

Decoding cognitive control in human parietal cortex

Michael Esterman¹, Yu-Chin Chiu, Benjamin J. Tamber-Rosenau, and Steven Yantis

Department of Psychological and Brain Sciences, Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218-2686

Edited by Richard M. Shiffrin, Indiana University, Bloomington, IN, and approved August 24, 2009 (received for review April 2, 2009)

Efficient execution of perceptual-motor tasks requires rapid voluntary reconfiguration of cognitive task sets as circumstances unfold. Such acts of cognitive control, which are thought to rely on a network of cortical regions in prefrontal and posterior parietal cortex, include voluntary shifts of attention among perceptual inputs or among memory representations, or switches between categorization or stimulus-response mapping rules. A critical unanswered question is whether task set shifts in these different domains are controlled by a common, domain-independent mechanism or by separate, domain-specific mechanisms. Recent studies have implicated a common region of medial superior parietal lobule (mSPL) as a domain-independent source of cognitive control during shifts between perceptual, mnemonic, and rule representations. Here, we use fMRI and event-related multivoxel pattern classification to show that spatial patterns of brain activity within mSPL reliably express which of several domains of cognitive control is at play on a moment-by-moment basis. Critically, these spatio-temporal brain patterns are stable over time within subjects tested several months apart and across a variety of tasks, including shifting visuospatial attention, switching categorization rules, and shifting attention in working memory.

fMRI | pattern classification | task switching | working memory

Effective goal-directed behavior requires organisms to manage and coordinate the perceptual, cognitive, and motor processes required by ongoing tasks. Efforts to unravel the neural circuits responsible for rapid and efficient cognitive control have revealed that regions of dorsolateral and dorsomedial prefrontal cortex and posterior parietal cortex subserved different aspects of control, including the maintenance of working memory representations of task sets (1, 2), detection and management of perceptual, cognitive, or response conflict (3) and the initiation of task set reconfiguration (4, 5).

Frequent reconfiguration of the mind/brain in one or more domains is required as goals and environmental demands change. These domains may include, for example, states of attention that select task-relevant sensory or mnemonic information; categorization rules that can be applied to an attended perceptual input; and/or stimulus-response mapping rules that specify a task-appropriate response after perception and categorization. Transitions between task sets (e.g., shifting attention or switching categorization rule) are accomplished through voluntary acts of cognitive control. The psychological and neurophysiological basis of task set shifting has been investigated with behavioral and neuroimaging techniques (4–8).

A critical unresolved question concerns whether task set reconfiguration in different cognitive domains is associated with multiple domain-specific control mechanisms or with a common, domain-independent control mechanism. According to the domain-specific account, distinct dedicated brain regions control task set shifts in different cognitive domains (9). According to the domain-independent account, a common cortical source of control signals initiates shifts in multiple (or all) cognitive domains (10). Shifting task set across several domains recruits a network of brain regions including the prefrontal cortex, superior parietal cortex, and the basal ganglia (4–6, 10). Studies have revealed at least partially distinct neural substrates for different domains of cognitive control, such as the intraparietal sulcus and frontal eye fields for shifts of

visuospatial attention (5, 10, 11), and prefrontal cortex for switching categorization rule (4, 6, 12, 13). Thus, a third possibility is that a common cortical locus initiates task set shifts within many domains of control by coordinating activity in brain regions that are specialized for representing and processing information within specific domains.

Several recent functional neuroimaging studies have reported that voluntary shifts of attention in multiple sensory domains are associated with transient shift-evoked activity in the medial wall of the superior parietal lobule (mSPL) and superior precuneus (5, 11, 14–17). This brain activity is time-locked to the initiation of the shift and is transient, suggesting that it plays a unique role in reconfiguring, rather than maintaining the state of attention. A recent meta-analysis of studies investigating different domains of control suggested that a common set of cortical regions, including posterior parietal cortex, tend to be associated with cognitive control (broadly construed) in each studied domain (18). Individual differences in behavioral switching performance correlate across tasks requiring both internal (mnemonic) and external (perceptual) loci of attentional control (19). Together, these lines of evidence implicate a domain-independent transient reconfiguration process associated with a neural signature that is consistently observed in mSPL. A recent neurophysiological investigation revealed that neurons in macaque area 7a, in the posterior parietal cortex, responded transiently to categorization rule reconfigurations (20); this finding is qualitatively similar to the signal observed in human mSPL with fMRI (e.g., ref. 10). However, few studies have investigated multiple domains of control simultaneously. One recent study that directly compared the neural basis of cognitive control in multiple domains reported that activation in a common region in mSPL was associated with both spatial shifts of attention and switches of categorization rule (10); no other brain region exhibited this domain-independent pattern of activity.

The functional properties of mSPL are not well understood. This cortical region is virtually absent from the neurophysiology or human focal lesion literature (21). Furthermore, although several functional imaging studies have implicated mSPL as a domain-independent locus of cognitive control, other interpretations of this activity are possible. For example, mSPL may contain distinct domain-specific subpopulations of neurons that are recruited for different types of reconfigurations (10), a pattern that conventional univariate statistical techniques may fail to detect. No neurophysiological data have been collected to address this question.

