timepoints during task stimulus presentation (Button press v Visual stimuli). Stimulus type (color, motion, or shape) of each cue was intentionally collapsed to capture parietal activity that is non-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).

ROI selection.

All regions of interest (ROIs) were selected from a 1000-area parcellation (Schaefer et al., 2018) derived from a well-established 7-network resting-state fMRI (rs-fMRI) analysis (Yeo et al., 2011) that was obtained at 1mm resolution and resampled to 3 x 3 x 3mm, the same spatial resolution as the task data (https://github.com/ThomasYeoLab/CBIG/tree/master/stable_projects/brain_parcellation/Schaefer2018_LocalGlobal). The Schaefer et al parcel that included MNI coordinate [-30, -10, 68] in dorsal premotor cortex identified by Badre & D’Esposito (2007) was selected as the PMd ROI (parcel #144, MNI centroid [-28,-12,60]). Likewise, the parcel that included MNI coordinate [-38, 10, 34] in pre-premotor cortex from Badre & Deposito 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 rsfMRI) was then used to identify parietal parcels with maximal functional connectivity to pre-PMd and PMd

parcels, respectively. These 100 HCP subjects' data 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 timeseries averages were obtained by fitting HCP resting state data to the Schaefer et al 1000-area parcellation using MATLAB and AFNI. Pearson's correlations between all parcels were z-scored using Fisher's R-to-Z transform via AFNI's 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 [-32,-70,50], parcel #321), hereafter referred to as "caudal IPS", or cIPS; and a more rostral IPS subregion with greatest connectivity to PMd (MNI centroid [-38,-42,56], parcel #211), hereafter referred to as "rostral IPS", or rIPS.

To determine whether these parietal ROIs had differential connectivity to pre-PMd compared to 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 $p = 0.025$.

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 standard deviation 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" MNI centroid coordinates [-42 -40,42], "IP Sb" [-52,-34, 44], "IP Sc" [-32,-48,46], "IP Sd" [-48,-42,50], and "IP Se" [-38,-52,58]); for further ROI analysis, see Extended data table 1. We evaluated whether these IPS subregions with relatively nonspecific connectivity had frontal connectivity profiles distinct from the selectively connected rIPS and cIPS using a 1-way ANOVA and post-hoc Tukey HSD multiple comparisons test. A vector of each parietal ROI's correlation 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.

fMRI-based ROI analysis.

After selecting the parcel used for each ROI, an ROI mask was created by resampling from the Schaefer et al. 1000-area map. The mask was then applied to each subject's univariate whole-brain contrast dataset (Second-order v First-Order, First-

order v Irrelevant, High v Low, and Button press v Visual stimuli). We then took the mean beta value from all voxels within the selected parcel for each subject and performed ANOVAs and paired t-tests at the group level using custom MATLAB scripts.

Graded activity for second-order abstraction in IPS.

Lastly, we conducted an exploratory analysis to investigate whether selectivity for second-order abstraction varied systematically across the multiple subregions of IPS defined above. To do so, we initially computed the Pearson correlation coefficient between the mean beta value obtained from the Second-order v 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 v Irrelevant contrast, then calculated the significance of the difference between the Second-order v First-order correlation and the First-order v Irrelevant correlation.




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 represent between-subject SEM. Asterisks indicate $p < 0.05$.

Results

Behavior.

Subjects completed a task in which factors including the level (order) of abstraction, type of visual stimulus, sensory salience, and cue position were manipulated (Figure 1). As expected, subject performance improved with increased sensory salience: accuracy was higher for pertinent high coherence compared to low coherence cues except when the cue was irrelevant (second-order cue at high coherence versus second-order cue at low coherence, $t(10) = 7.96$, $p = 1.00 \times 10^{-5}$; first-order cue at high coherence versus first-order cue at low coherence, $t(10) = 13.99$, $p = 7.01 \times 10^{-8}$; irrelevant cue at high coherence versus irrelevant cue at low coherence, $t(10) = 0.64$, $p = 0.54$ (*ns*); Figure 2). This result confirmed that subjects attended to appropriate task-relevant stimuli and were behaviorally affected by 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 comprised it, we compared accuracy for trials in which motion, color, and shape stimuli, respectively, served as the second-order cue. No differences were seen ($F(2,30) = 0.16$, $p = 0.85$ (*ns*)), suggesting

that abstract decisions were not dependent on low level visual feature modality. As all cue types were presented equally often in positions 1, 2, and 3 across all trials, we next confirmed that position of the second-order cue did not affect accuracy ($F(2,30) = 1.53$, $p = 0.23$ (*ns*)).

Furthermore, accuracy did not differ significantly between trials for the six corresponding permutations of position for the second-order, first-order, and irrelevant cues ($F(5,60) = 0.81$, $p = 0.55$ (*ns*)). 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.




Figure 3. Univariate whole brain analysis depicting the positive main effect of task 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 Human Connectome Project (HCP) data. Abbreviations: pre-PMd, pre-premotor cortex; PMd, dorsal premotor cortex; rIPS, rostral intraparietal sulcus; cIPS, caudal intraparietal sulcus.

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 (Figure 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, in order 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 Materials and Methods; Table 1), we identified the parietal parcels with the highest positive raw z-scored correlation for pre-PMd (parcel 321 in caudal IPS (cIPS)) and PMd (parcel 211 in rostral IPS (rIPS); Figure 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: $t(99) = -20.57$, $p = 1.61 \times 10^{-37}$; cIPS: $t(99) = 14.80$, $p = 7.79 \times 10^{-27}$). 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).

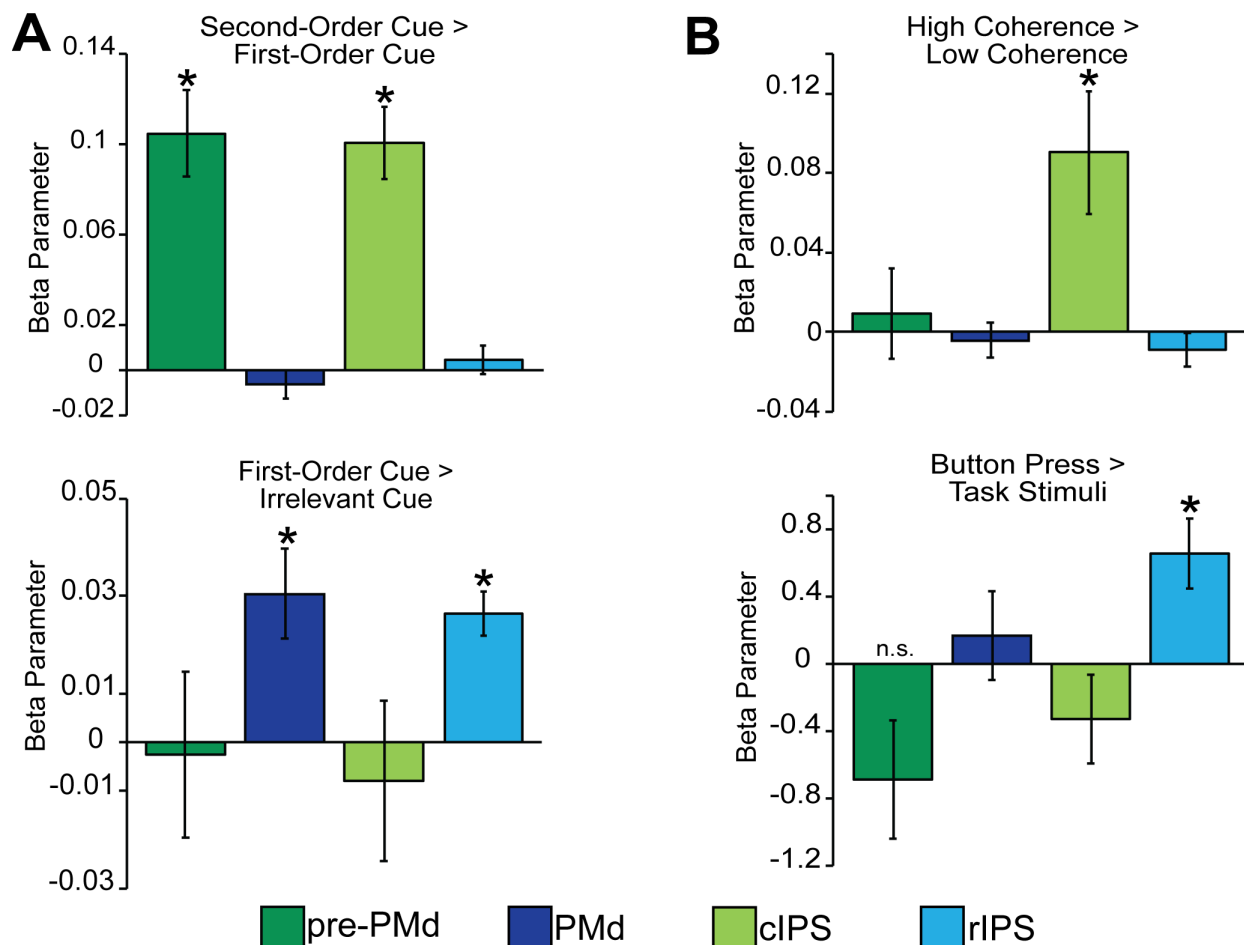


Figure 4. ROI-based results. **A.** Within-ROI activity for the univariate contrast of all second-order cues versus first-order cues (top) and all first-order cues versus irrelevant cues (bottom). **B.** Activity for the same ROIs in the contrast of all high coherence versus low coherence cues (top) and the timepoint of button press versus task stimulus displays (bottom). Beta parameter scaled to represent percent signal change (see Methods). Error bars represent between-subject SEM. Asterisks indicate $p < 0.05$; n.s. denotes non-significance. See also Figure 6: Extended data table.

Functional imaging ROI analysis.

Using these four ROIs – pre-PMd, PMd, cIPS and rIPS – we performed a 3-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 ($F(1.78, 17.81) = 32.27, p = 1.99 \times 10^{-6}$) and a significant interaction between ROI and coherence ($F(1.48, 14.76) = 7.21, p = 0.01$), though no 3-way interaction between ROI, abstraction, and coherence ($F(3.00, 30.00) = 0.359, p = 0.78$ (*ns*)).

The interaction between ROI and abstraction was driven by differences in the responses of pre-PMd and cIPS to second-order stimuli. Pre-PMd had significantly greater activity for the second-order versus the first-order cue (Second-order v First-order, $t(10) = 5.45, p = 2.81 \times 10^{-4}$), as did cIPS ($t(10) = 6.29, p = 9.07 \times 10^{-5}$), but neither PMd ($t(10) = -1.00, p = 0.339$) nor rIPS ($t(10) = 0.71, p = 0.493$) showed such changes (Figure 4A, upper panel). 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 (Figure 4A, lower panel). Specifically, PMd and rIPS strongly differentiated relevant from irrelevant first-order cues (First-order v Irrelevant, PMd: $t(10) = 3.26, p = 0.009$; rIPS: $t(10) = 5.77, p = 1.80 \times 10^{-4}$), but pre-PMd ($t(10) = -0.15, p = 0.880$ (*ns*)) and cIPS ($t(10) = -0.48, p = 0.641$ (*ns*)) 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 second-order policy abstraction (Badre & D'Esposito, 2007; Badre & Nee, 2018), these data demonstrate that PMd and rIPS are dissociably responsive to attentional manipulations for first-order 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 (Figure 4B, upper panel). When high and low coherence stimuli were compared, collapsed across order of abstraction, cIPS responded more strongly to high coherence (High v Low, $t(10) = 2.93, p = 0.015$). Notably, as indicated by the absence of a 3-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, $F(10,1) = 1.08, p = 0.324$ (*ns*)). To determine whether cIPS might therefore have a specific role in sensorimotor transformations (Erickson & 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 v Visual stimuli, $F(40,3) = 4.54, p = 0.008$). cIPS did not distinguish between the motor output and visual stimuli ($t(10) = -1.26, p = 0.237$ (*ns*)), but rIPS was significantly more active for the button press ($t(10) = 3.18, p = 0.0098$). Neither of the two prefrontal ROIs showed significantly different activity for this contrast (pre-PMd $t(10) = -1.96, p = 0.079$ (*ns*); PMd $t(10) = 0.64, p = 0.536$ (*ns*)). 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 sensory salience, while first-order rIPS was differentially active during motor output.

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 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 rs-fMRI data (see Materials and Methods). We identified five parcels: parcel 197 (IPSa, MNI centroid [-42,-40,42]), parcel 196 (IPSc [-52,-34, 44]), parcel 204 (IPSc [-32,-48,46]), parcel 326 (IPSc [-48,-42,50]), and parcel 212 (IPSe [-38,-52,58])(Extended data table 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 ($F(6,693) = 16.63, p = 5.712 \times 10^{-18}$; see Materials and Methods), allowing us to draw comparisons between subregions with less specific connectivity to pre-PMd/PMd and those with more selective connectivity.




Figure 5. The gradient of abstraction along the rostral-caudal axis in IPS. Pearson’s correlation of beta parameter estimate with MNI y-coordinate of seven IPS subregions defined by independent connectivity measures (see Methods). Correlation of second-order information (Second-order cue versus first-order cue contrast, circular symbols, black solid line) plotted with attentional control (First-order cue versus irrelevant cue, square symbols, black dashed line). A significant difference between the two correlations is represented by the solid gray bracket, left. The color of each datapoint corresponds to the color-coded ROI shown on the axial view of the brain diagram, right (Lime green, cIPS; Pink, IPSe; Purple, IPSc; Orange, IPSd; Bright blue, rIPS; Red, IPSa; Brown, IPSb). Asterisks indicate $p < 0.05$, n.s. not significant.

Confirming our predictions, these IPS subregions showed task activations that were both largely similar to each other and also dissimilar to those of rIPS and cIPS (Extended data table 1). A 3-way ANOVA with factors of ROI, coherence, and order produced an interaction of coherence and ROI ($F(6,60) = 6.92$, $p = 1.28 \times 10^{-5}$) 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 (IPSa-e non-significant results in Extended data table 1). With respect to attentional control of the first-order cues, only rIPS and IPSc ($t(10) = 2.42$, $p = 0.036$) had significantly greater activity for first-order relevant versus irrelevant cues. However, post-hoc analyses of the interaction of order and ROI ($F(25.36, 2.54) = 9.34$, $p = 0.004$) revealed that all five of the nonspecific connectivity subregions had significantly greater activity for the second-order cue compared to the first-order cue (Extended data table 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 rostral/caudal axis (see 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 second-order versus first-order stimuli might vary with rostrocaudal position given preliminary evidence from Choi, Drayna, and Badre (2018). We therefore correlated the anterior-posterior position of all seven IPS subregions (IPSa-e, cIPS, rIPS) with their respective beta values for the Second-order v First-order contrast (Figure 5). We found a significant correlation between the anterior-posterior position, as indexed by the MNI y-coordinate of the centroid, and beta value ($r(5) = -0.7671$, $p = 0.044$; Figure 5). Interestingly, the increase in selectivity for the second-order cue in more caudal IPS subregions was opposite to that of frontal cortex, where more rostral regions show greater responsiveness to more abstract stimuli. Notably, if we instead correlated anterior-posterior position with strength of resting state functional connectivity to pre-PMd ($r(5) = -0.4694$, $p = 0.288$ (*ns*)) or to PMd ($r(5) = 0.670$, $p = 0.100$ (*ns*)), the correlations were not significant.

To test whether attentional control of the relevant first-order cue was similarly represented in spatially-graded activity, we next correlated beta values of the First-order v Irrelevant dataset with the MNI y-coordinate of the seven IPS subregions (IPSa-e, rIPS, cIPS). We found no correlation between these variables ($r(5) = 0.549$, $p = 0.202$, (*ns*)). However, a comparison of the Second-order v First-order spatial relationship and the First-order v Irrelevant spatial relationship revealed that these two correlations were significantly different ($z = -2.31$, $p = 0.010$; Figure 5). These results suggest that a representation of higher-order abstract information within a spatial gradient may be a topographic feature of IPS.

