

ORIGINAL ARTICLE

The Behavioral Relevance of Task Information in Human Prefrontal Cortex

Michael W. Cole^{1,2}, Takuya Ito^{1,2}, and Todd S. Braver¹

¹Department of Psychology, Washington University, St Louis, MO 63130, USA, and ²Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ 07120, USA

Address correspondence to Prof. Michael W. Cole, Center for Molecular and Behavioral Neuroscience, 197 University Ave, Suite 212, Newark, NJ 07102, USA. Email: mwcole@mwcole.net

Abstract

Human lateral prefrontal cortex (LPFC) is thought to play a critical role in enabling cognitive flexibility, particularly when performing novel tasks. However, it remains to be established whether LPFC representation of task-relevant information in such situations actually contributes to successful performance. We utilized pattern classification analyses of functional MRI activity to identify novelty-sensitive brain regions as participants rapidly switched between performance of 64 complex tasks, 60 of which were novel. In three of these novelty-sensitive regions—located within distinct areas of left anterior LPFC—trial-evoked activity patterns discriminated correct from error trials. Further, these regions also contained information regarding the task-relevant decision rule, but only for successfully performed trials. This suggests that left anterior LPFC may be particularly important for representing task information that contributes to the cognitive flexibility needed to perform successfully in novel task situations.

Key words: functional magnetic resonance imaging, instructional control, multivariate pattern analysis, prefrontal cortex, task rules

Introduction

Decades of neuroscientific research has established that human lateral prefrontal cortex (LPFC) is more active when a high degree of cognitive flexibility is required, such as during novel task learning (Cole, Laurent, et al. 2013), decision making (McGuire and Botvinick 2010), and rapid task-switching (Braver et al. 2003; Ruge et al. 2013). Further, more recent work has utilized advances in neuroimaging pattern classification methods (Norman et al. 2006) to establish that LPFC activity patterns contain information regarding the currently relevant task-set (Cole et al. 2011; Woolgar et al. 2011; Zhang et al. 2013; Waskom et al. 2014). Moreover, Cole et al. (2011) demonstrated that PFC activity patterns represent task relevant rule information not only during practiced task performance, but also when performing tasks with novel rule combinations. However, it remains unclear whether such task representations—and LPFC representations generally—actually contribute to behavioral performance (but see Etzel et al.

(2015) for a recent example of this type of finding). The present study sought to determine whether flexible representations within LPFC are behaviorally relevant, particularly during novel task performance, for which cognitive flexibility is most critical. Such a link between LPFC activity patterns and task performance would suggest that a critical means by which LPFC contributes to successful cognitive flexibility is through task rule representation, particularly in novel circumstances.

A recent study found that most aspects of a complex task could be decoded from LPFC in nonhuman primates (Rigotti et al. 2013). They found that the ability to decode task information was diminished during error trials, indicating a collapse in task “dimensionality” within LPFC. In light of this finding, we performed decoding of multiple dimensions of complex cognitive tasks, testing decoding of task rules on correct versus error trials. Importantly, our paradigm required extraordinary behavioral flexibility—most trials involved learning to perform tasks that

involved the novel integration of multiple rules. Thus, identifying behavior-sensitive activity patterns within LPFC would expand upon previous work to show that representations within human LPFC contribute to successful behavioral flexibility in the complex cognitive task scenarios indicative of human cognition.

Our primary approach was to utilize a previously developed cognitive paradigm that requires rapid instructed task learning (RITL) (Cole et al. 2010). A key feature of this paradigm is that it involves a large set of distinct tasks (64) each composed of 3 integrated rules (Fig. 1). In prior work, we demonstrated successful decoding of multiple task rule representations within LPFC (Cole et al. 2011). In the current study, we limited preparation time to induce errors. This allowed us to obtain enough error trials per subject to observe relationships between brain activity and task performance. We hypothesized that portions of LPFC's primary functional network—the cognitive control network (Cole and Schneider 2007; Duncan 2010)—would involve four properties indicative of important contributions to flexible cognitive control: 1) sensitivity to task novelty (novel vs. practiced tasks), 2) sensitivity to task performance (correct vs. error trials); 3) representation of decision rules (rule-specific activity patterns); and 4) a collapse of task rule representation on error trials (suggesting dimensionality reduction). Together these properties would strongly support a key role for LPFC (and possibly its distributed network) in producing flexible behavior.