Methods in multivariate analyses and machine learning have recently been deployed to investigate the voxelwise spatial pattern of activity within functional regions of interest (22). Multivoxel pattern classification (MVPC) applied to BOLD fMRI data supports inferences about functional organization at the subregion level, as follows. If a given cortical region contains two or more intermixed subpopulations of neurons with distinct functional se-

Author contributions: Y.-C.C., B.J.T.-R., and S.Y. designed research; Y.-C.C. and B.J.T.-R. performed research; M.E. and Y.-C.C. contributed new reagents/analytic tools; M.E., Y.-C.C., and B.J.T.-R. analyzed data; and M.E. and S.Y. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: esterman@jhu.edu.

This article contains supporting information online at www.pnas.org/cgi/content/full/0903593106/DCSupplemental.

lectivity, any given voxel will contain a mixture of neurons from these different populations. Each voxel's activity will reflect any slight over-representation of one subpopulation of neurons. Slight differences in functional selectivity between voxels can be detected as spatial patterns of activity across a region of interest unique to each cognitive or perceptual state.

This logic has been successfully applied to decode the contents of perceptual representations, including orientation selectivity (23, 24) and motion selectivity (25, 26) in visual cortex, and individual object selectivity in the lateral occipital complex (27). Other studies have used MVPC to link perceptual representations to those held and retrieved in memory (28, 29), with a particular focus on sensory cortices. More abstract representations of task rules have also been decoded in prefrontal cortex (30). Here, we report an application of MVPC to examine the underlying neural basis in human mSPL of a cognitive process—the initiation of task set switches—rather than a perceptual or mnemonic representation. We analyzed data from two recent studies in which subjects were required to apply cognitive control in two distinct domains: either (i) shifting visual attention between objects in different spatial locations and switching between two digit categorization rules (10); or (ii) shifting visual attention between objects in external space and shifting internal attention between elements held in working memory (31). Within regions of mSPL identified (using conventional univariate general linear model conjunction analyses) as exhibiting domain-independent cortical activity, we report evidence for transient and distinct spatiotemporal patterns of activity for different domains of cognitive control, within the same cortical region. Pattern classification was carried out in an event-related fashion at each time point surrounding a task shift. Critically, these patterns are stable over time within subjects tested several months apart and support reliable neural decoding of both spatial and nonspatial acts of control.

Results

Experiment 1: Spatial Shifts of Attention and Categorization Rule Switching. We investigated the role of mSPL in shifting attention or switching categorization rule with event-related fMRI. Participants viewed a computer display containing multiple rapidly changing visual streams of alphanumeric characters. Instructional cues and target digits were embedded in streams of distractor letters (Fig. 1*A*, see *Methods*). At any given moment, the participant occupied one of four preparatory states: they covertly attended to either the left or right rapid serial visual presentation (RSVP) stream and were prepared to perform either a parity (odd/even) or magnitude (high/low) digit categorization task. Cues instructed participants to either shift spatial attention between the left and right visual streams while continuously fixating the center of the screen, or to maintain the current state of attention and categorization rule. Univariate multiple regression analyses of the neuroimaging data revealed a common region of mSPL (Brodmann area 7) that exhibited transient activity in response to shift cues of either type compared to hold cues. The results of these univariate analyses have been reported elsewhere (10).

Decoding Reconfiguration Signals Across Domains of Cognitive Control. The domain-independent region of mSPL was defined separately for each subject using data from the remaining 15 subjects (see *Methods* for details of voxel selection). Pattern classification was carried out separately for each subject using a linear support vector machine (SVM) learning algorithm and a leave-one-out cross-validation procedure (see *Methods*). First, BOLD responses of individual voxels were z-transformed relative to the entire time course within each run to minimize differences across runs. We focused on the BOLD response at individual time points (i.e., TRs) ranging from 2 s before the cue to 12 s after the cue. Separate classifiers were trained and tested for each TR. This enabled us to