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 that is separable from attentional control, sensory features, and motor responding. These findings include dissociations between specific regions and their representations of different cognitive processes: activity within pre-PMd and cIPS is modulated by hierarchical order, for example, whereas activity in PMd and rIPS is strongly modulated by attentional control of relevant stimulus-response relationships. Notably, representation of higher abstraction in IPS exists in a gradient from rostral to caudal, opposite the caudal to rostral organization of increasing abstraction seen in lateral frontal cortex.

The above results follow previous work on policy abstraction establishing that parietal cortical activity may vary with increasing abstraction. Nee & D'Esposito (2016) demonstrated that parietal cortex is strongly active across different levels of control demands (feature, contextual, and temporal control, from less to more abstract), though they did not evaluate whether such activity in parietal cortex varied parametrically with abstraction. More recently, in a pioneering rs-fMRI study by Choi and colleagues (Choi et al., 2018) that further analyzed previously published task data (Badre & D'Esposito, 2007), lateral frontal regions whose activity correlated with progressively more abstract decisions in caudal-to-rostral fashion were shown to participate in intrinsic networks whose parietal components were consistent with a rostral-to-caudal gradient. These authors also found that networks in parietal cortex responded differentially to the order of policy abstraction. However, their network analysis did not investigate more granular parietal activations, which we accomplished using a multimodal parcellation to yield specific subdivisions of IPS. We also linked fronto-parietal connectivity directly to task data, in finding that cIPS could be identified by strongest connectivity to pre-PMd and also by strongest activity for the second- versus first-order cue in a ranking of all parietal parcels (not reviewed in results). Furthermore, their task was designed to understand policy abstraction while controlling for confounds critical in frontal cortex, such as the number of response mappings. Here we build upon this work to explicitly control for processes represented in parietal cortex that could confound interpretation: for example, by equating the featural complexity of our visual stimuli across orders of policy abstraction, including explicit manipulations of visual attention and sensory salience, and isolating motor responses from decision timepoints.

In keeping with the potential importance of specific confounds for parietal cortex, these data demonstrate that the functions of frontal and parietal areas within intrinsic networks should not be viewed as homogeneous (Figure 4). Instead, these areas also appear to make unique contributions to decision making. Here we find that cIPS (and not pre-PMd) responds strongly to coherence differences, while rIPS (and not PMd) is sensitive to the motor response. The former result is in keeping with findings that a region within right inferior parietal lobule correlates with perceptual capacity, while an area within left middle frontal gyrus correlates with cognitive control capacity (Eayrs & Lavie, 2019). These findings are also consistent with previous theoretical and empirical work that has identified other significant frontal/parietal dissociations. For example, parietal cortex, which is well situated to serve as an episodic buffer in working memory (Gelastopoulos, Whittington, & Kopell, 2019), might have a role in the manipulation but not the monitoring of items within working memory, unlike lateral frontal cortex (Chamod & Petrides, 2010). Likewise, in a spatial working memory task performed in

patients with either frontal or parietal lesions, frontal lesions reduced resistance to distraction, while parietal lesions impaired the ability to benefit from predictable spatial positions of targets (Saj, Verdon, Hauert, & Vuilleumier, 2018). Frontal and parietal cortical-spatial relationships are further distinguished by our results. An analogous graded selectivity for second-order abstraction that exists in the cortical space between PMd and pre-PMd has not been discovered in frontal studies, despite observing this spatial relationship between caudal and rostral IPS subregions (Figure 5).

Less consistent with specializations within parietal cortex is the distinction— i.e., between cIPS and rIPS – in the response to coherence. Behaviorally, subjects performed better when both second- and first-order cues were high rather than low coherence (Figure 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 responding in parietal cortex (Xu, 2018), and IPS in particular. The location of cIPS in this study (MNI centroid -32, -70, 50) closely approximates the location of IPS1 (Bray, Arnold, Iaria, & MacQueen, 2013), a region previously shown to contain a retinotopic map of spatial attention (Silver & Kastner, 2009), among other functions (Kastner, S., Chen, Q., Jeong S.K., Mruczek, 2017; Xu, 2018). That cIPS responds differentially to high and low coherence across orders may therefore reflect its responsiveness to differences in stimulus salience, independent of its distinction between orders of policy abstraction. The spatial conjunction of this response to both perceptual and higher-order features might enable cIPS, in particular, to access both cues as needed, though we did not detect a corresponding change in univariate activity reflecting this conjunction.

More generally, these findings build upon previous work demonstrating that IPS responds to stimuli that are more abstract by definitions other than the rule complexity of policy abstraction. Left IPS has been shown to respond in load-dependent fashion for both visual and verbal stimuli, thereby generalizing across sensory modality (Cowan et al., 2011). Moreover, when working memory load and context binding are both manipulated, IPS activity during the delay period more strongly distinguishes context than load, consistent with accounts in which IPS tracks salient features of the task (Gosseries et al., 2018). Similarly, when objects must be either categorized as part of a group or identified as a specific exemplar, IPS connectivity with object-sensitive regions in temporal cortex distinguishes the task conditions (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, such 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 and colleagues found that IPS represented response complexity but not instructions (Palenciano, González-García, Arco, Pessoa, & Ruz, 2019). Similarly, during a switch task, TMS 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, Andres, Brass, & Andres, 2014). Together these different findings are consistent with a view in which the posterior parietal cortex, including the IPS, might more broadly represent “action-independent, nonspatial visual information” (Xu, 2018).

Related to this question of what IPS encodes is the problem of where IPS encodes it. In seeking to define the broader spatial organization of the parietal cortex, and of IPS specifically, previous work has identified topographically organized IPS subregions, including attentional maps with systematic representations of retinotopic space (Silver, Ress, & Heeger, 2005) that extend along the sulcus (Silver & Kastner, 2009). This organization supports theories that IPS includes “priority maps” corresponding to the attentional weighting of a spatial location (Jerde & Curtis, 2013), such that lesions in these areas disrupt saccades to remembered locations (Mackey, Devinsky, Doyle, Golfinos, & Curtis, 2016). In contrast, more anterior-superior regions of IPS may respond more to features than to locations (Bettencourt & Xu, 2016), with the most anterior regions integrating feature and action to enable tool use (Kastner, S., Chen, Q., Jeong S.K., Mruczek, 2017). Here we confirm an additional localized rostral to caudal gradient in the representation of policy abstraction, mirroring that of the lateral frontal cortex. Yet the exact relationship between intrinsic connectivity and task representations within IPS remains unclear, as frontal connectivity strength did not correlate with rostral-caudal organization in the same manner that selectivity for higher task abstraction did in this study. How abstract representations are organized with respect to other feature responses in the previously-defined subdivisions of IPS (e.g., retinotopic IPS0-5) remains an additional mystery.

Altogether these findings demonstrate that the IPS contains topographic information related to policy abstraction, both distinct from representations of attentional control, sensory features, and motor responses, and organized into specific networks identifiable by their frontal connectivity. This work builds upon 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 & Kayser, 2013; Kastner, S., Chen, Q., Jeong S.K., Mruczek, 2017; Xu, 2018). While both frontal and parietal areas are generally co-active in cognitive control tasks such as the one implemented here, our results emphasize that their functions are not identical, but complementary. Building on this and previous research, future work should thus continue to define how the representations and connectivity of IPS contribute to human performance of abstract, hierarchically-organized tasks.

Supplemental Materials

See Figure 6 data table on the following page.

ROI	Parcel num	Centroid	7-Network Membership	pre-PMd Connectivity	PMd Connectivity	Num voxels	Second-order v First-order	First-order v Irrelevant	Button press v Task stimuli	High v Low
rIPS	211	-38,-42,56	Dorsal Attention, Posterior Parietal	-0.085	2.44	26	t(10) = 0.7124, p = 0.492 (n.s.)	t(10) = 5.7697, p = 0.0002*	t(10) = 3.1826, p = 0.009*	t(10) = -1.0251, p = 0.329 (n.s.)
cIPS	321	-32,-70,50	Control, Parietal	1.479	-0.73	44	t(10) = 6.2865, p = 9.066 x 10 ⁻⁵ *	t(10) = -0.4807, p = 0.641 (n.s.)	t(10) = -1.2564, p = 0.238 (n.s.)	t(10) = 2.9326, p = 0.015*
PMd	144	-28,-12,60	Somato-Motor	-0.17	n/a	55	t(10) = -1.0034, p = 0.339 (n.s.)	t(10) = 3.2614, p = 0.009*	t(10) = 0.6408, p = 0.536 (n.s.)	t(10) = -0.4584, p = 0.657 (n.s.)
pre-PMd	355	-40,10,32	Control, Lateral Prefrontal Cortex	n/a	-0.373	15	t(10) = 5.4495, p = 0.0003*	t(10) = -0.1549, p = 0.880 (n.s.)	t(10) = -1.9597, p = 0.079 (n.s.)	t(10) = 0.4119, p = 0.689 (n.s.)
IPSa	197	-42,-40,42	Dorsal Attention, Posterior Parietal	0.781	1.193	40	t(10) = 3.0804, p = 0.012*	t(10) = 1.3662, p = 0.202 (n.s.)	t(10) = -0.4524, p = 0.285 (n.s.)	t(10) = -0.5176, p = 0.616 (n.s.)
IP Sb	196	-52,-34,44	Dorsal Attention, Posterior Parietal	0.729	1.074	41	t(10) = 2.4135, p = 0.036*	t(10) = 2.0209, p = 0.071 (n.s.)	t(10) = 1.3244, p = 0.215 (n.s.)	t(10) = 0.8898, p = 0.395 (n.s.)
IP Sc	204	-32,-48,46	Dorsal Attention, Posterior Parietal	0.7	0.716	28	t(10) = 4.3708, p = 0.001*	t(10) = 2.4224, p = 0.036*	t(10) = -0.9575, p = 0.361 (n.s.)	t(10) = 0.1888, p = 0.854 (n.s.)
IP Sd	326	-48,-42,50	Control, Parietal	1.272	0.475	31	t(10) = 3.1437, p = 0.010*	t(10) = 1.5558, p = 0.151 (n.s.)	t(10) = 0.9303, p = 0.374 (n.s.)	t(10) = 0.6846, p = 0.509 (n.s.)
IP Se	212	-38,-52,58	Dorsal Attention, Posterior Parietal	0.468	1.352	49	t(10) = 3.5282, p = 0.005*	t(10) = 1.8582, p = 0.093 (n.s.)	t(10) = -0.1591, p = 0.877 (n.s.)	t(10) = 0.6038, p = 0.559 (n.s.)

Figure 6. Extended data table: Comparative results from IPSa-e. “Parcel num” references the parcel identifier number from the Schaefer et al (2018) 1000-area, 7-network parcellation used as the ROI. Centroids are provided in MNI coordinates. “7-Network membership” refers to the Yeo et al (2011) network that the parcel has been previously classified into. Connectivity values are raw z-scored correlations that have been standardized within frontal region (pre-PMd, or PMd) for comparison (see Materials and Methods). “Num voxels” is the number of voxels within each parcel used as the ROI for further testing after resampling to task data space (3 x 3 x 3mm). Results from paired t-tests of each ROI are reported with significant results denoted by an asterisk at p < 0.05 and non-significant results represented by (n.s.).

Chapter 3: Does activity in parietal areas involved in decision-making vary as a function of strategy?

Abstract

One can perform the same goal-oriented task using different strategies (Morrison et al., 2016). Studies have shown dissimilarities in neural activity for strategies that depend on unique sensory or mnemonic systems (Iaria et al., 2003), but few have investigated if different strategies in performing a cognitive control task produce activation differences within the fronto-parietal control network if those strategies differ in the form of abstraction. An abstract decision-making task where half of all trials required use of a hierarchically-defined “policy strategy” and the other half permitted a heterarchical match/nonmatch “comparison strategy” (Badre, 2008) was used to investigate whether frontal and parietal control regions implicated in policy abstraction differentially represent abstraction strategy. To explore this question, previously collected fMRI data from healthy female and male subjects performing a second-order policy abstraction task was re-analyzed. Accuracy was higher for trials allowing the comparison strategy, and a significant behavioral interaction was observed between sensory coherence of the second-order cue and strategy trial type. Across the brain, known decision-making regions for hierarchical policy abstraction tasks had increased activity for the policy strategy, including frontal regions pre-PMd and PMd. A region in the superior frontal sulcus covaried activity with individual differences in subjects’ strategy use. However, parietal activity did not strongly differentiate between the strategies, suggesting that IPS subregions specifically important for policy abstraction decision information may not preferentially represent control of hierarchical policy abstraction over other non-hierarchical types of task abstraction.

Introduction

The same task can be completed using different strategies. For example, a common cognitive task that investigates short term memory called the Sternberg task (Sternberg, 1966) can be completed using either a visuo-spatial strategy or a verbal strategy (Sanfratello et al., 2014). Previous research has further observed that the same task can produce different neural correlates depending on the strategy employed. This is seen in the finding that reading out loud can recruit a phonological pathway dependent on left PMd, or a semantic pathway that engages left ventral inferior frontal cortex, depending on whether a verbal or mnemonic strategy is used, respectively (Kherif et al., 2009). Additional distinctions between mental imagery strategies that engage the precuneus (Cavanna & Trimble, 2006), visuo-spatial methods dependent on scene-selective visual areas (Ramanoël, Durteste, Bécu, Habas, & Arleo, 2020), and navigation methods that recruit hippocampal versus striatal systems (Iaria et al., 2003) have been delineated across various tasks, but differentiations of the frontoparietal network in strategies for cognitive control tasks have not been as thoroughly studied.

One possibility is that frontoparietal control regions are entirely invariant to different strategies of performing an abstract decision-making task as long as they are equivalently difficult, and in fact, many studies assert this as their defining characteristic (Cole et al., 2013; Crittenden & Duncan, 2014; Fedorenko et al., 2013). However, if

different strategies elicit different types of cognitive control demands, then they may produce differential activation within hierarchically-organized control regions. This is seen in other studies of frontal cortex that have found evidence that control regions are selective to the type of demand required, producing differential activation in caudal and rostral frontal regions for task aspects such as stimulus-domain sensitivity and temporal control over action (Nee & D'Esposito, 2016, 2017; Ranti, Chatham, & Badre, 2015). The frontal and parietal literature further disagrees about the relative impact that the frontal versus parietal lobes have in cognitive control relating to the flexible use of different strategies, with some studies asserting that posterior parietal cortex mediates cognitive control task set switching across perceptual, rule-based, and mnemonic domains (Chiu, Esterman, Gmeindl, & Yantis, 2012) and many other studies claiming that this function is unique to the frontal lobe (Brass, Ullsperger, Knoesche, Cramon, & Phillips, 2005).

Unlike previous studies, the decision-making task developed for this dissertation generates specific hypotheses about which regions within lateral prefrontal and parietal cortices should covary with a particular strategy of performance. Here I build upon Chapter 2 of this dissertation where connectivity-defined regions of parietal cortex, specifically subregions within IPS, showed activity reflecting selectivity for varying orders of policy abstraction using a hierarchical decision-making task. The precise nature of the involvement of cIPS and rIPS for the employment of task strategies dependent on policy abstraction explicitly, however, was left unknown. Due to the unique, counter-balanced design of the task developed for this thesis, the outcome of the two first-order cues' sensory discriminations on half of all trials permitted usage of an alternative strategy that did not require use of the second-order policy. As a result, the same neural data could be divided and re-analyzed for correct trials in which second-order policy use was absolutely necessary for performance ("policy strategy") and trials in which subjects could compare the two first-order rules ("comparison strategy") to arrive at the correct response.

Generally speaking, if there are differences in neural representations of strategy within the same cognitive control task, then a policy strategy that relies on the use of hierarchical second-order cues is likely to engage brain networks implicated in hierarchical control during policy abstraction (Badre & D'Esposito, 2007; Badre et al., 2010) more strongly than an alternate strategy that utilizes a different form of abstraction. The comparison strategy in this study uses a match/nonmatch rule that is similar to relational abstraction, another type of task abstraction that is heterarchical rather than hierarchical in design (Badre, 2008; Badre & D'Esposito, 2009; Badre & Nee, 2018). Previous work has attempted to define differences in neural activation between types of task abstraction, but there has been no consensus (Badre & Nee, 2018; Bahlmann, Blumenfeld, & D'Esposito, 2015; Nee et al., 2014).