Materials and Methods

Participants

Participants were recruited from Washington University in St Louis and the surrounding area. During screening, potential participants were excluded if they had any medical, neurological or psychiatric illness, any contraindications for MRI scans, were non-native English speakers, or were left-handed. Additionally, participants were excluded from further analysis due to poor data yield during scanning, either for technical reasons or below-criteria task performance (below 70% accuracy during the practice session, or below 60% accuracy on practiced tasks during the test session). Five scanned participants were excluded in this manner, one due to ending the fMRI session early and the others due to low performance accuracy (<60%) on practiced tasks. This left 21 participants (9 females), aged 18–32 years (mean age = 22.7 years) included in the analysis. All participants gave informed written consent prior to study participation.

MRI Data Collection

Image acquisition was performed on a 3T Siemens Trio MRI scanner. 36 transaxial slices were acquired every 2000 ms (field of view = 256 mm, echo time = 25 ms, flip angle = 90°, voxel dimensions = 4.0 mm³) with a total of 258 gradient echo-planar imaging volumes collected per run (across 10 runs). Three-dimensional anatomical magnetization prepared rapid acquisition

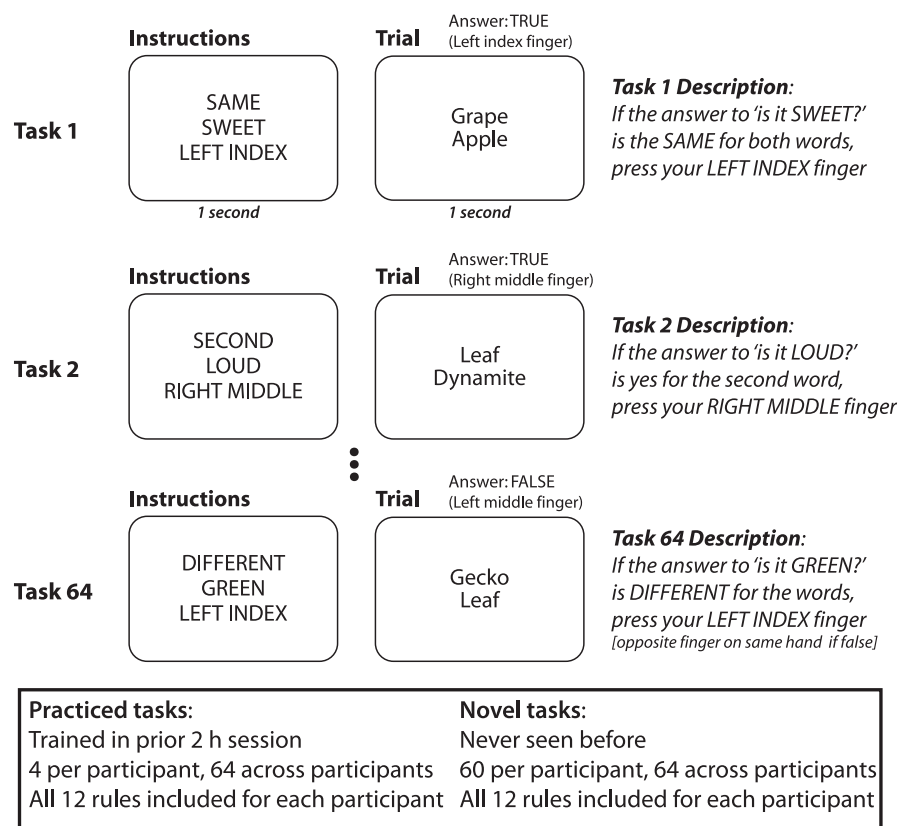


Figure 1. Cognitive task paradigm. The permuted rule operations (PRO) paradigm was used to test highly flexible cognitive control. Four tasks were practiced in a previous session, while 60 tasks were novel during the neuroimaging session. Preparation time was shortened from previous versions of the paradigm in order to induce enough errors to identify brain representations contingent upon behavioral performance. The paradigm was used to isolate 3 forms of flexible cognitive control: 1) novel versus practiced task performance, 2) correct versus error behavioral performance, and 3) decision rule identity (SAME, DIFFERENT, SECOND, or NOTSECOND rules).

gradient echo images and T_2 structural in-plane images were collected for each subject before fMRI data collection.

Task Paradigm

The permuted rule operations (PRO) cognitive paradigm was used (Fig. 1), as described previously (Cole et al. 2010, 2011; Cole, Reynolds, et al. 2013). Briefly, 12 task rules of three types (logical decision, sensory semantic, and motor response) were combined into 64 distinct combinations, producing 64 unique tasks (i.e., instruction sets). Tasks were defined as unique combinations of rules, such that the same stimuli would require a distinct set of cognitive operations for correct performance across distinct tasks. Four of these 64 tasks (counterbalanced across subjects) were practiced in a prior 2-h behavioral session (outside the scanner), with the remaining tasks presented for the first time during the test (fMRI) session.