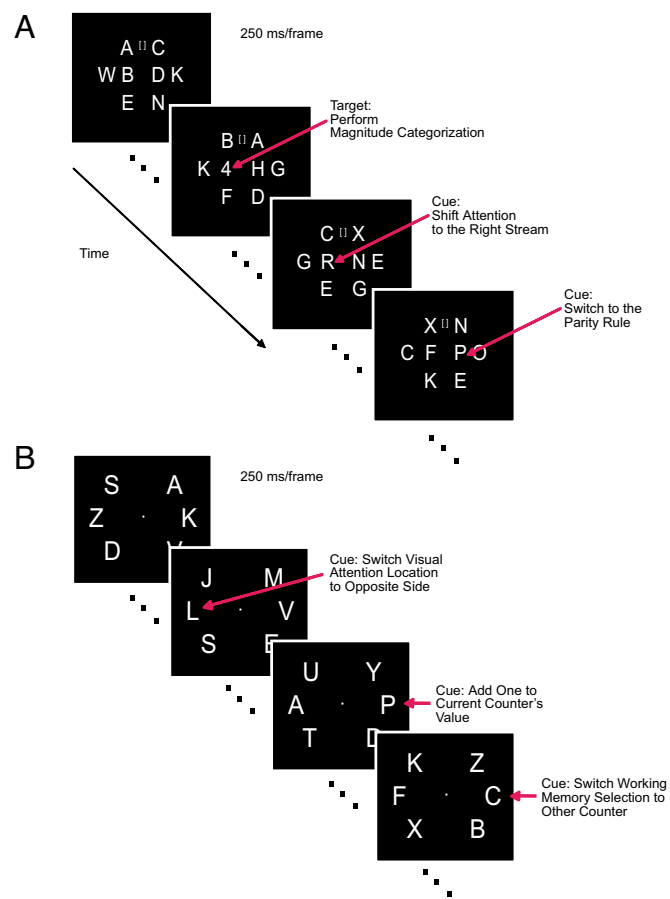


Fig. 1. Behavioral paradigms. (A) Experiment 1. In this example, subjects were instructed to begin by attending to the left central RSVP stream and to prepare for the magnitude Task. When a digit (e.g., “4”) is presented, they perform the high/low judgment on this digit. They shift attention covertly to the right central RSVP stream when the “R” cue is presented in the attended stream and continue to prepare for the magnitude task. When the “P” cue is presented in the attended stream, they switch to the parity task and wait for a digit to perform the odd/even judgment. [Reproduced with permission from Chiu Y-C, Yantis S (A domain-independent source of cognitive control for task sets: Shifting spatial attention and switching categorization rules) (Copyright 2009, *J Neurosci*).] (B) Experiment 2. In this example, subjects were instructed to begin by attending to the left central RSVP stream. The “L” cue (“location”) instructs them to shift spatial attention to the right central RSVP stream. The “P” cue (“plus”) indicates the currently active counter should be incremented. When the “C” cue (“counter”) is presented, they shift internal attention to the other counter.

explore classification performance across time and to identify the most informative time points in the BOLD signal. Given the delay of the hemodynamic response, we anticipated that classification performance would peak ≈ 4 –8 s after cue onset.

Our primary goal was to determine whether different spatial patterns of cortical activity were evoked by switches of categorization rule (magnitude-to-parity or parity-to-magnitude) compared to shifts of spatial attention (left-to-right or right-to-left), to test the hypothesis that different neuronal subpopulations respond selectively to different domains of task set reconfiguration. Classification accuracy for categorization rule shifts versus spatial attention shifts was significantly greater than chance at 4–10 s after the cue to switch (Fig. 2; for time points 4 and 10 s, $t(15) = 2.99$ and 3.1, $P < 0.05$, respectively; for time points 6 and 8 s, $t(15) = 5.10$ and 5.52, $P < 0.001$, respectively; two-tailed t test vs. chance performance of 50%). Averaging the magnitude of the BOLD signal across the two peak time points of 6–8 s led to a classification rate of 60.3% [$t(15) = 5.52$, $P < 0.0001$; see Fig. S1 for individual subject results.]




Fig. 2. Event-related classification performance in Experiment 1: decoding attention shift vs. categorization rule switch. Classification performance is presented for each time point relative to the onset of a switch cue, based on voxels in medial SPL (see *Methods*). *, $P < 0.05$; **, $P < 0.001$.

MVPC supports inferences about the relative contribution, or importance, of individual voxels to successful classification (22). This allows us to explore the response properties of the most important voxels, as well as the spatial pattern of voxel selectivity. To this end, mSPL voxels were rank-ordered according to how much they contributed to the classification accuracy (using the SVM weights derived from classifying the mean of the BOLD signal at time points 6 and 8 s after the relevant shift cue), separately for each subject (see *SI Methods*). MVPC was carried out repeatedly within subsets of voxels, starting with a subset of the two most important voxels and then with increasingly large subsets of the N most important voxels ($n = 2, 10, 25, 50, 100, 150, 200,$ and 300). Classifier performance peaked when MVPC was restricted to the most important 100 voxels, and sharply decreased with 200 voxels or more (Fig. S2). The reduction in classification performance that emerges when increasing numbers of voxels are added reflects the addition of low-importance voxels that contain noise.

Approximately 54% of the 100 voxels that yielded the best classification performance were assigned positive weights (i.e., when these voxels were highly active, they contributed to an ‘attention shift’ classification) and 46% were assigned negative weights (i.e., when these voxels were highly active, they contributed to a ‘rule shift’ classification). To visualize the response properties of these voxels, the mean shift-cue-evoked BOLD response was examined. The distribution of the most discriminating voxels reveals subpopulations that respond preferentially to attention shifts and

rule switches, respectively (Fig. 3A shows a single representative subject). This illustrates that the classifier is relying on mSPL voxels that are recruited selectively for each type of cognitive reconfiguration. The spatial pattern of a subset of these 100 voxels (on a single slice plane) is displayed in Fig. 3B for four representative subjects.