Given that the previously identified task-active subnetwork comprised of pre-PMd and cIPS preferentially processed second-order policy abstraction information, cIPS activity may preferentially represent the policy strategy, and this preference would be apparent at the timepoint of the second-order cue on a trial during policy use. However, if the role of cIPS is to process higher-order abstract information regardless of the type of abstraction, then activity would not differentiate between the policy and comparison strategies. This latter hypothesis would hold true if subjects encoded the second-order

cue as task-relevant information in both of the studied task strategies, even if it was not primarily used for selecting a response on comparison strategy trials. With respects to rIPS, a parietal subregion that formerly showed activation for relevant first-order cues and action initiation may not differentiate between a policy strategy and a comparison strategy. This hypothesis would apply if both strategies required equivalent activation for attentional control over relevant first-order information as well as motor output for a single button press (see Chapter 2).

Outside of parietal cortex, frontal regions strongly implicated in policy abstraction are likely to exhibit increased activity for the policy strategy in accordance with the associated hierarchical level of policy abstraction (Badre & D’Esposito, 2007). Thus, I predicted that pre-PMd would have stronger activation on policy strategy trials versus comparison strategy trials, and this strategy dissociation would be particularly apparent during second-order cue presentation but not during the first-order cue. I further predicted that PMd would show increased activity for the policy strategy overall and explicitly during first-order cue presentation, as a region implicated in response competition (Badre & D’Esposito, 2007). This PMd hypothesis was derived from the fact that the policy strategy was necessitated on trials where the comparison strategy was insufficient to resolve response-level competition between nonmatching button presses (Badre & D’Esposito, 2007; Silvia A. Bunge et al., 2002; O’Shea, Sebastian, Boorman, Johansen-Berg, & Rushworth, 2007). As a result, policy strategy trials may require increased activity in PMd for selection between competing motor options (“different-button” trials, see Figure 7) that did not occur on comparison strategy trials (“same-button” trials).

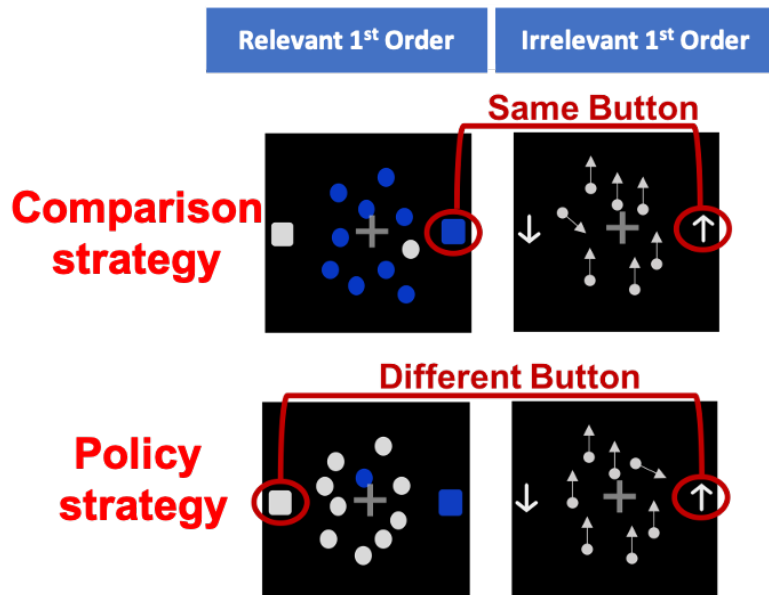


Figure 7. Two example trials showing the relevant and irrelevant first-order cues (second-order cue and response display not pictured). When the outcome of both first-order cue perceptual discriminations refer to the same button (SB), then the comparison strategy was permitted

Results address each of these hypotheses, overall lending evidence to a theory of cognitive control where pre-PMd may specialize in strategies demanding hierarchical policies. Similarly, results in PMd show trend-level specificity for a policy strategy. Negative results in parietal cortex may further suggest that intrinsically-connected IPS

subregions pertinent for task performance may represent policy abstraction strategies in addition to other types of abstraction strategies.

Materials and Methods

All task data collection and methods up to and including “*fMRI pre-processing*” for this chapter were identical to Chapter 2, as re-analyzed data (see Chapter 2 Materials and Methods). Additions to analysis and novel methods are described as follows.

Behavioral measure of strategy use.

On each trial of the abstract decision-making task used in this study, one stimulus served as the second-order cue and two other stimuli served as first-order cues that the subject used for responding by pressing either button 1 (left button) or button 2 (right button) on a button pad. Subjects were instructed to select one of the first-order cues as relevant by using the second-order cue and to ignore the remaining first-order cue as irrelevant, but if the perceptual judgment for both of the first-order cues indicated to press the same button (SB trial, e.g., both cued to press button 2), then the subject could utilize a strategy in which they consistently compared the irrelevant and relevant first-order cues (comparison strategy) to reach the correct answer. This meant that all SB trials implicitly permitted the comparison strategy. Conversely, on trials where the perceptual judgement for the first-order cues indicated to press different buttons (DB trial, e.g., one first-order cue indicated to press button 1 while the other indicated to press button 2), then the strategy that required use of the second-order cue (policy strategy) was absolutely necessary to employ in order to reach a correct answer on the trial. Thus, all DB trials necessitated the policy strategy. Because the task design was balanced to have equal probability of button 1 and button 2 as correct answers, half of all trials had first-order cues indicating to press the same button (SB trials) and the other half of all trials had first-order cues indicating to press different buttons (DB trials) (Figure 7). This allowed retention of all correct trials from the 660 total trials for this re-analysis of data.

Each subject’s behavioral measure of policy strategy use was assessed as the difference in accuracy on DB trials compared to SB trials: $DB_{accuracy} - SB_{accuracy}$. It can be inferred that subjects who exhibited a larger accuracy difference on SB trials were likely using the comparison strategy more frequently. Therefore, as an inference, subjects who rarely used the comparison strategy and therefore used the policy strategy more frequently on both trial types had a value closer to zero; likewise, subjects who frequently used the comparison strategy and therefore used the policy strategy comparatively less overall had increased accuracy on SB trials and thus more negative values on this behavioral measure.

Univariate analysis.

The AFNI function 3dDeconvolve was used to create a general linear model that evaluated 19 different conditions as separate regressors. All regressors evaluated correct trials only, except for one regressor of lesser interest which included all three stimulus displays on error trials. The 18 regressors of interest comprised all permutations of a single cue type (the second-order, first-order, or irrelevant cue),

position (1, 2, or 3 within a trial), and response trial type (same button (SB) or different button (DB) responses of the relevant and irrelevant first-order cues). Thus, the first six regressors included the following stimulus presentation conditions: (1) second-order cue in position 1 on SB trials, (2) second-order cue in position 1 on DB trials, (3) second-order cue in position 2 on SB trials, (4) second-order cue in position 2 on DB trials, (5) second-order cue in position 3 on SB trials, and (6) second-order cue in position 3 on DB trials. Regressor (7) started with the first-order cue in position 1 on SB trials, and so on for the remaining 11 permutations. It is important to emphasize that the same data was not reused across regressors, as regressors only took into consideration brain activity during a single cue on a trial. However, the same trial was used across regressors, seen, for example, on a Second-First-Irrelevant-SB trial, where regressor (1) included only activity during the second-order cue on that trial subtype, while regressor (9) included only activity during the first-order cue on that same trial subtype. Each regressor was subsequently created from the convolution of a gamma probability density function that peaked at 6-sec with a condition-specific vector of stimulus onsets and durations of 2-sec for stimulus presentation. Resultant estimated beta coefficients for each subject were mapped to the MNI template for use in group-level analyses. Whole-brain voxelwise contrasts were subsequently created for all cues on different button trials versus all cues on same button trials (DB v SB), second-order cue on different button trials versus second-order cue on same button trials (Second-order DB v Second-order SB), and the first-order cue on different button trials versus first-order cue on same button trials (First-order DB v First-order SB). To remove areas that deactivated during task performance, each contrast was masked by the positive main

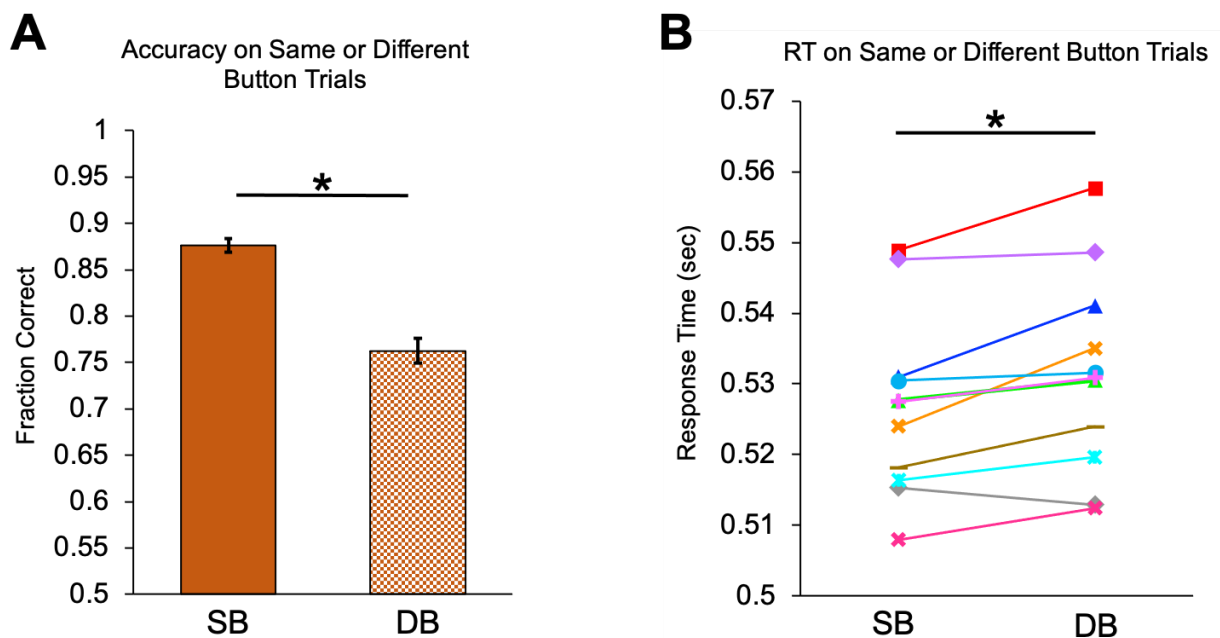


Figure 8. Behavioral effect of strategy. Performance on Same Button (SB) compared to Different Button (DB) trial types. **A)** Group-level effect of strategy trial type on accuracy. Error bars represent between-subject SEM. **B)** Individual subject response times plotted as unique color/symbol pairs compared across strategy trial type. Black asterisks indicate $p < 0.05$.

effect of task, created by contrasting all conditions versus baseline at a loose threshold ($p < 0.05$, uncorrected). All voxels in whole-brain group level maps were then thresholded at uncorrected $p < 0.001$, then cluster-corrected to produce whole-brain maps at $p < 0.05$ using AFNI function “3dClustSim”, with bi-sided voxel thresholding and second-nearest neighbor clustering settings ($nn = 2$).

fMRI-based ROI analysis.

The same four ROIs from Chapter 2 were used for this study (pre-PMD, PMd, cIPS, rIPS), selected from the Schaefer et al 1000-area parcellation resampled to task data voxel dimensions. The identical ROI mask from Chapter 2 was applied to each subject’s univariate whole-brain contrast dataset (DB v SB, Second-order DB v Second-order SB, First-order DB v First-order SB) (see Chapter 2 methods). The mean beta value from all voxels within the selected parcel were then computed for each subject and paired t-tests were performed at the group level using custom MATLAB scripts. For the four ROIS, significance was assessed using a Bonferroni type adjustment with significance threshold of $p = 0.0125$ for multiple comparisons correction.

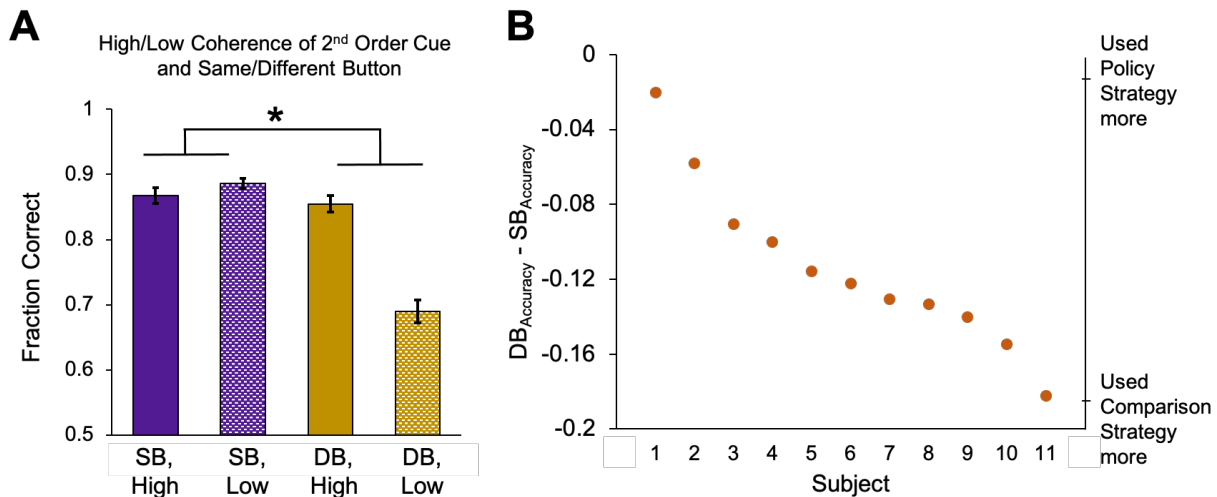


Figure 9. A) Coherence of the second-order cue mattered more when policy strategy is required on DB trials. Error bars represent between-subject SEM. Asterisk indicates $p < 0.05$ for interaction effect. **B)** Individual subjects’ behavioral measure of strategy use assessed as accuracy on DB trials subtracted from accuracy on SB trials. Subjects with less negative values used the Policy strategy more consistently during task performance, while subjects with more negative values used the Comparison strategy more frequently.

Behavioral covariate analysis.

In an exploratory test to discover which voxels the BOLD signal significantly covaried with use of the policy strategy, subjects’ behavioral measure of strategy use was used as a covariate factor in a separate whole-brain t-test of DB v SB at the group level. Individual voxels were thresholded at uncorrected $p < 0.001$, then cluster-corrected to produce whole-brain maps at $p < 0.05$.

Results

Behavior.

Subjects performed an abstract decision-making task that dissociated factors known to be implicated in various parietal cortex functions (Figure 1; see Chapter 2 Methods). Due to the equal probability of the correct response being button 1 or button 2, subjects could respond correctly on a trial using either a hierarchical policy strategy or a first-order cue comparison strategy. On trials where both of the first-order cues (relevant and irrelevant) referred to the same button press response (SB trial), subjects had increased accuracy compared to trials where first-order cues referred to different buttons (DB trial) ($t(10) = 8.34$, $p = 8.139 \times 10^{-6}$) (Figure 8A). This reflected behavioral facilitation from potential use of a comparison strategy that considers the outcome of both first-order cue discriminations. In addition to accuracy improvement, response time on SB trials was significantly faster ($t(10) = -3.60$, $p = 0.005$) (Figure 8B).