As described previously (Cole et al. 2010), the semantic rules consisted of sensory semantic decisions (e.g., “is it sweet?”). The logical decision rules specified how to respond based on the two semantic decision outcomes for each trial. The motor response rules specified which button to press based on the logical decision outcome. The task instructions made explicit reference to the motor response for a “true” outcome, while participants knew (from the practice session) to use the other finger on the same hand for a “false” outcome. Stimuli and motor responses were counterbalanced overall and with respect to all rules to reduce the chance of biased responding.

Unlike previous versions of this paradigm, two key changes were made for the current study: 1) the task rules relevant for the current trial were only available for a short period before the trial began (1900 ms vs. the previous 5–9 s), thus limiting encoding and preparation time; and 2) only one trial was presented for a given task, thus increasing the frequency of task-switching. Additionally, the rule that was previously cued as “JUST ONE” was now referred to as “DIFFERENT” to ease comprehension.

Each subject completed 10 runs comprised 36 trials each during the imaging session, totaling 360 trials per session. Of the 360 trials per session, 240 trials were considered novel tasks (i.e., tasks that were encountered for the first time during the imaging session, and repeated no more than 4 times throughout the session), and 120 trials were considered practiced tasks (i.e., the 4 tasks that were practiced in a prior 2-h behavioral session; encountered 30 times each in the fMRI session). Within each trial, the timing was as follows: 1000-ms instruction presentation, 900-ms delay (between instruction and trial stimuli), 1500-ms trial stimulus presentation, followed by a variable intertrial delay (6.6–14.6 s).

Preprocessing

Preprocessing was performed using AFNI (version 2011-12-21) (Cox 1996). Images were slice-time corrected, motion-corrected (realigned), detrended by run, aligned to a Talairach template image, and temporally compressed using general linear model (GLM) estimates—fitting a canonical hemodynamic response function to each of the 360 trials separately.

We used PyMVPA (version 0.6) for classification analyses (Hanke et al. 2009). We averaged the data across trials within each run (by task condition), such that there were 10 data points per condition per subject (one per run). Note that this was true regardless of the particular condition of interest (e.g., correct trials, SAME rule trials, incorrect DIFFERENT rule trials). There were some exceptions where certain subjects performed entire runs

without any errors, resulting in fewer than 10 data points for a particular error-trial condition. In such cases, we used a standard balancing algorithm that equalized the number of samples between conditions to be included in the classification process. This ensured that the classifier would not be biased toward classifying a particular condition due to an unequal number of samples per condition. As part of the data normalization process prior to classification, we z-normalized each voxel, subtracting the across-sample mean and dividing by the across-sample standard deviation. Note that results were highly similar with and without regressing out trial-to-trial reaction times, an approach suggested to deal with potential task difficulty confounds (Todd et al. 2013; Woolgar et al. 2014).