Pattern classification reveals that the mSPL contains voxels that respond selectively during acts of control in different cognitive domains. The overall mean BOLD response of mSPL was greater after attention shift cues than after rule switch cues [$t(15) = 4.23$, $P < 0.01$], despite the fact that the distributions of peak BOLD response magnitudes for each type of shift were highly overlapping (see Figs. S3 and S4). A large mean difference in activity between conditions within an area could provide a partial basis for accurate classification. We therefore took several steps to determine whether the observed mean difference was a significant source of classification accuracy.

First, we compared multivoxel classification to classification performance based only on the mean activity within the entire ROI (in effect, treating it as a single feature). Classification performance was significantly worse than MVPC (60.3% for MVPC vs. 54.5% for the mean activity, paired t test, $P < 0.01$). Removing the mean difference (mean-centering) by normalizing rule- and attention-shift BOLD signal (see *Methods*) (24) had no significant effect on classifier performance (60.9% correct classification with and 60.3% without normalization). Together, these controls validate the sensitivity of multivariate pattern analyses in detecting different voxel patterns for different reconfiguration signals and reveal the pattern specificity of the transient signal in mSPL after acts of task set reconfiguration.

Decoding Reconfiguration Signals Within Each Domain. In addition to assessing the degree to which MVPC could accurately classify which of two different domains of cognitive control was at play (spatial attention shift versus rule switch), we also examined the degree to which MVPC could accurately classify different types of shifts within a single domain. The SVM classifier successfully decoded shifts of spatial attention from left-to-right vs. right-to-left (59.5% correct overall, $P < 0.0001$). However, MVPC failed to correctly classify shifts between the two digit categorization rules (shifts from the magnitude categorization rule to the parity rule and vice versa; 49.8% correct classification overall). This suggests that the neurons in mSPL encode the direction of spatial shifting but not necessarily the nature of the specific task rule reconfiguration. However, it is also possible that mSPL voxels represent the cue location or the location to which attention is directed (opposite of the cue), rather than the direction of the shift itself. These are not mutually exclusive

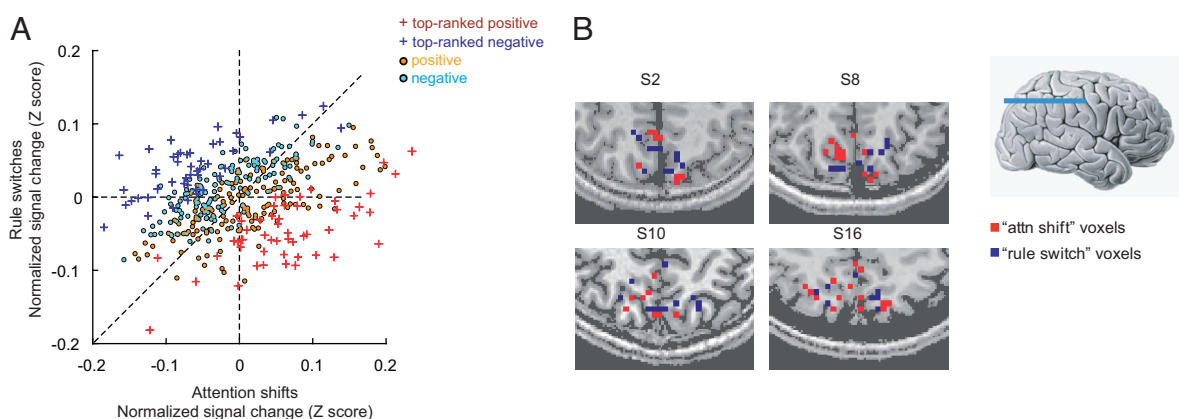


Fig. 3. Illustration of response properties and spatial pattern of important voxels (A) Scatterplot (data from Subject 8) of mean normalized BOLD response evoked by attention shift or rule switch cues in each of 402 mSPL voxels. Plus signs are the 100 most important voxels, and circles are the rest of the voxels. Dotted line indicates equivalent relative response to attention and rule reconfigurations. (B) Voxels projected on brain by importance value, in four subjects.

alternatives, and the resolution of fMRI and current experimental design does not allow us to distinguish between these alternatives.

The voxels selective for left-to-right versus right-to-left shifts did not correlate with left-right position in the brain (correlation of importance values and left-right position, $r = 0.01$), which constitutes evidence against a lateralized organization in this medial brain region. Furthermore, the classification of shift direction is not entirely driven by the highest magnitude “attention-weighted voxels” from the attention shifts vs. rule shifts classifier. Specifically, the most important 50 or 100 positively weighted “attention shift voxels” classified direction of shift (left-to-right vs. right-to-left) no better than the most important 50 or 100 negatively weighted “rule shift voxels” (50 voxels: 55.0% vs. 54.5%; 100 voxels: 57.5% vs. 54.8%; $P > 0.1$). However, among these subsets of voxels, only classification with the best 100 “attention shift voxels” (57.5%) was not significantly worse than classification using all SPL voxels (59.5% accuracy). This leaves open the possibility that “attention shift voxels” preferentially contain direction-specific information.

Experiment 2: Spatial Shifts of Attention and Switching in Working Memory.