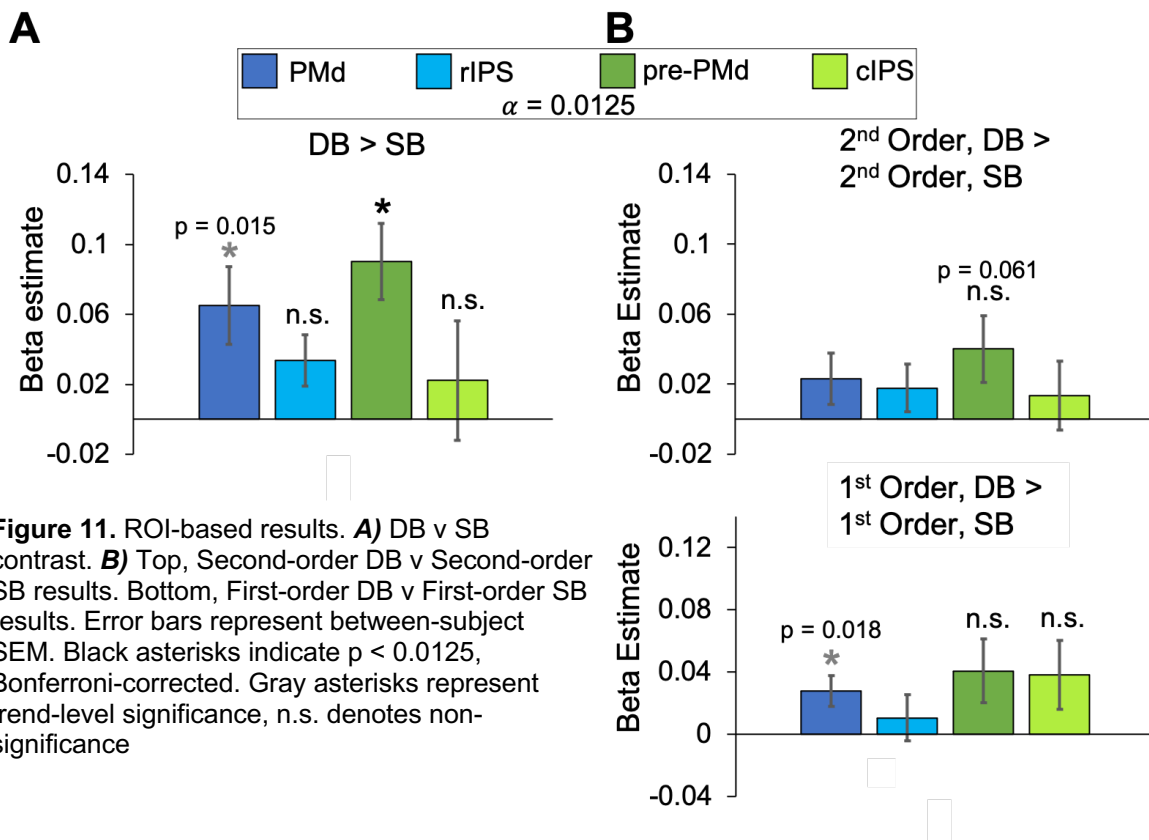


Figure 10. Group-level whole brain univariate results for contrast of DB v SB depicting activation in **A)** left frontal eye field, **B)** bilateral ventral premotor, **C)** bilateral pre-PMd, **D)** right intraparietal area/middle occipital gyrus, **E)** right prefrontal subdivision of thalamus, and **F)** right ventral occipital cortex. Results voxel-wise corrected at $p < 0.05$ after cluster correction of $p < 0.001$. Significant clusters overlaid to MNI standard anatomical.

While stimulus coherence of cue types was shown previously to affect accuracy (Figure 2; see Chapter 2 Results), here the effect of coherence was further demonstrated as dependent on the strategy employed. A significant interaction between

coherence level of the second-order cue (high, low) and SB/DB trial type was observed in a 2-way ANOVA ($F(1,43) = 50.67$, $p = 1.285 \times 10^{-8}$) (Figure 9). Subject accuracy was lowest on DB trials that required explicit use of the second-order cue at low coherence for the policy strategy. This served as further evidence of a strategy shift that occurred when subjects utilized the comparison strategy on SB trials.

A behavioral measure was next assessed for each subject that represented how often the subject utilized the comparison strategy based on their accuracy advantage on SB trials (see Methods)(Figure 9B). A less negative value meant that the subject rarely used the comparison strategy and thus depended on the policy strategy more frequently; a more negative value meant that the subject often used the comparison strategy.



Univariate functional imaging maps.

Group-level whole brain results show greater activation in known policy abstraction decision-making regions for the policy strategy during DB trials, as predicted (Figure 10). In the DB v SB contrast, increased activation was observed bilaterally in pre-PMd, suggesting consistent use of the second-order cue during policy strategy trials. Other neural activation was seen in bilateral ventral premotor cortex, right fusiform gyrus/ventral occipital cortex, right middle occipital gyrus/intraparietal area, right prefrontal subdivision of the thalamus, and left frontal eye field.

Functional imaging ROI analysis.

An ROI analysis was conducted using the four Schaefer et al (2018) parcels that were previously identified as important for policy abstraction in frontal and parietal cortex (pre-PMd, MNI centroid [-40,10,32]; PMd [-28,-12,60]; cIPS [-32,-70,50]; rIPS [-38,-42,56]; see Chapter 2 Methods). Univariate results generally support hypotheses in which frontal ROIs selectively represent a policy strategy compared to a match/nonmatch comparison strategy during abstract decision-making (Figure 11). In the DB v SB contrast, parietal regions cIPS and rIPS did not have greater activation for one strategy over the other after significance was assessed with Bonferroni correction ($\alpha = 0.0125$) (cIPS $t(10) = 0.65$, $p = 0.531$ (*ns*); rIPS ($t(10) = 2.31$, $p = 0.043$ (*ns*)) (Figure 11A). Frontal region PMd showed trend-level significance after correction ($t(10) = 2.93$, $p = 0.015$ (*ns*)), with increased activation during DB trials that required the policy strategy. However, frontal region pre-PMd, shown previously to be important for contextual control based on an active second-order policy, strongly differentiated the policy strategy from the comparison strategy ($t(10) = 4.11$, $p = 0.002$).




Figure 12. Region of superior frontal sulcus that covaries with policy strategy. Voxel-wise corrected at $p < 0.05$ with a cluster correction of $p < 0.001$.

Trial timepoints were further broken down to investigate neural signals specifically during the second-order cue on DB compared to SB trials, when effects of the policy strategy might be more pronounced due to use of context-specific policy (Figure 11B, Top). Neither parietal subregion cIPS ($t(10) = 0.68$, $p = 0.513$ (*ns*)) nor rIPS ($t(10) = 1.30$, $p = 0.221$ (*ns*)) differentiated in activity for strategy at the time when second-order information was presented. As expected based on previous studies of hierarchical control (Badre & D'Esposito, 2007), frontal region PMd was also not sensitive to strategy use at the timepoint of the second-order cue ($t(10) = 1.56$, $p = 0.149$ (*ns*)). However, activity in pre-PMd additionally did not significantly differ between the strategies during the second-order cue ($t(10) = 2.11$, $p = 0.061$ (*ns*)). Overall, neither parietal nor frontal ROIs important for hierarchical policy abstraction differed based on strategy use while second-order policy abstraction information was provided.

Neural activity was next investigated at the timepoint of the relevant first-order cue (First-order DB v First-order SB) (Figure 11B, Bottom). As in previous comparisons of strategy, neither cIPS ($t(10) = 1.73$, $p = 0.115$ (*ns*)) nor rIPS ($t(10) = 0.70$, $p = 0.500$ (*ns*)) significantly changed activity based on strategy use during presentation of the first-order cue. However, trend-level increased activity was observed in PMd selectively during the first-order cue for the policy strategy ($t(10) = 2.82$, $p = 0.018$ (*ns*)). As predicted from previous evidence (Badre & D'Esposito, 2007), there were no significant changes in pre-PMd for policy versus comparison strategies exclusively at the timepoint of first-order rule use ($t(10) = 1.97$, $p = 0.077$ (*ns*)).

Regions that covary with strategy use

Lastly, an exploratory analysis was conducted that investigated which, if any, regions across the brain covary activity with utilization of a particular strategy. It can be inferred that subjects who exhibited a greater accuracy advantage on SB trials may have been using the comparison strategy more frequently. As such, a behavioral measure was assessed for each subject that subtracted mean accuracy on DB trials from SB trials (Figure 9B). A less negative value meant that the subject experienced less of an advantage on SB trials, and thus potentially depended on the policy strategy more frequently; a more negative value meant that the subject had a greater advantage on SB trials, which could infer comparison strategy use. Univariate results that used the behavioral measure of strategy use (see Materials and Methods, Behavior results) as a covariate revealed that a region within the right superior frontal sulcus (centroid MNI [-18, -45, 39], Figure 12) increased its activity with less use of the comparison strategy, and therefore corresponded with more consistent policy strategy use across task performance. The more that a subject solely depended on the policy strategy without advantage from the comparison strategy on SB trials, the greater activation seen in this region of the right superior frontal sulcus.

Discussion

In the present study, task performance was broken down into trials that either required a second-order policy abstraction strategy (policy strategy) or trials that allowed a match/nonmatch strategy dependent only on evaluation of first-order cues (comparison strategy). Results overall indicated that the frontal region known to be important for performance of second-order policy abstraction decisions, pre-PMd, is more responsive to a policy strategy that is dependent on second-order policy abstraction. Frontal region PMd showed trend-level significance in activation differences between strategies over all trial time points, as well as at the timepoint of the first-order cue. Both parietal subregions cIPS and rIPS did not significantly differ in activation between strategies. Taken altogether, results support hypotheses that activity in frontal ROIs implicated in policy abstraction reflect specificity for policy strategies at decision-relevant timepoints, while parietal ROIs may not be preferential to policy versus other types of task abstraction that differ in strategies.

Although cIPS was selective for performing second-order versus first-order policy abstraction decisions (see Chapter 2), activation in this subregion did not differ between the policy strategy and the comparison strategy (Figure 11). The policy strategy used the hierarchical rules of the designed policy abstraction task, but the comparison

strategy used a heterarchical rule similar to those employed in a relational abstraction task, in which subjects assessed the relationship between the first-order action outcomes to determine if they were matching or not. If the first-order cue button associations matched (SB trials, Figure 7), then subjects could refrain from performing second-order policy abstraction, but the higher-order rule of relational abstraction was nevertheless still performed on those trials (Badre, 2008; Badre & Nee, 2018). Thus, cIPS activity could reflect higher-order abstract task information specifically for both hierarchical policy abstraction and heterarchical relational abstraction. Previous studies have not discovered consistent differences in frontal activity across different types of task abstraction (Badre & Nee, 2018; Bahlmann et al., 2015; Nee et al., 2014), but those studies have notably lacked investigation of highly implicated parietal regions such as cIPS.

Alternatively, the lack of activation differences within cIPS could have been observed here because subjects encoded the second-order cue as relevant higher-order abstract sensory information across both strategy trial types. Some subjects self-reported that they used both strategies within a single trial (though this was anecdotal feedback not quantified for every subject), suggesting that second-order cue information may have been treated as decision-relevant on most if not all trials. Subjects may have been performing policy abstraction decisions during the second-order cue on both SB and DB trials more often than not, which may explain why pre-PMd activity was not significantly higher specifically during second-order cue presentation on policy trials. Because caudal IPS has strong structural connectivity to retinotopically defined visual cortices (Bray et al., 2013; Silver & Kastner, 2009), and was previously shown to discriminate activity based on sensory salience (Figure 4B, see Chapter 2), it is possible that cIPS preferentially tracks sensory information for higher-order abstract decisions without any consideration whatsoever as to whether the task is hierarchical policy abstraction or heterarchical relational abstraction. Ultimately, additional research is necessary to tease apart the specifics of cIPS activity during different types of task abstraction as well as during visual control processes.

Within rIPS, results support the hypothesis that this subregion would not change activity if attentional control over the first-order cue did not differ between the two strategies and the cue maintained its relevance for motor output (Figure 11). This parietal subregion could be probed further by investigating activity changes at the timepoint of the irrelevant cue. Since comparison strategy trials upgraded the irrelevant cue to equally relevant in the match/nonmatch decision of the first-order stimuli, I predict that rIPS would show increased relative activity for comparison strategy versus policy strategy during irrelevant cue presentation.

In this study, subjects could flexibly use either strategy within a given trial (with varying success based on trial type), and thus subjects used them to varying degrees across their total performance (Figure 9B). With subject-initiated switches between strategies, the finding that a region in the anterior portion of right superior frontal sulcus (SFS) covaries activity based on use of a particular strategy may be supported by studies showing its role in task-set switching (Cutini et al., 2008; De Baene, Duyck, Brass, & Carreiras, 2015; Sakai & Passingham, 2006) (Figure 12). Using fMRI and computational modeling, Donoso and colleagues reported that an anatomically close region identified as right frontopolar cortex (MNI [28,52,20]) increased activity when

subjects switched to a new strategy to be employed (Donoso, Collins, & Koechlin, 2014). Further evidence for the significance of anterior SFS in abstract task strategies was shown in a patient study where impairments in specific rule-based behaviors were associated with lesions to the frontopolar cortex (P. W. Burgess, Veitch, De Lacy Costello, & Shallice, 2000). Given this evidence, anterior SFS in this study may facilitate switching to the policy strategy on DB trials when the comparison strategy is not sufficient to make the decision, though further analysis on this finding was outside of the scope of the current study.

It has not been unnoticed that the qualitative difference in performing the policy strategy versus comparison strategy in this study could be likened to a congruency task. When the irrelevant and first-order cues both referred to the same button on SB trials, one could consider this as congruency with the stimulus-response mappings; DB trials that necessitated the policy strategy could be considered incongruent. Response time and accuracy results (Figure 8) support this as a possibility. Under this framework, the abstract decision-making task could be reconceptualized as a form of the temporal flanker task (Hazeltine, Lightman, Schwarb, & Schumacher, 2011), where the irrelevant and first-order cues are separated temporally in a trial and the presentation position determines if the irrelevant cue primes response to the first-order cue (Note: in-depth look at how position of cue types might effect relevant versus irrelevant decision items differently is pursued next in Chapter 4). Increased activity in pre-PMd and trend-level in PMd could be explained apart from specificity for a hierarchical policy abstraction strategy due to a similar congruency effect as shown in this study (Hazeltine et al., 2011), or a combination of the policy abstraction strategy and congruency effects. Previous studies have shown greater activation broadly across the DLPFC during higher response conflict due to incongruent response mappings (M. M. Botvinick, Braver, Barch, Carter, & Cohen, 2001; Matsumoto & Tanaka, 2004). Other work has directly examined the overlap between congruency effects and policy abstraction, concluding that first-order rule nonmatch trials (potentially analogous to DB trials that require the policy strategy) produced increased activation in PMd and left rostralateral prefrontal cortex, in support of frontal results in the current study (Figure 10) (Wolfensteller & von Cramon, 2011; Wolfensteller & Von Cramon, 2010). Interestingly, the study by Wolfensteller and von Cramon additionally found that congruency effects were seen only in a subset of subjects that completed their task using a particular strategy.

In distracter interference tasks, the congruency sequence effect is a smaller measured congruency effect (reduced reaction time and accuracy advantages) after incongruent trials than after congruent trials (M. M. Botvinick et al., 2001; Gratton, Coles, & Donchin, 1992). Importantly, the present study did not find any post-hoc congruency sequence effects, affirmed by the fact that task design was not optimized to capture this effect and precise reaction time measures were not possible without motor confounds during first-order cue presentation.

Finally, the observed behavioral shift to a comparison strategy away from the instructed policy strategy on SB trials reviewed here (Figure 9A) may reflect subjects minimization of cognitive effort (M. Botvinick & Braver, 2015; Kurzban, Duckworth, Kable, & Myers, 2013; Shenhav et al., 2017). It is interesting to note that this shift to the comparison strategy was only useful for obtaining a correct answer half of the time (on

SB trials only), while the policy strategy could be successfully used on both SB and DB trials consistently to reach a correct answer. One possible explanation for this sub-optimal strategy shift is that the effort required to recruit supplemental brain regions for implementing the second-order and then also the first-order policies with the policy strategy (Figure 10) may have been greater than that required to focus on the action outcome of the first-order cues with the comparison strategy. Whether the minimization of cognitive effort in this case was due to a cost-benefit analysis that weighs abstract reasoning in the policy strategy as too costly (Székely & Michael, 2020), or due to a facilitatory match/nonmatch behavior for the comparison strategy that was adopted as habitual over the course of performing 660 trials (Miller & Cohen, 2001), are future directions for further research on this effect.

In summary, results indicate that subregions of the IPS track information about higher order abstract decisions based on the associated level of abstraction, but do not respond preferentially for strategies that rely on policy abstraction compared to other forms of task abstraction. Instead, frontal regions may take the role of distinguishing between policy abstraction strategies and alternative match/nonmatch comparison strategies.