Searchlight-Based Classification Analysis

We defined and analyzed regions of interest (ROIs) using a two-stage approach, in which voxel clusters identified by the searchlight analysis were subsequently re-trained and analyzed as independent ROIs (Etzel et al. 2013).

To define the ROIs, we used PyMVPA's (Hanke et al. 2009) searchlight algorithm, which employs linear support vector machines to classify fMRI activity patterns in the whole brain using two-voxel radius “searchlight spheres” surrounding each voxel (corresponding to a 33-voxel cluster per searchlight). Searchlights were restricted to a whole-brain gray matter mask that was dilated by one voxel. A novel versus practiced classification was performed, where novel and practiced conditions were collapsed across correct and incorrect behavioral performance trials.

We spatially smoothed each subject's statistical map using an iterative algorithm (AFNI's 3dBlurToMask) that ensured the final empirical smoothness was 6-mm FWHM. Note that each subject's statistical map was estimated (using AFNI's 3dFWHMx) to be at or below 6-mm FWHM prior to smoothing. We then compared each voxel's accuracy to chance (50%) using one-way one-sample *t*-tests. The group map was then cluster thresholded to correct for multiple comparisons. The cluster size was estimated using AFNI's 3dClustSim with 10 000 Monte Carlo simulations using $P < 0.01$ as the uncorrected threshold and a smoothing parameter of 6-mm FWHM. For a corrected $P < 0.05$, the cluster size was calculated to be 34.1 voxels.

ROI-Based Classification Analyses

We performed three basic classifications with the ROIs: 1) two-way classification of novel and practiced task conditions; 2) two-way classification of Correct and Incorrect task conditions; 3) four-way classification of logical decision rules.

The ROI-based classifications were carried out with support vector machines using a linear kernel with the cost parameter scaled to the norm of the data (i.e., the norm of the input matrix), as is standard in PyMVPA. We obtained group accuracies by classifying each subject's data separately and subsequently averaging the classification rates across subjects. We employed a leave-one-run-out cross-validation scheme to avoid circularity in the classifications. For each cross-validation, we excluded all samples from a single run across all conditions, trained using data from the remaining runs, and finally tested the classifier on the samples from the excluded run. The 4-way rule-based classification was performed in the same manner, but with four samples (one per decision rule) per run. We ensured that each region's across-subject accuracy distribution was approximately normally distributed prior to running *t*-tests. Further, we reran all *t*-tests after arcsine transforms (a common method to make

accuracy data more normally distributed), obtaining highly similar results as before. For each analysis, we corrected for multiple comparisons on the 8 searchlight-defined ROIs using false discovery rate (FDR) (Genovese et al. 2002).

Permutation Testing

We attempted to increase the number of error trials from previous implementations of the PRO paradigm, while simultaneously ensuring that participants were performing with above chance accuracy. It was therefore necessarily the case that all subjects produced more correct than error trials, which added additional complexity to the logical decision rule correct versus error classification analysis. In particular, we needed a way to ensure classification was not biased toward greater accuracy for correct trials simply due to the greater number of correct trials for classifier training. We used permutation testing to produce unbiased P -values, based on null distributions produced with the same degree of bias as the real classification results. In other words, since both the accuracies produced by the permutation test and the true classification test had the same ratio of correct to error trials, evaluating the true classification rate relative to the null distribution produced an unbiased P -value, regardless of the different frequencies of correct and error trials.