The degree to which these spatial patterns are consistent within subjects over time and across a range of tasks is of critical importance to their interpretation. One possibility is that these patterns are entirely task- and context-specific, such that changes in stimulus properties or configuration, for example, would alter this neural signature. Alternatively, these patterns may reflect stable properties of the neural subpopulations in different dimensions of cognitive control. In Experiment 1, the switches were spatial or nonspatial (i.e., rule) and involved an external or an internal locus of selection. To examine the stability of the multivoxel patterns within subjects across time and context, we examined data from a second experiment in which some of the participants in Experiment 1 subsequently participated (31).

Participants in Experiment 2 had to maintain two different counters in working memory (32) while viewing multiple rapidly changing RSVP streams (Fig. 1B, see *Methods*). One counter was “active” and the other “inactive,” but both counter values had to be maintained in WM at all times. Instructional cues appeared in the attended stream, which required the subject to increment the active counter, shift internal attention to the currently inactive counter (i.e., make it active), shift visuospatial attention to the other RSVP stream, or hold the current counter values, activation states, and focus of attention. A counter switch operation is both a nonspatial switch and one with an internal locus in WM rather than in external perceptual space. If the switch-related spatial patterns in Experiment 1 are able to accurately predict the switches in Experiment 2 (within an individual subject), it would indicate robust stability in the underlying neural subpopulations along some common dimension(s).

We first used SVM to train a classifier to discriminate between spatial attention and mnemonic shifts (Experiment 2) and found that the classifier was able to decode 70% of shifts correctly as either spatial attention or WM shifts at the peak time point of 6 s after the cue [$t(2) = 5.87, P < 0.05$; Fig. 4], using the same region and voxels as identified in Experiment 1 (see *Methods*). Thus, spatial attention shifts and working memory switches exhibited distinct spatial patterns of cortical activity in Experiment 2. In addition, we could classify shifts of left-to-right vs. right-to-left (77%, $P < 0.05$), but could not classify the two types of WM shifts (54%, $P > 0.2$).

To test the stability and generalizability of the multivoxel patterns, using the data from each of the three subjects who participated in both experiments, a classifier was trained to predict shifts of spatial attention vs. shifts of digit categorization rules using the patterns of activity in mSPL collected during Experiment 1, and was then tested with the data obtained from the same subject in Experiment 2. We decoded across experiments with the BOLD response at 6 s, since it represented the most informative time point for both experiments (see Figs. 2 and 4). Within-subject random-




Fig. 4. Event-related classification performance in Experiment 2: decoding attention shift vs. working memory switch. Classification performance is presented for each time point relative to the onset of a switch cue, based on voxels in medial SPL (see *Methods*). *, $P < 0.05$; **, $P < 0.001$.

ization tests were performed to evaluate statistical significance with an alpha of .05 (see *SI Methods*). The cross-experiment classification performance was significantly greater than chance (59.7%) for all three subjects (Fig. 5, Experiment 1- \rightarrow Experiment 2). The reverse classification (training on Experiment 2 and testing on Experiment 1) was also significantly greater than chance for two of three participants (56.8%, Fig. 5, Experiment 2- \rightarrow Experiment 1). In Subject 3, who failed to classify from Experiment 2 to Experiment 1 above chance (52.1%), within-Experiment 1 classification was the weakest of the three subjects (55.6%). Overall, the distribution of the domain-specific neural subpopulations recruited for task set reconfiguration exhibits robust stability within subjects across time, as five of six attempts to classify across experiments were significantly better than chance. This pattern of results is preserved when accounting for any mean-difference across the two conditions (see Fig. S5).

Discussion

We investigated the neural basis of task set reconfiguration in three distinct domains. Transient neural activity in mSPL is uniquely




Fig. 5. Within and across experiment classification. In all three subjects, classification performance is significantly greater than chance (red lines indicate the critical value from a randomization test, $P = 0.05$; see *SI Methods*) for classification within Experiment 1 (light green) and within Experiment 2 (light blue). When data from Experiment 1 were used to train a classifier to predict switches in Experiment 2, decoding was significantly greater than chance in all three subjects (dark blue). When data from Experiment 2 were used to train a classifier to predict switches in Experiment 1, decoding was significantly greater than chance in subjects A and B, but not C (dark green).

associated with reconfiguration in multiple domains, including shifts of external attention between objects, locations, and features; shifts of attention between WM representations; and shifts of categorization rule (5, 10, 11, 14–17, 31). These results suggest that mSPL plays a domain-independent role in the initiation of task set reconfiguration. However, the conventional univariate statistical analyses used in these studies left open the possibility that mSPL contains two or more neural subpopulations that exhibit domain-specific patterns of activity. We used event-related multivoxel pattern classification (er-MVPC) to reliably decode shifts of spatial attention, switches of categorization rule, and switches between working memory representations. MVPC extracts information at a subregion scale of functional cortical organization and supports the inference that different neural populations respond selectively to different acts of cognitive control within this confined region of parietal cortex.