Chapter 4:

Are retained versus removed working memory items reflected in parietal activity during decision-making?

Abstract

Every discrete decision requires selection from at least two unique options, and in the absence of consistent external representations, those options are held in working memory. As a result, relevant working memory decision items are retained for use in a future motor response while irrelevant items must be removed to reduce retrieval interference (Lewis-Peacock et al., 2012). Parietal cortex and IPS specifically are associated with working memory item representation and control (Chatham et al., 2014; Cowan et al., 2011; M. D'Esposito et al., 1998), implicating it in possible retention and removal processes during decision-making. To test whether an active process of removal and/or retention exists in parietal cortex, fMRI data collected from the same abstract decision-making task developed for this thesis work were re-analyzed. A subset of the data were selected where second-order cues were presented in position 2 within a trial, effectively acting as a retro-cue for the relevance or irrelevance of the previously presented cue in position 1. Investigating activity at this timepoint would reveal working memory item removal if the previous cue was irrelevant, and retention if the cue was relevant. Hypothesis 1 predicted that parietal cortex would reflect an active process of content unbinding during irrelevant working memory item removal, but results instead suggested that removal may be a latent process that is undetectable via univariate imaging. Hypothesis 2 investigated working memory item retention and predicted that rIPS would increase activity for retaining response-relevant first-order contents while cIPS would not change activity for working memory control of first-order information. Contrary to these hypotheses, results revealed that rIPS did not significantly change activity for working memory item retention of relevant first-order information, while cIPS increased activity at this timepoint. This finding supports an alternative hypothesis that cIPS and rIPS should instead be considered to be primarily engaged with decision-making aspects of the policy abstraction task, and not necessarily implicated in aspects of working memory item retention. Continued research is necessary to further expose the neural underpinnings of working memory item removal and retention.

Introduction

Every discrete decision requires selection from at least two unique options. Working memory is required to internally hold decision options while they are being weighed or updated over time (Mark D'Esposito & Postle, 2015; Nee & Brown, 2013). Injury to brain regions important for working memory produces deficits in decision-making (Smith, Xiong, Elkind, Putnam, & Cohen, 2015) and, conversely, increased working memory capacity is associated with improvements in decision-making performance (Klingberg, 2010). At the time of deciding, a single relevant option is selected from working memory. The relevant item must be retained in working memory if the associated behavioral output is to be performed at a future timepoint. On the other hand, irrelevant decision options must be removed from working memory. If irrelevant items are not removed, then they can cause interference that results in increased errors

and delayed reaction time during decision-making tasks (LaRocque, Lewis-Peacock, Drysdale, Oberauer, & Postle, 2013; Oberauer, 2001). Thus, removal of irrelevant working memory items during decision-making serves a significant function. Despite this fact, the underlying neural processes that subservise working memory item removal and retention is still largely unknown.

Several theories exist regarding the mechanisms that support generalized working memory item removal. Under a framework that conceptualizes working memory items as sensory content (e.g. visual display of the letter “A”) that is bound to a context (e.g. third in serial position of a string of other letters), some researchers propose that a complete removal of irrelevant information would entail an “active unbinding” of the content from the context (Lewis-Peacock et al., 2018). This unbinding is a deliberate process for removing irrelevant information due to a limited working memory capacity in which retrieval is impaired when there is competition between items (Buschman et al., 2011; Unsworth & Engle, 2007). While neural correlates of this process remain elusive, decoding studies have revealed that irrelevant item contents become undetectable shortly after being designated as unnecessary (Lewis-Peacock et al., 2012) and representations of those contents can be reinstated if items are re-categorized as relevant (Sprague, Ester, & Serences, 2016).

A logical place to begin studying working memory item removal and retention would be to investigate regions of the brain known to be involved in the control of working memory items. Along with the PFC, parietal cortex has been strongly implicated in working memory tasks (Friedman & Goldman-Rakic, 1994). In nonhuman primates, neurons in posterior parietal cortex exhibit increased firing rate activity during memory delay periods (Chafee & Goldman-Rakic, 1998). Human fMRI studies have shown equivalent evidence of working memory delay activity in posterior parietal cortex (M. D’Esposito et al., 1998; Pessoa, Gutierrez, Bandettini, & Ungerleider, 2002). The IPS specifically has been shown to maintain sensory details about an item held in visual working memory in an unbiased state even during distraction (Lorenc, Sreenivasan, Nee, Vandenbroucke, & D’Esposito, 2018).

As such, IPS and parietal cortex as a whole are pertinent to investigate how the neural representation of working memory items might change for removal or retention during decision-making. Since parietal cortex has been implicated in the binding of sensory information to contextual information (Gosseries et al., 2018) and the maintenance of both currently attended and unattended information in working memory (Christophel, Iamshchinina, Yan, Allefeld, & Haynes, 2018), there may also be an explicit parietal mechanism for unbinding during removal of information that is determined to be irrelevant, and/or retention of relevant information in service of an upcoming behavioral response. Certainly related to removal or retention, previous research has implicated parietal cortex in the selection of working memory content through persistent oscillatory neural activity with prefrontal and occipital regions (Quentin et al., 2019), as well as through gating mechanisms with corticostriatal systems (Chatham et al., 2014). Therefore, if the removal process is an active unbinding of sensory information to a context as hypothesized elsewhere, then parietal cortex may reflect this in BOLD activity when a working memory item is determined to be irrelevant via a retro-cue.

Due to the counterbalanced design of the policy abstraction task used in Chapters 2 and 3, the second-order, first-order, and irrelevant cues appeared equally likely in serial position 1, 2, or 3 across all trials. With 660 total trials per subject, specific trial subtypes of cue position could be selected from the dataset for re-analysis with sufficient power. On trials where the second-order cue is in serial position 2, decision information pertaining to the cue that appeared in position 1 is held in working memory and assessed whether it is relevant or irrelevant. On these trials, the second-order cue serves as the retro-cue for the previously encoded cue. At the time of position 2 then, prior cue information in working memory is either retained and used for an upcoming motor response if it is relevant, or removed from working memory if it is irrelevant. Thus, both the retention and removal of working memory items could be successfully explored using this task design.

In Chapter 2 of this thesis, two subregions of IPS were identified that were strongly implicated in the representation of decision items crucial for performance during policy abstraction. Of those two subregions, a caudal area (cIPS) was sensitive to the sensory salience of abstract cues as well as to second-order decision information, while a rostral area (rIPS) was sensitive to the decision motor output as well as attentional control of first-order information. Since cIPS and rIPS were implicated in encoding this task-relevant information, they may also be important for controlling working memory representations of cue information for a removal and/or retention process. If this is the case, then IPS subregions would remove/retain item representations at their associated order of abstraction, meaning that rIPS activity may increase when relevant first-order stimulus-motor mapping information is removed during active unbinding or retained in working memory for an upcoming decision, while cIPS activity may not differentiate between removal or use of first-order information as a region that does not track lower-order policy (see Chapter 2). On the other hand, if cIPS and rIPS are not involved with working memory item removal/retention, then both of these regions may instead differentiate activity based on specific policy abstraction decision aspects unrelated to working memory content. In this alternate case, cIPS and rIPS may both increase activity for the decision timepoint when all necessary information is present, and thus both cIPS and rIPS would increase activity when the first-order information is retro-cued as relevant compared to irrelevant.

The present Chapter 4 investigates two hypotheses regarding working memory item removal and use, respectively. Hypothesis 1 posits that if there are parietal regions that actively unbind sensory information during working memory item removal, then they will have greater activity during irrelevant cue information removal. Switching focus to working memory item retention, Hypothesis 2 next proposes that if IPS represents working memory item retention of relevant information during decision-making, then rIPS would increase activity at the timepoint of working memory item retention but cIPS would not be responsive to changes in first-order information held in working memory at that same time. Results of this study are not able to lend evidence towards BOLD activity signatures of working memory removal, but instead may support the existence of activity silent representations of removal (Sprague et al., 2016; Stokes, 2015). In addition, results in cIPS and rIPS point to more exclusive roles in the policy abstraction decision process that may not involve representations of working memory item removal and/or retention.




Figure 13. Policy abstraction decision-making task with two trial subtypes based on serial position of the second-order, first-order, and irrelevant cues within the trial. **A)** Hypothesis 1 trial subtype. At the timepoint of the second-order cue (indicated by the black star) on trials where the irrelevant cue is in position 1 and the second-order cue is in position 2 (ISF-S), activity in parietal cortex would reflect the removal of the irrelevant cue from working memory as it was determined to be not useful for the trial. **B)** Hypothesis 2 trial subtype. At the timepoint of the second-order cue on trials where the first-order cue is position 1 and the second-order cue is position 2 (FSI-S), activity in parietal cortex would reflect the retention of the first-order cue in working memory and its subsequent use for determining a response.

Materials and Methods

In this re-analyzed dataset, 11 subjects (eight women and three men, ages 18-45 yrs) completed a policy abstraction task with conditions that dissociated various known parietal functions (see Chapter 2 Materials and Methods for details). The same methods up to and including “*fMRI pre-processing*” in Chapter 2 were used for this study.

Univariate analysis.

With the fully pre-processed data, a new general linear model was created that evaluated 36 different conditions as separate regressors. Only correct trials were included, and they were grouped with consideration for the position of all three cue types within a trial while capturing activity during only one cue display, different from the organization described in Chapter 3. The regressors evaluated all permutations of cue type and position within a trial (second-first-irrelevant, second-irrelevant-first, first-second-irrelevant, etc.), as well as response trial type (same button or different button responses of the relevant and irrelevant first-order cues). As such, the first six regressors included the following stimulus presentation conditions: (1) second-order cue in position 1, first-order cue in position 2, and irrelevant cue in position 3 on SB trials during the duration of just the second-order cue display (SFI-S), (2) second-order cue in position 1, first-order cue in position 2, and irrelevant cue in position 3 on SB trials during the duration of the first-order cue (SFI-F) (3) second-order cue in position 1, first-order cue in position 2, and irrelevant cue in position 3 on SB trials during the irrelevant cue (SFI-I) (4) First-order cue in position 1, second-order cue in position 2, and irrelevant cue in position 3 on SB trials during the first-order cue (FSI-F) (5) First-order cue in position 1, second-order cue in position 2, and irrelevant cue in position 3 on SB trials during the second-order cue (FSI-S), (6) First-order cue in position 1, second-order cue in position 2, and irrelevant cue in position 3 on SB trials during the irrelevant cue (FSI-I), and so on for the 30 remaining permutations in this manner. Identical to methods in previous chapters, each regressor was created from the convolution of a gamma probability density function that peaked at 6-sec with a condition-specific vector of stimulus onsets and durations of 2-sec for stimulus presentation. Resultant estimated beta coefficients for each subject were mapped to the MNI template for use in group-level analyses. Response trial types (SB or DB trials) were collapsed for the remainder of this analysis to focus on specific cue/position timepoints. Whole-brain voxelwise contrasts were subsequently created for the timepoint of the second-order cue on ISF trials compared to the second-order cue on FSI trials (ISF-S v FSI-S), and the timepoint of the second-order cue on FSI trials compared to second-order cue on ISF trials (FSI-S v ISF-S)(Figure 13). 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). All voxels in whole-brain group level maps were then thresholded at uncorrected $p < 0.001$, then cluster-corrected to produce whole-brain maps at $p < 0.05$ using AFNI function “3dClustSim”, with bi-sided voxel thresholding and second-nearest neighbor clustering settings ($nn = 2$).




Figure 14. Group-level behavioral effect of cue position permutation on trial accuracy for the second-order, first-order, and irrelevant cues. “S” abbreviates second-order cue, “F” abbreviates first-order cue relevant on the trial, “I” abbreviates irrelevant cue. Error bars represent between-subject SEM, while n.s. denotes non-significance.

Because the primary hypothesis regarding the removal of irrelevant decision information (Hypothesis 1) was specific to parietal cortex, the ISF-S v FSI-S contrast was masked further by anatomical parietal cortex defined using cytoarchitectonic macro labels from the standard AFNI atlas “CA_ML_18_MNI”. The other contrast, FSI-S v ISF-S, was next analyzed for results in specific ROIs for Hypothesis 2.

fMRI-based ROI analysis.

Hypothesis 2 predicted that parietal subregions important for policy abstraction decisions would differentially represent the use of relevant decision information held in working memory based on its level of abstraction. Therefore, a subregion sensitive to attentional control of relevant first-order cue information would activate when that relevant stimulus-response information is retained in working memory for performing an upcoming behavior, while a subregion that selectively encoded second-order cue information would not change activity. To test this hypothesis, the same two parietal ROIs from Chapter 2 were used for this study (cIPS, rIPS), selected from the Schaefer et al 1000-area parcellation and resampled to task data voxel dimensions (3 x 3 x 3mm). An ROI mask of the cIPS and rIPS parcels was applied to each subject’s univariate whole-brain FSI-S v ISF-S contrast dataset (see Chapter 2 for detailed ROI analysis methods). The mean beta value from all voxels within the selected parcel were then computed for each subject and paired t-tests were performed at the group level using custom MATLAB scripts.

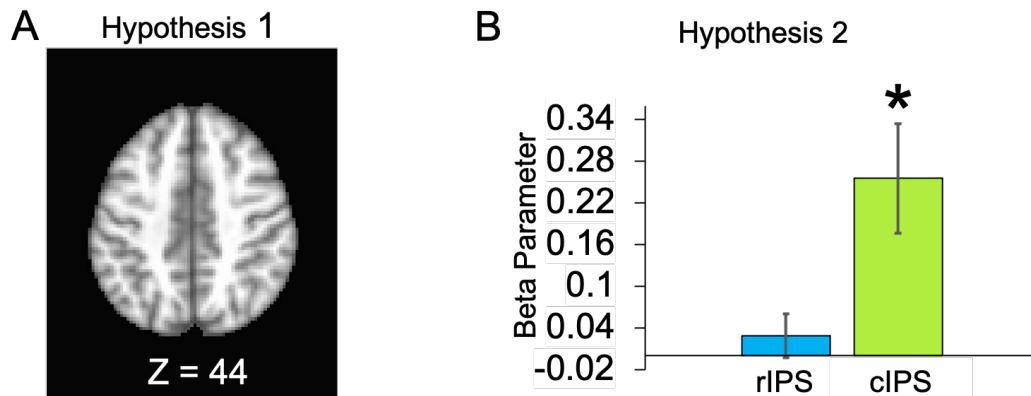


Figure 15. Univariate analysis. **A.** Hypothesis 1 results. No significant activity differences detected throughout the brain for the removal of irrelevant decision information in the ISF-S v FSI-S contrast. Whole-brain voxel-wise corrected at $p < 0.05$ with a cluster correction of $p < 0.001$. **B.** Hypothesis 2 results. Within-ROI activity for the univariate contrast of FSI-S v ISF-S, with significant activity representing greater activation when relevant information in working memory is retained for deciding a future motor response compared to removed. Beta parameter estimate scaled to represent percent signal change (see Chapter 2 Methods). Error bars represent between-subject SEM. Asterisk indicates $p < 0.05$.

Results

This study investigated two independent hypotheses relating to parietal representations of retained versus removed decision items held in working memory. This was accomplished by assessing brain activity during specific trial subtypes of an abstract decision-making task that sequentially displayed a second-order, a relevant first-order, and an irrelevant first-order cue in all permutations of serial position (Figure 13).

Behavior.

Collapsing all other task factors across serial position of cue types confirmed that behavioral performance on the task was not impacted by a particular position permutation ($F(5,60) = 1.01$, $p = 0.422$ (*ns*); Figure 14). Response time was also not significantly different across cue type position permutations ($F(5,60) = 1.97$, $p = 0.096$). Details and results for other behavioral controls for this task can be found in Chapters 2 and 3.

Univariate functional imaging maps.