Using PyMVPA's AttributePermutator with 10 000 Monte Carlo simulations, we created 2 null distributions (correct and error) for each of the classification types by permuting the rule encodings across all data samples within either condition. After constructing the null distributions for correct and error conditions separately, we then subtracted the error null distribution from the correct null distribution to create a new null distribution of correct versus error condition differences. Each value in one distribution was randomly paired with another value in the other distribution for subtraction. To test the significance between the true accuracies of correct and error conditions relative to the null distribution, we computed the difference of the correct and error trials separately using the original (nonpermuted) data labels. By comparing the nonpermuted correct versus error differences relative to the differences of the 2 null distributions, we were able to determine if there was a significant difference between the correct and error accuracies. Since our permutation approach was not standard, we also conducted a version of this analysis closer to standard classification approaches to determine if it mattered in practice which approach was used. This involved randomly sampling categories with replacement, ensuring an equal number of samples across classes during training and testing. Each class was sampled 1000 times—this high number was necessary to achieve a stable solution over multiple runs of the analysis.

Results

Behavioral Performance

Participants performed the PRO cognitive paradigm (Fig. 1), which involved 2 sessions. Participants were 85% accurate on average during the (behavior-only) practice session, in which they learned and performed four of the 64 tasks repeatedly. During the neuroimaging session the average accuracy rate was 81%, which reflected the intermixed performance of the 4 practiced tasks and 60 additional novel tasks. Looking at the 2 task types separately, practiced task accuracy was 83% ($t_{(20)} = 18.0$, $P < 0.00001$), and novel task accuracy was 80% ($t_{(20)} = 19.0$, $P < 0.00001$). The performance difference between the 2 task

types was significant ($t_{(20)} = 2.2$, $P = 0.04$). Note that this was not due to a speed-accuracy trade off, as error trials were slower than correct trials (correct mean RT: 953 ms, error mean RT: 1051 ms, $t_{(20)} = 11.0$, $P < 0.00001$). (These results were similar when computed separately for novel and practiced trial types. Novel trials = correct mean RT: 962 ms, error mean RT: 1051 ms, $t_{(20)} = 8.8$, $P < 0.00001$. Practiced trials = correct mean RT: 936 ms, error mean RT: 1058 ms, $t_{(20)} = 8.1$, $P < 0.00001$.) This result indicates the increased difficulty associated with performing novel tasks. On the other hand, performance of both task types was relatively high and the difference between them was relatively small, indicating that participants were able to achieve successful task performance even on novel tasks. However, we achieved a key goal of our modifications of the paradigm—introducing a substantial number of performance errors relative to our prior work (for which accuracy was over 90% for both novel and practiced tasks; Cole et al. (2010)). This was likely due to the reduction in encoding and preparation time (from 5–9 s to 1900 ms) and using a trial-by-trial task switching design (which may have also increased the relative difficulty of practiced task performance).

Localizing Flexible Cognitive Processing: Novel Versus Practiced Tasks

Our first goal was to functionally localize brain regions potentially contributing to cognitive flexibility. In particular, we tested for regions sensitive to novel versus practiced task performance, given that RITL requires especially flexible cognition. Thus, we identified brain regions with differential fMRI activity across novel and practiced tasks. It is possible to identify regions in this manner using either traditional univariate GLM analyses, or multivariate pattern classification approaches (Etzel et al. 2013). In prior work with this same paradigm (Cole et al. 2010), a univariate approach was used that identified novelty-sensitive activity in 2 portions of prefrontal cortex: right dorsolateral prefrontal cortex (DLPFC) and left anterior prefrontal cortex (aPFC). However, in that study, we capitalized on the slower timing that separated task instructional cues and task trials to reveal these regions in terms of their distinct activity dynamic signatures. In the current study, because these 2 trial phases were in closer temporal proximity (to increase error rates and total trial numbers), we expected that the multivariate analysis approach would have greater sensitivity to detect novelty-related effects.