The spatiotemporal patterns we observed are stable over time and generalize across different types of reconfigurations. In Experiment 1, a classifier trained on attention shifts vs. categorization rule switches successfully classified attention shifts vs. working memory switches in the same subjects participating in Experiment 2 several months later. The functions of these neural subpopulations may vary along several possible dimensions. Given that both experiments examined spatial shifts of attention vs. another domain, the critical dimension could potentially be whether the reconfiguration is spatial or nonspatial. Although the importance of SPL in all manner of switches, including nonspatial ones, has been amply documented (5, 11, 14–17), spatial information may be specially represented in the parietal cortex (33, 34), and this explanation is bolstered by the accurate classification of shift direction (left-to-right vs. right-to-left), together with the failure to classify the nonspatial domains (e.g., magnitude-to-parity vs. parity-to-magnitude), as well as the overall mean activity difference between attention shift evoked response and rule shift evoked response. According to this explanation, classification, especially in the cross-experiment case, may rely on voxels that are either tuned to spatial shifts of attention, or “other,” nonspatial cognitive shifts.

Another possible axis of organization is external versus internal locus of selection. It has been demonstrated that even within the spatial domain, attention can be oriented to the external world (perceptually) or to internal representations in spatial working memory, and that these acts of control recruit only partially overlapping activity in the parietal and frontal cortex (35). In our experiments, shifting of counter selection and rule selection both require a change in the internal focus of attention to purely mental representations (counters or rules), while the spatial task requires attention toward external perceptual inputs that are continuously present (left or right visual stream). Therefore, the counter and rule switches may both require a change in priority of goals held in working memory, thought to be represented in prefrontal cortex (1, 2), compared to shifts of spatial attention toward perceptual inputs, which require change in top-down modulations in ventral visual cortex (5).

The spatial vs. nonspatial and external vs. internal possibilities are not mutually exclusive, and future studies considering the multiple internal or external domains of control within the same experiment will be required to further clarify these functional axes. In the first case (spatial vs. “other”), nonspatial switches of feature- or object-based attention should evoke patterns of activity that are similar to those evoked by categorization rule and memory switches. The second possibility (internal vs. external) would predict that spatial and feature- or object-based attention switches would evoke similar patterns of activity, because their locus is also external.

While cross-experiment decoding was significantly better than chance in most cases, it was not as accurate as within-experiment decoding. Several factors contribute to the attenuation of classification performance across experiments, including the possibility that the different contexts of these two experiments did yield

different multivoxel patterns and that the patterns changed with learning and experience in the task. Nevertheless, it is striking that classification was successful even in the face of the effective spatial smoothing that results from imperfections in coregistration of multiple fMRI sessions. The degree to which plasticity and learning influence the neural recruitment within mSPL remains unknown.

A challenge in the current study was that classification of transient shift events was attempted in a rapid event-related fMRI design. Most previous applications of MVPC have used blocked designs in which a single “trial” is the average BOLD response over several seconds of repeated stimulation (24) or slow event-related designs in which the trials are separated in time to prevent temporal overlap in the sluggish BOLD hemodynamic response (refs. 30 and 36; however, see ref. 28 for a notable exception and see ref. 37, which used a novel classification method). Successful application of er-MVPC in the face of a transient and non-deconvolved BOLD signal suggests the possibility of online, real-time decoding of subjects’ intentions and cognitive states during more naturalistic and overlapping cognitive demands.

Some voxels are assigned small weights (and therefore may be said to have low importance) in classification. There are several possible functional interpretations of these low-importance voxels. Voxels that contain redundant information (i.e., that are highly correlated with other voxels) may be assigned low weights. However, at least some voxels have little or no selectivity to these types of shifts (see Fig. 3A), and many only serve to impair classification performance (see Fig. S2). These unselective voxels may contain approximately equal numbers of neurons with selectivity to each type of shift, leading to no net selectivity. Alternatively, they could be dominated by neurons that are truly domain-independent. Last, they could contain neurons that are selective for other domains of control not explored in the current experiments.

While these results reveal that content of mSPL activity is process-specific, it is not yet known how this region communicates with other brain regions thought to be part of a cognitive control network, including lateral parietal and prefrontal cortex, and possibly subcortical structures such as the basal ganglia. EEG measurements have shown that, at least for spatial shifts of attention, posterior parietal activity precedes activity in the frontal cortex, suggesting a critical role for this region in the initiation of task set reconfiguration (38). Functional connectivity may reveal how the unique patterns of activity observed in the current study are related to domain-specific cortical activity. For example, Chiu and Yantis (10) found that spatial shifts of attention recruited mSPL and the inferior parietal cortex and frontal eye fields, a result that echoes other studies of visual spatial attention (11, 39). Conversely, they found that shifts of categorization rule were associated with activity in mSPL and in the left IPS, a region associated with alphanumeric categorization task switching (40). This analysis, like most fMRI analyses, does not support the inference of causality or necessity of this parietal region in acts of cognitive control. Because lesions of mSPL are so rare (21), cognitive neuropsychological methods are not well-suited to answer this question. Transcranial magnetic stimulation offers one possible method that could be used to explore the necessity and temporal properties of this region for cognitive reconfigurations.