The first hypothesis (Hypothesis 1) addressed in this study examined neural correlates in parietal cortex when irrelevant working memory items are determined to be unnecessary for a decision and are removed from working memory as a result. Investigating brain activity at the timepoint of the second-order cue on trials where the irrelevant cue was presented in position 1 and the second-order cue was presented in position 2 (ISF-S) captured responses related to the identification that the irrelevant cue was, indeed, irrelevant and was thus removed from working memory (Figure 13A). A contrast of this activity with separate conditions when the cue in position 1 was useful for the decision (FSI-S) would therefore leave only neural signals from parietal cortex for the process of working memory item removal. This univariate contrast (ISF-S v FSI-S)

restricted to anatomical parietal cortex was assessed at the group level for all subjects, without any significant results surviving cluster-correction (Figure 15A). As an exploratory follow-up, the same contrast was assessed outside of parietal cortex at the whole-brain level. There were also no significant clusters across the brain for this contrast. Any potential changes in neural activity related to working memory item removal during decision-making were not detected, at least at the whole-brain univariate level.

Functional imaging ROI analysis.

The second hypothesis considered in this study (Hypothesis 2) investigated explicit activity in IPS during timepoints when relevant first-order information held in working memory is retained for subsequent motor output during a specified response period (Figure 13B). Two ROIs identified in Chapter 2 of this thesis, cIPS and rIPS, were used for this analysis of the FSI-S v ISF-S contrast dataset. Results revealed that cIPS increased activity at this timepoint ($t(10) = 3.24$, $p = 0.009$), while rIPS did not ($t(10) = 0.91$, $p = 0.382$ (*ns*))(Figure 15B).

Discussion

This study investigated two separate hypotheses regarding the neural correlates of decision information held in working memory that is removed versus retained for a future action output. Results indicate a lack of detectable activity for working memory item removal of irrelevant information, and do not clearly support the involvement of IPS subregions important for policy abstraction in activities for relevant working memory item retention during decision-making.

Hypothesis 1 explicitly sought to identify regions within parietal cortex that change activity for the removal process, under a framework that theorized working memory item removal as an active unbinding of sensory information to context (Lewis-Peacock et al., 2018). Here, the current study found no significant neural responses for this process in parietal cortex nor in any region across the brain (Figure 15A). This negative result, while difficult to interpret, may support an existent modeling theory of activity silent working memory representations that are enacted at the level of dynamic synaptic weights (Oberauer, 2018; Pals, Stewart, Akyürek, & Borst, 2020; Stokes, 2015). This theory models removal as an instantiation of synaptic Hebbian anti-learning, which detaches remembered items from their bound contexts in a manner that is undetectable by univariate fMRI methods (Oberauer, Lewandowsky, Farrell, Jarrold, & Greaves, 2012).

Lack of observable univariate activity for removal of irrelevant cue information in this task may also be in agreement with another survey of discarded memory items. In a study using a computational model that utilized inverted encoding model (IEM) simulations of fMRI data, Lorenc et al (2020) found that discarded working memory items produced a negative IEM reconstruction in visual cortical regions with item representations (Lorenc, Vandenbroucke, Nee, de Lange, & D'Esposito, 2020). The authors concluded that this finding may be caused by suppression of the irrelevant information within implicated brain regions. Representation suppression would still require cellular activity in those regions to engage inhibitory mechanisms, thus univariate fMRI methods might not detect when irrelevant items are discarded as a

significant decrease in BOLD signal. Notably, suppression is a different theorized mechanism for how irrelevant working memory items might be represented, in contrast to the active unbinding process of features to contexts during removal proposed elsewhere (Lewis-Peacock et al., 2018; Sahan, Sheldon, & Postle, 2019).

In a closer-targeted examination of parietal cortex, Hypothesis 2 predicted that IPS subregions important for policy abstraction decisions would show activity that reflected relevant working memory item retention contingent on the level of abstraction that the information conveyed. Within this hypothesis, rIPS would increase activity for working memory item retention at the time when relevant first-order stimulus-response information is retained compared to removed, but this same timepoint would not produce activity differences in cIPS because it was previously shown to be invariant to first-order information attentional control (see Chapter 2). Contradictory to this primary hypothesis, results revealed that rIPS did not significantly change activity for working memory item retention of relevant first-order information, while cIPS increased activity at this timepoint (Figure 15B). This finding supports an alternative hypothesis that cIPS and rIPS should instead be considered to be primarily engaged with decision-making aspects of the policy abstraction task, and not necessarily implicated in aspects of working memory item retention. Certainly, as these parietal regions were deliberately identified by their sub-network membership with hierarchically-organized abstract decision-making frontal regions (Badre & D'Esposito, 2007), their involvement with explicit decision-related mechanisms might be expected to be stronger than other functions such as working memory item retention.

Instead of retention of relevant first-order items held in working memory, it seems more likely that the increased activity in cIPS during the second-order cue seen on FSI trials reflects increased activity for the decision itself, since all necessary information for performing the task is available at the timepoint of second-order cue presentation on these trials but is not yet available on ISF trials. This interpretation is in agreement with Chapter 2's finding that cIPS had increased activity for cues with high versus low coherence stimuli (Figure 4B, Top; See Chapter 2), when the task decision can be made more quickly compared to low coherence stimuli that require evidence collected over longer temporal scales. This explanation is further supported by work from Cookson and Schumacher that utilized a hierarchical pre-cuing task to show that frontal activity in previously defined hierarchical regions tracks with the amount of information available for processing at the timepoint of investigation (Cookson & Schumacher, 2021). This may mirror the current finding in that at the timepoint of the second-order cue on ISF trials, subjects can only process the second-order cue information for the decision, but on FSI trials when all decision-relevant information has been presented, subjects can process the whole policy abstraction task.

This interpretation further delineates rIPS from cIPS as a subregion implicated in the action generation of sensorimotor transformations (Erickson & Kayser, 2013). Activity did not increase in rIPS when relevant first-order decision contents stored in working memory were retained for use, but rIPS may show increased activity at trial timepoints closer to actual motor output at response phase onset. A future analysis to investigate BOLD time course signatures leading up to the response period would confirm the role of rIPS in the conversion of first-order information into an action.

Additional research to disentangle not only the differential contributions of IPS subregions but also the extent of cognitive control processes implicated in those subregions with higher-order abstract representations is necessary to fully determine the role of IPS in working memory item retention and removal.

Chapter 5: Dissertation Conclusions

This thesis investigated the multifaceted role of parietal cortex in the performance of executive functions that take place during decision-making. When search space was defined only by anatomical parietal cortex, frontal regions implicated in second- and first-order policy abstraction decisions had strongest intrinsic connectivity with subdivisions of the left IPS, previously suggested as an important region for cognitive control networks (Yeo et al., 2011). Chapter 2 lent evidence to an organization of hierarchical frontoparietal subnetworks with IPS based on level of policy abstraction, where second-order decision information activated pre-PMd and caudal IPS (cIPS) selectively, while relevance of first-order decision information activated PMd and rostral IPS (rIPS) (Figure 4A). Parietal cIPS and rIPS also differentially represented sensory and motor aspects of decision information, suggesting that their contributions to abstract decision-making networks may subserve parallel processes to integrate visual and motor information streams (Figure 4B). Across IPS, a caudal-rostral gradient of selectivity for higher order abstraction was discovered, spatially mirrored to that seen in frontal cortex (Figure 5).

Those findings emphasized that IPS processes perceptual, motor, and attentional aspects of abstract decisions differently when it is segregated into network-defined functional subdivisions, but how specific is the activity in cIPS and rIPS for different types of task abstraction? When comparing the performance of a hierarchical policy strategy and a heterarchical comparison strategy within the same abstract task, Chapter 3 revealed that cIPS and rIPS did not show activity suggesting specificity for the type of abstract decision being made in each strategy (Figure 11A). Instead, Chapters 2 and 3 showed that cIPS represents second-order policy and may also equivalently represent higher-order abstract information across abstraction types as long as that information is encoded as task-relevant. Similarly, rIPS may not distinguish the type of abstraction when attentional control of cue information for the purpose of action transformation does not differ. On the other hand, frontal activity in pre-PMd and PMd was sensitive to the abstraction strategy, highlighting their roles in explicitly hierarchically-organized policy decisions and underscoring unique yet complementary activity between parietal and frontal regions.

Finally, do parietal cortex and IPS subregions implicated in abstract decisions represent decision information that is removed or retained in working memory? Chapter 4 examined parietal activity when decision information held in working memory is either removed or retained for a subsequent action output after being identified as irrelevant or relevant, respectively. There were no observed activity differences in parietal cortex nor the rest of the brain for working memory item removal (Figure 15A), which may lend evidence for a model where removal is an “activity silent” process not detectable using univariate fMRI methods, potentially through the change of synaptic weights associated with item representations (Lewis-Peacock et al., 2018; Stokes, 2015). Conclusive evidence for working memory item retention was also not found within rIPS and cIPS, instead suggesting that rIPS activity may reflect externally-directed attentional control of first-order information for action transformation and cIPS may increase activity when second-order information is utilized for reaching a decision threshold.

Overall, this thesis describes how IPS contributes to abstract decision-making, but a number of questions still remain. Just how modular are the discovered IPS subregions, cIPS and rIPS, for a given order of abstraction? In order to bridge this research with foundational findings of hierarchical control in frontal cortex (Badre & D'Esposito, 2007), future research should examine if third- and fourth-order policy decisions elicit distinct task activity within IPS. Previous research suggests that parietal cortex may only be responsive to lower level context updates, such as that at the second-order policy level, and that PFC interactions with the striatum resolve higher-order context updates (Nee & Brown, 2013). If so, then cIPS may be selective only to second-order abstraction and the gradient of activity in IPS for second-order decisions may be consistent regardless of increase in order. However, if IPS flexibly adapts based on task demands, then cIPS may actually shift selectivity to whatever the highest level of abstraction is for a given task.

Along those lines, Chapter 3 would be supplemented by a study with consistent use of one abstraction strategy for a defined set of trials, rather than the subject-discretionary switching between strategies that the current task paradigm supported. This would strengthen conclusions that cIPS and rIPS are truly flexible across multiple types of higher order task abstraction, rather than the alternative that they may represent task-specific activity in policy abstraction decisions. Additional tasks to probe the boundaries of sensory versus abstraction representations in cIPS, as well as motor versus attentional modulation in rIPS, would provide clarity about these additional functions of cognitive control observed in these subregions.

Lastly, there is much to be learned regarding working memory item removal and retention during abstract decision-making. Results underscored that univariate fMRI methods alone are insufficient to probe the neural instantiation of the observed behavioral effects. Multivariate methods would assist in detecting finer changes in activity for regions that represent working memory items, such as the sensory regions that encoded the stimuli (Mark D'Esposito & Postle, 2015). Decoding of the middle temporal (MT) visual area during first-order cues represented by motion stimuli and visual area V4 during first-order color stimuli would reveal additional insight into removal and retention. With respects to previously defined IPS subregions, additional analyses should elaborate on the role of cIPS in reaching abstract task decision thresholds, and the role of rIPS in response generation. Investigation at the timepoint of position 3 of ISF-F trials versus FSI-I trials would likely reflect greater activity in cIPS for reaching the decision threshold during presentation of the first-order cue, as well as greater activity in rIPS for upcoming motor execution of the decision. Timeseries analyses would further confirm the temporal dependencies of activity in these subregions at applicable points in trials.

These additional analyses would help define effects seen in cIPS and rIPS, ultimately providing clarity for the complex contributions of parietal cortex in decision-making and its relationship with frontal cortex.

References

- Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences*, 12(5), 193–200.
<https://doi.org/10.1016/j.tics.2008.02.004>
- Badre, D. (2012). Opening the gate to working memory, 109(49), 19878–19879.
<https://doi.org/10.1073/pnas.1216902109>
- Badre, D., & D'Esposito, M. (2007). Functional Magnetic Resonance Imaging Evidence for a Hierarchical Organization of the Prefrontal Cortex. *Journal of Cognitive Neuroscience, Early Access*(Early Access), 080219115128817–080219115128818.
<https://doi.org/10.1162/jocn.2007.91201>
- Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nature Reviews. Neuroscience*, 10(9), 659–669.
<https://doi.org/10.1038/nrn2667>
- Badre, D., Kayser, A. S., & D'Esposito, M. (2010). Frontal cortex and the discovery of abstract action rules. *Neuron*, 66(2), 315–326.
<https://doi.org/10.1016/j.neuron.2010.03.025>
- Badre, D., & Nee, D. E. (2018). Frontal Cortex and the Hierarchical Control of Behavior. *Trends in Cognitive Sciences*, 22(2), 170–188.
<https://doi.org/10.1016/j.tics.2017.11.005>
- Bahlmann, J., Blumenfeld, R. S., & D'Esposito, M. (2015). The Rostro-Caudal Axis of Frontal Cortex is Sensitive to the Domain of Stimulus Information. *Cerebral Cortex*, 25(7), 1815–1826. <https://doi.org/10.1093/cercor/bht419>
- Bettencourt, K. C., & Xu, Y. (2016). Decoding the content of visual short-term memory under distraction in occipital and parietal areas, 19(1).
<https://doi.org/10.1038/nn.4174>
- Binder, J. R., Medler, D. A., Desai, R., Conant, L. L., & Liebenthal, E. (2005). Some neurophysiological constraints on models of word naming. *NeuroImage*, 27(3), 677–693. <https://doi.org/10.1016/j.neuroimage.2005.04.029>
- Bjork, R. A. (1970). Positive forgetting: The noninterference of Items intentionally forgotten. *Journal of Verbal Learning and Verbal Behavior*, 9(3), 255–268.
[https://doi.org/10.1016/S0022-5371\(70\)80059-7](https://doi.org/10.1016/S0022-5371(70)80059-7)
- Botvinick, M., & Braver, T. (2015). Motivation and Cognitive Control: From Behavior to Neural Mechanism. *Annual Review of Psychology*, 66(1), 83–113.
<https://doi.org/10.1146/annurev-psych-010814-015044>
- Botvinick, M. M. (2008). Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, 12(5), 201–208. <https://doi.org/10.1016/j.tics.2008.02.009>
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3), 624–652.
<https://doi.org/10.1037/0033-295X.108.3.624>
- Brass, M., Ullsperger, M., Knoesche, T. R., Cramon, D. Y. von, & Phillips, N. A. (2005). Who Comes First? The Role of the Prefrontal and Parietal Cortex in Cognitive Control. *Journal of Cognitive Neuroscience*, 17(9), 1367–1375.
<https://doi.org/10.1162/0898929054985400>
- Braver, T. S. (2012). The variable nature of cognitive control: A dual mechanisms framework. *Trends in Cognitive Sciences*, 16(2), 106–113.