Indeed, this was the case. Using ROIs in the same locations as the previous study (the spherical regions in Fig. 2; each with a radius of 9 mm), we found no difference in mean activation amplitude between novel and practiced trials for either PFC region (treated as a univariate GLM analysis; aPFC mean beta difference: -0.13 , $t_{(20)} = -0.88$, $P = 0.4$; DLPFC mean beta difference: -0.06 , $t_{(20)} = -0.60$, $P = 0.56$). Moreover, a standard whole-brain univariate GLM analysis failed to identify any regions that reliably differentiated novel and practiced tasks after correcting for multiple comparisons ($P < 0.05$, family-wise error corrected for multiple comparisons). In contrast, a multivariate pattern analysis of the 2 a priori PFC ROIs did show successful classification of novel versus practiced task trials, although the effect was only marginally significant in the aPFC (53%, $t_{(20)} = 1.48$, $P = 0.08$; DLPFC: 56%, $t_{(20)} = 2.85$, $P = 0.005$). Likewise, a whole-brain searchlight analysis identified an additional 8 ROIs (Table 1) that correspond with many portions of LPFC and its extended functional network, the cognitive control network (Cole and Schneider 2007; Duncan 2010). We conducted subsequent analyses using ROI-based classification with these 8 ROIs, as well as the 2 ROIs described above from the previous study (Fig. 2).

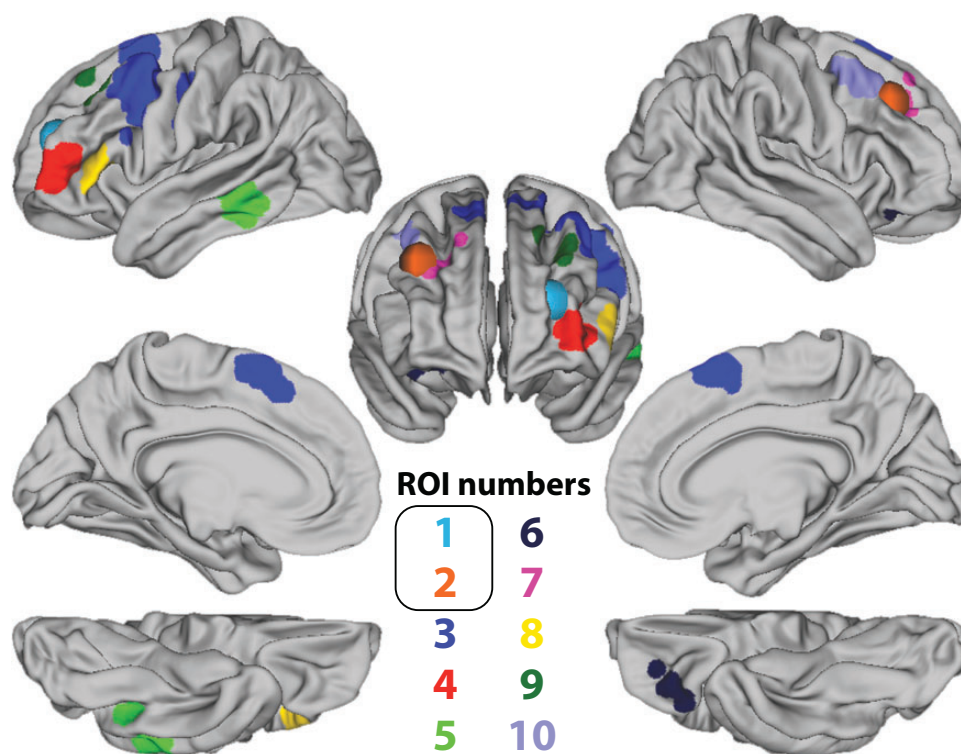


Figure 2. Regions of interest (ROIs). Regions defined based on are spherical (ROIs 1 and 2). The other regions/clusters were defined using a searchlight voxel activation pattern analysis decoding novel versus practiced tasks ($P < 0.05$, family-wise error corrected for multiple comparisons).