Unlike other applications of MVPC that decode representational contents of perception, attention, and memory, we examined the spatial patterns of brain activity associated with different cognitive processes, independent of their representational content. This application could be valuable for other fMRI research investigating the neural organization of a specified brain region involved in cognitive control, such the rostral-caudate gradients in the prefrontal cortex (e.g., ref. 41).

In summary, we used multivoxel pattern classification to show that mSPL supports the initiation of task set reconfiguration in multiple domains. Er-MVPC revealed that there are separable voxel-wise patterns of activity within mSPL evoked by different acts

of cognitive control. Decoding switches across multiple experiments conducted several months apart established the stability of these subpopulations across time and context. Together, these results support an emerging understanding of the neural basis of cognitive control involving both domain-specific and domain-independent mechanisms of task set reconfiguration.

Methods

Behavioral Task: Experiment 1. Subjects maintained central fixation while viewing two rapid serial visual presentation (RSVP) target streams of alphanumeric characters (4 Hz), each flanked by three distractor streams (Fig. 1A). At any given moment, the subject covertly attended to either the left or right stream and was prepared to categorize target digits that appeared in the attended stream according to one of two rules: odd vs. even (Parity Task) or high (6–9) vs. low (2–5) (Magnitude Task). Forty-eight critical events were randomly intermixed among filler items (non-cue letters) in each run; 24 target digits (2–9) and 24 cue letters ('L,' 'R,' 'M,' and 'P'). The cues respectively specified "attend Left," "attend Right," "prepare Magnitude Task," and "prepare Parity Task." See ref. 10 and *SI Methods* for further details.

Behavioral Task: Experiment 2. The display was similar to that in Experiment 1 (Fig. 1B). Subjects attended the left or right target stream, and maintained two counters in working memory, one active and ready to be incremented and the other inactive. The four cue letters were 'L,' 'C,' 'P,' and 'H' that respectively specified "shift the locus of spatial attention" "switch from one WM counter to the other," "increment the selected WM counter's value (plus)," and "hold the current state." At the end of the run, subjects verbally reported the values of the two counters. See *SI Methods* for further details.

mSPL Voxel Selection. The unsmoothed images were Talairach-transformed and resampled into 3-mm isotropic voxels for voxel selection. Sixteen separate random-effects group GLMs were performed, leaving each subject out. For each of these GLMs, a conjunction analysis (attSh>attHd AND rulSw>rulHd) defined an ROI in mSPL. For the purposes of voxel selection, the threshold was reduced to include a larger set of voxels with potentially weaker selectivity at the group level

($t = 2.00, P < 0.06$). This region was subsequently projected back into the native space of the subject left out of that GLM. Thus, the ROI selection process was independent of the data that were subsequently classified. We verified that these ROIs were still attention- and rule shift-selective (versus their respective hold trials; $t_{15} = 6.47$ and $3.66, P < 0.01$). These voxels were used for Experiment 2 as well.

Classification with Support Vector Machines. Decoding analyses were performed using OSU SVM toolbox (adaptation of libsvm: <http://www.csie.ntu.edu.tw/~cjlin/libsvm/>) for MATLAB and in-house MATLAB code. All analyses used linear SVM. The signal in each voxel was z-transformed for each of the runs. Initially, each TR after attSh and rulSw cues (WMSw for Experiment 2) was classified separately. The event-related fMRI signal associated with attSh and rulSw/WMSw cues starting 2 s before the event until 12 s after the event (i.e., TRs –1 to 6) was used for classification, as was the mean of the BOLD signal at 6 and 8 s after the event for follow-up analyses. In the mean-centered analysis, the mean of all of the voxels for each type of switch (attention or categorization rule) was subtracted from every voxel on each trial. This resulted in the overall voxel population mean being equated for the two types of shifts.

To estimate within-experiment classification rates, a leave-one-run-out cross-validation scheme was used. For cross-experiment decoding, all Experiment 1 or 2 data from a given subject were used as the training set, while Experiment 2 or 1 data, respectively, from that subject served as the testing set. The time point selected for cross-experiment SVM input was 6 s postcue. This was based on the observation that the event-related classification accuracy was greatest for that TR across the two experiments (Figs. 2 and 4).

Importance Map. With linear SVM, the absolute value of each voxel's weight provides a quantitative index of the importance of that voxel in the decision function (22); the most discriminating voxel has the largest absolute weight (See *SI Methods* for more details).

ACKNOWLEDGMENTS. We thank M. Stokes and K. Norman for helpful discussion regarding pattern classification analyses; J. Serences for comments on an earlier version of the manuscript; and H. Rosen for excellent technical support. This work was funded by U.S. National Institutes of Health Grant R01-DA13165 (to S.Y.).