- <https://doi.org/10.1016/j.tics.2011.12.010>
- Bray, S., Arnold, A. E. G. F., Iaria, G., & MacQueen, G. (2013). Structural connectivity of visuotopic intraparietal sulcus. *NeuroImage*, *82*, 137–145.
<https://doi.org/10.1016/j.neuroimage.2013.05.080>
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The Analysis of Visual Motion: A Comparison of Neuronal and Psychophysical Performance. *The Journal of Neuroscience*.
- Brody, C. D., & Hanks, T. D. (2016). Neural underpinnings of the evidence accumulator. *Current Opinion in Neurobiology*, *37*, 149–157.
<https://doi.org/10.1126/science.1233912>
- Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science*, *340*(6128), 95–98.
<https://doi.org/10.1126/science.1233912>
- Bunge, S. A. (2004). How we use rules to select actions: A review of evidence from cognitive neuroscience. *Cognitive, Affective, & Behavioral Neuroscience*, *4*(4), 564–579. <https://doi.org/10.3758/CABN.4.4.564>
- Bunge, Silvia A., Hazeltine, E., Scanlon, M. D., Rosen, A. C., & Gabrieli, J. D. E. (2002). Dissociable Contributions of Prefrontal and Parietal Cortices to Response Selection. *NeuroImage*, *17*(3), 1562–1571. <https://doi.org/10.1006/nimg.2002.1252>
- Bunge, Silvia A., Kahn, I., Wallis, J. D., Miller, E. K., Wagner, A. D., Silvia, A., ... Wagner, A. D. (2003). Neural Circuits Subserving the Retrieval and Maintenance of Abstract Rules. (2001), 3419–3428.
- Burgess, N., & Hitch, G. J. (2006). A revised model of short-term memory and long-term learning of verbal sequences. *Journal of Memory and Language*, *55*(4), 627–652.
<https://doi.org/10.1016/j.jml.2006.08.005>
- Burgess, P. W., Veitch, E., De Lacy Costello, A., & Shallice, T. (2000). The cognitive and neuroanatomical correlates of multitasking. *Neuropsychologia*, *38*(6), 848–863.
[https://doi.org/10.1016/S0028-3932\(99\)00134-7](https://doi.org/10.1016/S0028-3932(99)00134-7)
- Buschman, T. J., & Kastner, S. (2015). Perspective From Behavior to Neural Dynamics : An Integrated Theory of Attention. *Neuron*, *88*(1), 127–144.
<https://doi.org/10.1016/j.neuron.2015.09.017>
- Buschman, T. J., & Miller, E. K. (2007). of Attention in the Prefrontal and. *Science*, *315*(Ci), 1860–1862. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17395832>
- Buschman, T. J., Siegel, M., Roy, J. E., & Miller, E. K. (2011). Neural substrates of cognitive capacity limitations. *Proceedings of the National Academy of Sciences*, *108*(27), 11252–11255. <https://doi.org/10.1073/pnas.1104666108>
- Cavanna, A. E., & Trimble, M. R. (2006). The precuneus: a review of its functional anatomy and behavioural correlates. *Brain : A Journal of Neurology*, *129*(Pt 3), 564–583. <https://doi.org/10.1093/brain/awl004>
- Chafee, M. V., & Goldman-Rakic, P. S. (1998). Matching Patterns of Activity in Primate Prefrontal Area 8a and Parietal Area 7ip Neurons During a Spatial Working MemoryTask. *Journal of Neurophysiology*, *79*(6), 2919–2940.
<https://doi.org/10.1152/jn.1998.79.6.2919>
- Chamod, A. S., & Petrides, M. (2010). Dissociation within the frontoparietal network in verbal working memory: a parametric functional magnetic resonance imaging

- study. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 30(10), 3849–3856. <https://doi.org/10.1523/JNEUROSCI.0097-10.2010>
- Chatham, C. H., Frank, M. J., & Badre, D. (2014). Corticostriatal Output Gating during Selection from Working Memory. *Neuron*, 81(4), 930–942. <https://doi.org/10.1016/j.neuron.2014.01.002>
- Chen, X., Zirnsak, M., Vega, G. M., Govil, E., Lomber, S. G., & Moore, T. (2020). Parietal Cortex Regulates Visual Saliency and Saliency-Driven Behavior. *Neuron*, 106(1), 177–187.e4. <https://doi.org/10.1016/j.neuron.2020.01.016>
- Chiu, Y.-C., Esterman, M. S., Gmeindl, L., & Yantis, S. (2012). Tracking cognitive fluctuations with multivoxel pattern time course (MVPTC) analysis. *Neuropsychologia*, 50(4), 479–486. <https://doi.org/10.1016/j.neuropsychologia.2011.07.007>
- Choi, E. Y., Drayna, G. K., & Badre, D. (2018). Evidence for a Functional Hierarchy of Association Networks. *Journal of Cognitive Neuroscience*, 30(5), 722–736. https://doi.org/10.1162/jocn_a_01229
- Christophel, T. B., Jamshchinnina, P., Yan, C., Allefeld, C., & Haynes, J.-D. (2018). Cortical specialization for attended versus unattended working memory. *Nature Neuroscience*. <https://doi.org/10.1038/s41593-018-0094-4>
- Cohen, J. D., Barch, D. M., Carter, C., & Servan-Schreiber, D. (1999). Context-processing deficits in schizophrenia: Converging evidence from three theoretically motivated cognitive tasks. *Journal of Abnormal Psychology*, 108(1), 120–133. <https://doi.org/10.1037/0021-843X.108.1.120>
- Cole, M. W., Reynolds, J. R., Power, J. D., Repovs, G., Anticevic, A., & Braver, T. S. (2013). Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature Neuroscience*, 16(9), 1348–1355. <https://doi.org/10.1038/nn.3470>
- Cookson, S. L., & Schumacher, E. H. (2021). Dissociating the neural correlates of planning and executing hierarchical task sets. *IN PREP*.
- Corbetta, M., & Shulman, G. L. (2002). Control of Goal-Directed and Stimulus-Driven Attention in the Brain. *Nature Reviews Neuroscience*, 3(3), 215–229. <https://doi.org/10.1038/nrn755>
- Cowan, N., Li, D., Moffitt, A., Becker, T. M., Martin, E. A., Saults, J. S., & Christ, S. E. (2011). A Neural Region of Abstract Working Memory. *Journal of Cognitive Neuroscience*, 23(10), 2852–2863. <https://doi.org/10.1162/jocn.2011.21625>
- Cox J.S., R. W. . H. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29(29), 162–173. Retrieved from https://ac-els-cdn-com.ezp-prod1.hul.harvard.edu/S0010480996900142/1-s2.0-S0010480996900142-main.pdf?_tid=c22bae7a-b8f5-4a8b-9dac-91c1ed841d53&acdnat=1549393398_e37181b8933a2ac88c2d7dc0eab14413
- Crittenden, B. M., & Duncan, J. (2014). Task Difficulty Manipulation Reveals Multiple Demand Activity but no Frontal Lobe Hierarchy, (February), 532–540. <https://doi.org/10.1093/cercor/bhs333>
- Culham, J. C., & Valyear, K. F. (2006). Human parietal cortex in action. *Current Opinion in Neurobiology*, 16(2), 205–212. <https://doi.org/10.1016/j.conb.2006.03.005>
- Cutini, S., Scatturin, P., Menon, E., Bisiacchi, P. S., Gamberini, L., Zorzi, M., &

- Dell'Acqua, R. (2008). Selective activation of the superior frontal gyrus in task-switching: An event-related fNIRS study. *NeuroImage*, *42*(2), 945–955. <https://doi.org/10.1016/j.neuroimage.2008.05.013>
- D'Esposito, M., Aguirre, G. ., Zarahn, E., Ballard, D., Shin, R. ., & Lease, J. (1998). Functional MRI studies of spatial and nonspatial working memory. *Cognitive Brain Research*, *7*(1), 1–13. [https://doi.org/10.1016/S0926-6410\(98\)00004-4](https://doi.org/10.1016/S0926-6410(98)00004-4)
- D'Esposito, Mark, & Postle, B. R. (2015). The Cognitive Neuroscience of Working Memory. *Annual Review of Psychology*, *66*(1), 115–142. <https://doi.org/10.1146/annurev-psych-010814-015031>
- De Baene, W., Duyck, W., Brass, M., & Carreiras, M. (2015). Brain Circuit for Cognitive Control Is Shared by Task and Language Switching. *Journal of Cognitive Neuroscience*, *27*(9), 1752–1765. https://doi.org/10.1162/jocn_a_00817
- Dixon, M. L., De La Vega, A., Mills, C., Andrews-Hanna, J., Spreng, R. N., Cole, M. W., & Christoff, K. (2018). Heterogeneity within the frontoparietal control network and its relationship to the default and dorsal attention networks. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(7), E1598–E1607. <https://doi.org/10.1073/pnas.1715766115>
- Donoso, M., Collins, A. G. E., & Koechlin, E. (2014). Foundations of human reasoning in the prefrontal cortex. *Science*, *344*(6191), 1481–1486. <https://doi.org/10.1126/science.1252254>
- Dosenbach, N. U. F., Fair, D. A., Miezin, F. M., Cohen, A. L., Wenger, K. K., Dosenbach, R. A. T., ... Petersen, S. E. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proceedings of the National Academy of Sciences*, *104*(26), 11073–11078. <https://doi.org/10.1073/pnas.0704320104>
- Duncan, J., Emslie, H., Williams, P., Johnson, R., & Freer, C. (1996). Intelligence and the frontal lobe: The organization of goal-directed behavior. *Cognitive Psychology*, *30*(3), 257–303. <https://doi.org/10.1006/cogp.1996.0008>
- Eayrs, J. O., & Lavie, N. (2019). Individual differences in parietal and frontal cortex structure predict dissociable capacities for perception and cognitive control. *NeuroImage*, *202*(May), 116148. <https://doi.org/10.1016/j.neuroimage.2019.116148>
- Eichenbaum, A., Scimeca, J. M., & D'Esposito, M. (2020). Dissociable Neural Systems Support the Learning and Transfer of Hierarchical Control Structure. *The Journal of Neuroscience*, *li*(34), JN-RM-0847-20. <https://doi.org/10.1523/jneurosci.0847-20.2020>
- Erickson, D. T., & Kayser, A. S. (2013). The neural representation of sensorimotor transformations in a human perceptual decision making network. *NeuroImage*, *79C*, 340–350. <https://doi.org/10.1016/j.neuroimage.2013.04.085>
- Fedorenko, E., Duncan, J., & Kanwisher, N. (2013). Broad domain generality in focal regions of frontal and parietal cortex. *Proceedings of the National Academy of Sciences*, *110*(41), 16616–16621. <https://doi.org/10.1073/pnas.1315235110>
- Frank, M. J., & Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: Computational analysis. *Cerebral Cortex*, *22*(3), 509–526. <https://doi.org/10.1093/cercor/bhr114>
- Friedman, H., & Goldman-Rakic, P. (1994). Coactivation of prefrontal cortex and inferior parietal cortex in working memory tasks revealed by 2DG functional mapping in the rhesus monkey. *The Journal of Neuroscience*, *14*(5), 2775–2788.

- <https://doi.org/10.1523/JNEUROSCI.14-05-02775.1994>
- Gelastopoulos, A., Whittington, M. A., & Kopell, N. J. (2019). Parietal low beta rhythm provides a dynamical substrate for a working memory buffer. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(33), 16613–16620. <https://doi.org/10.1073/pnas.1902305116>
- Gillebert, C. R., 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*(6), 1694–1709. <https://doi.org/10.1093/brain/awr085>
- Gold, J. I., & Shadlen, M. N. (2003). The Influence of Behavioral Context on the Representation of a Perceptual Decision in Developing Oculomotor Commands. *The Journal of Neuroscience*, *23*(2), 632–651. <https://doi.org/10.1523/jneurosci.23-02-00632.2003>
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. <https://doi.org/10.1146/annurev.neuro.29.051605.113038>
- Goldman-Rakic, P. S. (1988). Topography of Cognition: Parallel Distributed Networks in Primate Association Cortex. *Annual Review of Neuroscience*, *11*(1), 137–156. <https://doi.org/10.1146/annurev.ne.11.030188.001033>
- Gosseries, O., Yu, Q., Larocque, J. J., Starrett, M. J., Rose, N. S., Cowan, N., & Postle, B. R. (2018). Parietal-occipital interactions underlying control-and representation-related processes in working memory for nonspatial visual features. *Journal of Neuroscience*, *38*(18), 4357–4366. <https://doi.org/10.1523/JNEUROSCI.2747-17.2018>
- Gratton, G., Coles, M. G. H., & Donchin, E. (1992). Optimizing the use of information: Strategic control of activation of responses. *Journal of Experimental Psychology: General*, *121*(4), 480–506. <https://doi.org/10.1037/0096-3445.121.4.480>
- Greicius, M. D., Supekar, K., Menon, V., & Dougherty, R. F. (2009). Resting-State Functional Connectivity Reflects Structural Connectivity in the Default Mode Network. *Cerebral Cortex*, *19*(1), 72–78. <https://doi.org/10.1093/cercor/bhn059>
- Grèzes, J., & Decety, J. (2000). Functional anatomy of execution, mental simulation, observation, and verb generation of actions: A meta-analysis. *Human Brain Mapping*, *12*(1), 1–19. [https://doi.org/10.1002/1097-0193\(200101\)12:1<1::AID-HBM10>3.0.CO;2-V](https://doi.org/10.1002/1097-0193(200101)12:1<1::AID-HBM10>3.0.CO;2-V)
- Hanks, T. D., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C., & Brody, C. D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*. <https://doi.org/10.1038/nature14066>
- Hazeltine, E., Lightman, E., Schwarb, H., & Schumacher, E. H. (2011). The Boundaries of Sequential Modulations: Evidence for Set-Level Control. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(6), 1898–1914. <https://doi.org/10.1037/a0024662>
- Heekeren, H. R., Marrett, S., Bandettini, P. A., & Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature*, *431*(7010), 859–862. <https://doi.org/10.1038/nature02966>
- Heekeren, H R, Marrett, S., Ruff, D. A., Bandettini, P. A., & Ungerleider, L. G. (2006). Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. *Proceedings of the National Academy*

- of Sciences of the United States of America*, 103(26), 10023–10028.
<https://doi.org/10.1073/pnas.0603949103>
- Heekeren, Hauke R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, 9(6), 467–479. <https://doi.org/10.1038/nrn2374>
- Iaria, G., Petrides, M., Dagher, A., Pike, B., & Bohbot, V. D. (2003). Cognitive Strategies Dependent on the Hippocampus and Caudate Nucleus in Human Navigation: Variability and Change with Practice. *The Journal of Neuroscience*, 23(13), 5945–5952. <https://doi.org/10.1523/JNEUROSCI.23-13-05945.2003>
- Jerde, T. A., & Curtis, C. E. (2013). Maps of space in human frontoparietal cortex. *Journal of Physiology-Paris*, 107(6), 510–516.
<https://doi.org/10.1016/j.jphysparis.2013.04.002>
- Karnath, H. O., & Perenin, M. T. (2005). Cortical control of visually guided reaching: Evidence from patients with optic ataxia. *Cerebral Cortex*, 15(10), 1561–1569.
<https://doi.org/10.1093/cercor/bhi034>
- Kastner, S., Chen, Q., Jeong S.K., Mruzek, R. E. B. (2017). A brief comparative review of primate posterior parietal cortex: a novel hypothesis on the human toolmaker. *Neuropsychologia*, 1–10. <https://doi.org/10.1109/EMBC.2016.7590696>
- Kayser, A. S., Buchsbaum, B. R., Erickson, D. T., & D’Esposito, M. (2010). The Functional Anatomy of a Perceptual Decision in the Human Brain. *Journal of Neurophysiology*, 103(3), 1179–1194. <https://doi.org/10.1152/jn.00364.2009>
- Kayser, A. S., & D’Esposito, M. (2013). Abstract rule learning: The differential effects of lesions in frontal cortex. *Cerebral Cortex*, 23(January), 230–240.
<https://doi.org/10.1093/cercor/bhs013>
- Kayser, A. S., Erickson, D. T., Buchsbaum, B. R., & D’Esposito, M. (2010). Neural representations of relevant and irrelevant features in perceptual decision making. *The Journal of Neuroscience*, 30(47), 15778–15789.
<https://doi.org/10.1523/JNEUROSCI.3163-10.2010>
- Kherif, F., Josse, G., Seghier, M. L., & Price, C. J. (2009). The main sources of intersubject variability in neuronal activation for reading aloud. *Journal of Cognitive Neuroscience*, 21(4), 654–668. <https://doi.org/10.1162/jocn.2009.21084>
- Kim, J. N., & Shadlen, M. N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature Neuroscience*, 2(2), 176–185.
<https://doi.org/10.1038/5739>
- Klingberg, T. (2010). Training and plasticity of working memory. *Trends in Cognitive Sciences*, 14(7), 317–324. <https://doi.org/10.1016/j.tics.2010.05.002>
- Koechlin, E., & Jubault, T. (2006). Broca’s Area and the Hierarchical Organization of Human Behavior. *Neuron*, 50(6), 963–974.
<https://doi.org/10.1016/j.neuron.2006.05.017>
- Koechlin, E., Ody, C., & Kouneiher, F. (2003). The Architecture of Cognitive Control in the Human Prefrontal Cortex. *Science*, 302(5648), 1181–1185.
<https://doi.org/10.1126/science.1088545>
- Konen, C. S., & Kastner, S. (2008). Representation of Eye Movements and Stimulus Motion in Topographically Organized Areas of Human Posterior Parietal Cortex. *Journal of Neuroscience*, 28(33), 8361–8375.
<https://doi.org/10.1523/JNEUROSCI.1930-08.2008>

- Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). Cost-benefit models as the next, best option for understanding subjective effort. *Behavioral and Brain Sciences*, 36(6), 707–726. <https://doi.org/10.1017/S0140525X13001532>
- LaRocque, J. J., Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2013). Decoding attended information in short-term memory: An EEG study. *Journal of Cognitive Neuroscience*, 25(1), 127–142. https://doi.org/10.1162/jocn_a_00305
- Lechner, H. A., Squire, L. R., & Byrne, J. H. (1999). 100 Years of Consolidation-Remembering Muller and Pilzecker. *Learning and Memory*, 77–87.
- Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012). Neural Evidence for a Distinction between Short-term Memory and the Focus of Attention. *Journal of Cognitive Neuroscience*, 24(1), 61–79. https://doi.org/10.1162/jocn_a_00140
- Lewis-Peacock, J. A., Kessler, Y., & Oberauer, K. (2018). The removal of information from working memory. *Annals of the New York Academy of Sciences*, 1424(1), 33–44. <https://doi.org/10.1111/nyas.13714>
- Lorenc, E. S., Sreenivasan, K. K., Nee, D. E., Vandenbroucke, A. R. E., & D’Esposito, M. (2018). Flexible coding of visual working memory representations during distraction. *Journal of Neuroscience*, 38(23), 5267–5276. <https://doi.org/10.1523/JNEUROSCI.3061-17.2018>
- Lorenc, E. S., Vandenbroucke, A. R. E., Nee, D. E., de Lange, F. P., & D’Esposito, M. (2020). Dissociable neural mechanisms underlie currently-relevant, future-relevant, and discarded working memory representations. *Scientific Reports*, 10(1), 1–17. <https://doi.org/10.1038/s41598-020-67634-x>
- Mackey, W. E., Devinsky, O., Doyle, W. K., Golfinos, J. G., & Curtis, C. E. (2016). Human parietal cortex lesions impact the precision of spatial working memory. *Journal of Neurophysiology*, 116(3), 1049–1054. <https://doi.org/10.1152/jn.00380.2016>
- Matsumoto, K., & Tanaka, K. (2004). Conflict and Cognitive Control. *Science*, 303(5660), 969–970. <https://doi.org/10.1126/science.1094733>
- McMenamin, B. W., Marsolek, C. J., Morseth, B. K., Speer, M. F., Burton, P. C., & Burgund, E. D. (2016). Conflicting demands of abstract and specific visual object processing resolved by frontoparietal networks. *Cognitive, Affective, & Behavioral Neuroscience*, 16(3), 502–515. <https://doi.org/10.3758/s13415-016-0409-4>
- Medina, J., Jax, S. A., & Coslett, H. B. (2020). Impairments in action and perception after right intraparietal damage. *Cortex*, 122(1988), 288–299. <https://doi.org/10.1016/j.cortex.2019.02.004>
- Mesulam, M. M. (1998). From sensation to cognition. *Brain*, 121(6), 1013–1052. <https://doi.org/10.1093/brain/121.6.1013>
- Miller, E. K., & Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, 24(1), 167–202. <https://doi.org/10.1146/annurev.neuro.24.1.167>
- Morrison, A. B., Rosenbaum, G. M., Fair, D., & Chein, J. M. (2016). Variation in strategy use across measures of verbal working memory. *Memory & Cognition*, 44(6), 922–936. <https://doi.org/10.3758/s13421-016-0608-9>
- Moussa, M. N., Steen, M. R., Laurienti, P. J., & Hayasaka, S. (2012). Consistency of

- Network Modules in Resting-State fMRI Connectome Data. *PLoS ONE*, 7(8).
<https://doi.org/10.1371/journal.pone.0044428>
- Muhle-Karbe, P. S., Andres, M., Brass, M., & Andres, M. (2014). Transcranial magnetic stimulation dissociates prefrontal and parietal contributions to task preparation. *Journal of Neuroscience*, 34(37), 12481–12489.
<https://doi.org/10.1523/JNEUROSCI.4931-13.2014>
- Murphy, A. P., Leopold, D. A., Humphreys, G. W., & Welchman, A. E. (2016). Lesions to right posterior parietal cortex impair visual depth perception from disparity but not motion cues. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1697). <https://doi.org/10.1098/rstb.2015.0263>
- Nee, D. E., & Brown, J. W. (2013). Dissociable frontal-striatal and frontal-parietal networks involved in updating hierarchical contexts in working memory. *Cerebral Cortex*, 23(9), 2146–2158. <https://doi.org/10.1093/cercor/bhs194>
- Nee, D. E., & D'Esposito, M. (2016). The hierarchical organization of the lateral prefrontal cortex. *ELife*, 5(MARCH2016), 1–26. <https://doi.org/10.7554/eLife.12112>
- Nee, D. E., & D'Esposito, M. (2017). Causal evidence for lateral prefrontal cortex dynamics supporting cognitive control. *ELife*, 6(Dcm), 1–19.
<https://doi.org/10.7554/eLife.28040>
- Nee, D. E., Jahn, A., & Brown, J. W. (2014). Prefrontal Cortex Organization: Dissociating Effects of Temporal Abstraction, Relational Abstraction, and Integration with fMRI. *Cerebral Cortex*, 24(9), 2377–2387.
<https://doi.org/10.1093/cercor/bht091>
- O'Shea, J., Sebastian, C., Boorman, E. D., Johansen-Berg, H., & Rushworth, M. F. S. (2007). Functional specificity of human premotor-motor cortical interactions during action selection. *European Journal of Neuroscience*, 26(7), 2085–2095.
<https://doi.org/10.1111/j.1460-9568.2007.05795.x>
- Oberauer, K. (2001). Removing irrelevant information from working memory: A cognitive aging study with the modified Sternberg task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(4), 948–957. <https://doi.org/10.1037/0278-7393.27.4.948>
- Oberauer, K. (2018). Removal of irrelevant information from working memory: Sometimes fast, sometimes slow, and sometimes not at all. *Annals of the New York Academy of Sciences*, 1424(1), 239–255. <https://doi.org/10.1111/nyas.13603>
- Oberauer, K., Lewandowsky, S., Farrell, S., Jarrold, C., & Greaves, M. (2012). Modeling working memory: An interference model of complex span. *Psychonomic Bulletin and Review*, 19(5), 779–819. <https://doi.org/10.3758/s13423-012-0272-4>
- Palenciano, A. F., González-García, C., Arco, J. E., Pessoa, L., & Ruz, M. (2019). Representational organization of novel task sets during proactive encoding. *Journal of Neuroscience*, 39(42), 8386–8397. <https://doi.org/10.1523/jneurosci.0725-19.2019>
- Pals, M., Stewart, T. C., Akyürek, E. G., & Borst, J. P. (2020). A functional spiking-neuron model of activity-silent working memory in humans based on calcium-mediated short-term synaptic plasticity. *PLOS Computational Biology*, 16(6), e1007936. <https://doi.org/10.1371/journal.pcbi.1007936>
- Pessoa, L., Gutierrez, E., Bandettini, P. A., & Ungerleider, L. G. (2002). Neural Correlates of Visual Working Memory. *Neuron*, 35(5), 975–987.

- [https://doi.org/10.1016/S0896-6273\(02\)00817-6](https://doi.org/10.1016/S0896-6273(02)00817-6)
- Petrides, M., Tomaiuolo, F., Yeterian, E. H., & Pandya, D. N. (2012). The prefrontal cortex: Comparative architectonic organization in the human and the macaque monkey brains. *Cortex*, *48*(1), 46–57. <https://doi.org/10.1016/j.cortex.2011.07.002>
- Pischedda, X. D., Go, X. K., Haynes, J., & Reverberi, X. C. (2017). Neural Representations of Hierarchical Rule Sets : The Human Control System Represents Rules Irrespective of the Hierarchical Level to Which They Belong, *37*(50), 12281–12296. <https://doi.org/10.1523/JNEUROSCI.3088-16.2017>
- Quentin, R., King, J.-R., Sallard, E., Fishman, N., Thompson, R., Buch, E. R., & Cohen, L. G. (2019). Differential Brain Mechanisms of Selection and Maintenance of Information during Working Memory. *The Journal of Neuroscience*, *39*(19), 3728–3740. <https://doi.org/10.1523/JNEUROSCI.2764-18.2019>
- Rahnev, D., Nee, D. E., Riddle, J., Larson, A. S., & D’Esposito, M. (2016). Causal evidence for frontal cortex organization for perceptual decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(21), 6059–6064. <https://doi.org/10.1073/pnas.1522551113>
- Ramanoël, S., Durteste, M., Bécu, M., Habas, C., & Arleo, A. (2020). Differential Brain Activity in Regions Linked to Visuospatial Processing During Landmark-Based Navigation in Young and Healthy Older Adults. *Frontiers in Human Neuroscience*, *14*(October), 1–16. <https://doi.org/10.3389/fnhum.2020.552111>
- Ranti, C., Chatham, C. H., & Badre, D. (2015). Parallel temporal dynamics in hierarchical cognitive control. *Cognition*, *142*, 205–229. <https://doi.org/10.1016/j.cognition.2015.05.003>
- Reynolds, J. R., Reilly, R. C. O., Cohen, J. D., & Braver, T. S. (2012). The Function and Organization of Lateral Prefrontal Cortex : A Test of Competing Hypotheses, *7*(2). <https://doi.org/10.1371/journal.pone.0030284>
- Rushworth, M. F. S., Hadland, K. A., Paus, T., & Sipila, P. K. (2002). Role of the human medial frontal cortex in task switching: A combined fMRI and TMS study. *Journal of Neurophysiology*, *87*(5), 2577–2592. <https://doi.org/10.1152/jn.2002.87.5.2577>
- Sahan, M. I., Sheldon, A. D., & Postle, B. R. (2019). The neural consequences of attentional prioritization of internal representations in visual working memory. *Journal of Cognitive Neuroscience*, *32*(5), 917–944. https://doi.org/10.1162/jocn_a_01517
- Saj, A., Verdon, V., Hauert, C. A., & Vuilleumier, P. (2018). Dissociable components of spatial neglect associated with frontal and parietal lesions. *Neuropsychologia*, *115*(March), 60–69. <https://doi.org/10.1016/j.neuropsychologia.2018.02.021>
- Sakai, K., & Passingham, R. E. (2006). Prefrontal set activity predicts rule-specific neural processing during subsequent cognitive performance. *Journal of Neuroscience*, *26*(4), 1211–1218. <https://doi.org/10.1523/JNEUROSCI.3887-05.2006>
- Sanfratello, L., Caprihan, A., Stephen, J. M., Knoefel, J. E., Adair, J. C., Qualls, C., ... Aine, C. J. (2014). Same task, different strategies: How brain networks can be influenced by memory strategy. *Human Brain Mapping*, *35*(10), 5127–5140. <https://doi.org/10.1002/hbm.22538>
- Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., ... Yeo, B. T. T. (2018). Local-Global Parcellation of the Human Cerebral Cortex from

- Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, 28(9), 3095–3114.
<https://doi.org/10.1093/cercor/bhx179>
- Schall, J. D. (2003). Neural correlates of decision processes: Neural and mental chronometry. *Current Opinion in Neurobiology*. [https://doi.org/10.1016/S0959-4388\(03\)00039-4](https://doi.org/10.1016/S0959-4388(03)00039-4)
- Shadlen, M. N., & Newsome, W. T. (2001). Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *Journal of Neurophysiology*, 86(4), 1916–1936. <https://doi.org/10.1152/jn.2001.86.4.1916>
- Shankar, S., & Kayser, A. S. (2017). Perceptual and categorical decision making : goal-relevant representation of two domains at different levels of abstraction, 2088–2103. <https://doi.org/10.1152/jn.00512.2016>
- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a Rational and Mechanistic Account of Mental Effort. *Annual Review of Neuroscience*, 40, 99–124. <https://doi.org/10.1146/annurev-neuro-072116-031526>
- Silver, M. A., & Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences*, 13(11), 488–495.
<https://doi.org/10.1016/j.tics.2009.08.005>
- Silver, M. A., Ress, D., & Heeger, D. J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *Journal of Neurophysiology*, 94(2), 1358–1371.
<https://doi.org/10.1152/jn.01316.2004>
- Smith, C. J., Xiong, G., Elkind, J. A., Putnam, B., & Cohen, A. S. (2015). Brain Injury Impairs Working Memory and Prefrontal Circuit Function. *Frontiers in Neurology*, 6(NOV), 1–13. <https://doi.org/10.3389/fneur.2015.00240>
- Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring Latent Visual Working Memory Representations in Human Cortex. *Neuron*, 91(3), 694–707.
<https://doi.org/10.1016/j.neuron.2016.07.006>
- Sternberg, S. (1966). High-Speed Scanning in Human Memory. *Science*, 153(3736), 652–654. <https://doi.org/10.1126/science.153.3736.652>
- Stokes, M. G. (2015). ‘Activity-silent’ working memory in prefrontal cortex: a dynamic coding framework. *Trends in Cognitive Sciences*, 19(7), 394–405.
<https://doi.org/10.1016/j.tics.2015.05.004>
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18(6), 643–662. <https://doi.org/10.1037/h0054651>
- Swinnen, S. P., Vangheluwe, S., Wagemans, J., Coxon, J. P., Goble, D. J., Van Impe, A., ... Wenderoth, N. (2010). Shared neural resources between left and right interlimb coordination skills: The neural substrate of abstract motor representations. *NeuroImage*, 49(3), 2570–2580. <https://doi.org/10.1016/j.neuroimage.2009.10.052>
- Székely, M., & Michael, J. (2020). The Sense of Effort: a Cost-Benefit Theory of the Phenomenology of Mental Effort. *Review of Philosophy and Psychology*.
<https://doi.org/10.1007/s13164-020-00512-7>
- Thiebaut de Schotten, M., Dell’Acqua, F., Valabregue, R., & Catani, M. (2012). Monkey to human comparative anatomy of the frontal lobe association tracts. *Cortex*, 48(1), 82–96. <https://doi.org/10.1016/j.cortex.2011.10.001>
- Toni, I. (1999). Signal-, Set- and Movement-related Activity in the Human Brain: An Event-related fMRI Study. *Cerebral Cortex*, 9(1), 35–49.

- <https://doi.org/10.1093/cercor/9.1.35>
- Unsworth, N., & Engle, R. W. (2007). The nature of individual differences in working memory capacity: Active maintenance in primary memory and controlled search from secondary memory. *Psychological Review*, *114*(1), 104–132.
<https://doi.org/10.1037/0033-295X.114.1.104>
- Vaidya, A. R., Jones, H. M., Castillo, J., & Badre, D. (2021). Neural representation of abstract task structure during generalization. *ELife*, *10*, 1–26.
<https://doi.org/10.7554/eLife.63226>
- Vernet, M., Quentin, R., Chanes, L., Mitsumasu, A., & Valero-Cabré, A. (2014). Frontal eye field, where art thou? Anatomy, function, and non-invasive manipulation of frontal regions involved in eye movements and associated cognitive operations. *Frontiers in Integrative Neuroscience*, *8*(AUG), 1–24.
<https://doi.org/10.3389/fnint.2014.00066>
- Wolfensteller, U., & von Cramon, D. Y. (2011). Strategy-effects in prefrontal cortex during learning of higher-order S–R rules. *NeuroImage*, *57*(2), 598–607.
<https://doi.org/10.1016/j.neuroimage.2011.04.048>
- Wolfensteller, U., & Von Cramon, D. Y. (2010). Bending the rules: Strategic behavioral differences are reflected in the brain. *Journal of Cognitive Neuroscience*, *22*(2), 278–291. <https://doi.org/10.1162/jocn.2009.21245>
- Xu, Y. (2009). Distinctive neural mechanisms supporting visual object individuation and identification. *Journal of Cognitive Neuroscience*, *21*(3), 511–518.
<https://doi.org/10.1162/jocn.2008.21024>
- Xu, Y. (2018). The Posterior Parietal Cortex in Adaptive Visual Processing. *Trends in Neurosciences*, *41*(11), 806–822. <https://doi.org/10.1016/j.tins.2018.07.012>
- Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., ... Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*, 1125–1165. <https://doi.org/10.1152/jn.00338.2011>.
- Zhang, J., Kriegeskorte, N., Carlin, J. D., & Rowe, J. B. (2013). Choosing the Rules: Distinct and Overlapping Frontoparietal Representations of Task Rules for Perceptual Decisions. *Journal of Neuroscience*, *33*(29), 11852–11862.
<https://doi.org/10.1523/JNEUROSCI.5193-12.2013>