- D'Esposito M, et al. (1995) The neural basis of the central executive system of working memory. *Nature* 378:279–281.
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Ann Rev Neurosci* 24:167–202.
- Carter CS, et al. (1998) Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science* 280:747–749.
- Braver TS, Reynolds JR, Donaldson DI (2003) Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron* 39:713–726.
- Yantis S, Schwarzbach J, Serences J, Carlson R, Steinmetz M (2002) Transient neural activity in human parietal cortex during spatial attention shifts. *Nat Neurosci* 5:995–1002.
- Cools R, Clark L, Robbins TW (2004) Differential responses in human striatum and prefrontal cortex to changes in object and rule relevance. *J Neurosci* 24:1129–1135.
- Monsell S (2003) Task switching. *Trends Cognit Sci* 7:134–140.
- Sudevan P, Taylor D (1987) The cuing and priming of cognitive operations. *J Exp Psychol: Hum Percept Perform* 13:89–103.
- Rushworth M, Paus T, Sipila P (2001) Attention systems and the organization of the human parietal cortex. *J Neurosci* 21:5262–5271.
- Chiu YC, Yantis S (2009) A domain-independent source of cognitive control for task sets: Shifting spatial attention and switching categorization rules. *J Neurosci* 29:3930–3938.
- Kelley T, Serences J, Giesbrecht B, Yantis S (2008) Cortical mechanisms for shifting and holding visuospatial attention. *Cereb Cortex* 18:114–125.
- Crone EA, Wendelken C, Donohue SE, Bunge SA (2006) Neural evidence for dissociable components of task-switching. *Cereb Cortex* 16:475–486.
- Dias R, Robbins TW, Roberts AC (1996) Dissociation in prefrontal cortex of affective and attentional shifts. *Nature* 380:69–72.
- Liu T, Slotnick S, Serences J, Yantis S (2003) Cortical mechanisms of feature-based attentional control. *Cereb Cortex* 13:1334–1343.
- Serences J, Schwarzbach J, Courtney S, Golay X, Yantis S (2004) Control of object-based attention in human cortex. *Cereb Cortex* 14:1346–1357.
- Shomstein S, Yantis S (2004) Control of attention shifts between vision and audition in human cortex. *J Neurosci* 24:10702–10706.
- Vandenberghe R, Gitelman D, Parrish T, Mesulam M (2001) Location- or feature-based targeting of peripheral attention. *NeuroImage* 14:37–47.
- Wager T, Jonides J, Reading S (2004) Neuroimaging studies of shifting attention: A meta-analysis. *NeuroImage* 22:1679–1693.
- Wager T, Jonides J, Smith E (2006) Individual differences in multiple types of shifting attention. *Mem Cognit* 34:1730–1743.
- Kamigaki T, Fukushima T, Miyashita Y (2009) Cognitive set reconfiguration signaled by macaque posterior parietal neurons. *Neuron* 61:941–951.
- Cavanna A, Trimble M (2006) The precuneus: A review of its functional anatomy and behavioural correlates. *Brain* 129:564–583.
- Norman K, Polyn S, Detre G, Haxby J (2006) Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends Cognit Sci* 10:424–430.
- Haynes J, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685.
- Kamitani Y, Tong F (2006) Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16:1096–1102.
- Serences J, Boynton G (2007) The representation of behavioral choice for motion in human visual cortex. *J Neurosci* 27:12893–12899.
- Eger E, Ashburner J, Haynes J, Dolan R, Rees G (2008) fMRI activity patterns in human LOC carry information about object exemplars within category. *J Cognit Neurosci* 20:356–370.
- Polyn S, Natu V, Cohen J, Norman K (2005) Category-specific cortical activity precedes retrieval during memory search. *Science* 310:1963–1966.
- Serences J, Ester E, Vogel E, Awh E (2009) Stimulus-specific delay activity in human primary visual cortex. *Psychol Sci* 20:207–214.
- Haynes J, et al. (2007) Reading hidden intentions in the human brain. *Curr Biol* 17:323–328.
- Rosenau BJ, Esterman M, Chiu Y-C, Yantis S (2009) A domain-independent source of cognitive control for shifting attention in vision and working memory. *Journal of Vision* 9(8):164, 164a.
- Garavan H (1998) Serial attention within working memory. *Mem Cognit* 26:263–276.
- Colby CL, Goldberg ME (1999) Space and attention in parietal cortex. *Annu Rev Neurosci* 22:319–349.
- Silver MA, Ress D, Heeger DJ (2005) Topographic maps of visual spatial attention in human parietal cortex. *J Neurophysiol* 94:1358–1371.
- Nobre AC, et al. (2004) Orienting attention to locations in perceptual versus mental representations. *J Cognit Neurosci* 16(3):363–373.
- Pessoa L, Padmala S (2007) Decoding near-threshold perception of fear from distributed single-trial brain activation. *Cerebral Cortex* 17:691–701.
- Kay K, Naselaris T, Prenger R, Gallant J (2008) Identifying natural images from human brain activity. *Nature* 452:352–355.
- Brignani D, Lepsien J, Rushworth MF, Nobre AC (December 15, 2008) The timing of neural activity during shifts of spatial attention. *J Cognit Neurosci*, 10.1162/jocn.2008.21176.
- Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3:201–215.
- Sohn MH, Ursu S, Anderson JR, Stenger VA, Carter CS (2000) Inaugural article: The role of prefrontal cortex and posterior parietal cortex in task switching. *Proc Natl Acad Sci USA* 97:13448–13453.
- Badre D (2008) Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends Cogn Sci* 12:193–200.