Table 1 Regions of interest (ROIs) with their corresponding cluster sizes and Talairach coordinates, which reflect the center of mass of each ROI

ROI	Name	Area	Size (voxels)	Talairach coordinates
1	Anterior lateral prefrontal cortex	10	99	22, -48, 19
2	Dorsolateral prefrontal cortex	9	111	30, 27, 35
3	Pre/supplementary motor areas	6	437	-28, 7, 49
4	Anterior lateral prefrontal cortex	10	73	-38, 50, 6
5	Posterior middle temporal cortex	20	63	-57, -40, -11
6	Lateral orbitofrontal cortex	11/47	58	33, 40, -6
7	Dorsolateral prefrontal cortex	8	45	21, 34, 37
8	Anterior lateral prefrontal cortex	45	41	-50, 30, 5
9	Superior prefrontal cortex	8	39	-27, 33, 44
10	Superior prefrontal cortex	6	35	43, 14, 48

*FDR corrected for multiple comparisons.

Testing for Task Rule Representations in Novelty-Sensitive Brain Regions

Consistent with regions that likely contribute critically to flexible cognitive control, we previously found that some LPFC regions represent specific task rules (Cole et al. 2011). We next sought to test if the novelty-sensitive regions we identified above also code for task rules. We again used multivariate pattern analysis, this time decoding logical decision rules in the 10 ROIs (4-way classification, chance = 25%). We found that activity patterns in the (Cole et al. 2010) left aPFC region coded for logical decision rules (accuracy = 30%, $t_{(20)} = 2.25$, $P = 0.02$). This was not the case for the Cole2010 right DLPFC region (accuracy = 28%, $t_{(20)} = 0.96$, $P = 0.2$, NS), though that region appeared to code for sensory semantic rules (accuracy = 28%, $t_{(20)} = 2.38$, $P = 0.01$). Looking across the other 8 ROIs, only regions within left anterior LPFC coded for

logical decision rules (ROIs 4 and 8; see Table 2). This suggests regions within left anterior LPFC code task-relevant information.

Behavioral Relevance of Representations in Novelty-Sensitive Brain Regions

We next tested whether the novelty sensitive regions might contribute to successful task performance. This property would be expected of a region that contributes critically to flexible behavior. We began by testing whether any of these regions showed differential activity patterns on correct and error trials, under the assumption that if task representations in these regions contributed to performance they should track performance differences. We used permutation tests for this analysis, due to the larger number of correct than error trials (see Materials and Methods).

Table 2 ROIs 2 through 10 are family-wise error corrected for multiple comparisons

ROI	Novel correct versus practiced correct	Correct logic rule classification	Novel correct versus novel error	Practiced correct versus practiced error
1	<i>P</i> = 0.08	P = 0.02	P < 0.0001	P = 0.004
2	P = 0.005	<i>P</i> = 0.17	P = 0.008	P = 0.002
3	P = 0.0002	<i>P</i> = 0.08	P < 0.0001	P = 0.0003
4	P = 0.02	P = 0.02	P < 0.0001	P = 0.0003
5	P = 0.04	<i>P</i> = 0.23	P = 0.02	P = 0.0005
6	P = 0.001	<i>P</i> = 0.14	P = 0.01	P = 0.02
7	P = 0.02	<i>P</i> = 0.17	P = 0.02	P = 0.0003
8	P = 0.05	P = 0.02	P = 0.03	<i>P</i> = 0.08
9	P = 0.05	<i>P</i> = 0.19	P = 0.02	P = 0.01
10	P = 0.008	<i>P</i> = 0.19	P = 0.01	P = 0.01

Values in bold indicate $P \leq 0.05$. Values in italic indicate P -values where $0.05 < P < 0.10$.

We found that activity in all of the 10 ROIs, both the a priori ones and the searchlight-identified ones, were sensitive to whether trial performance was accurate or associated with an error ($P < 0.05$, FDR corrected for multiple comparisons). This was true for novel tasks (Fig. 3A), and also true for practiced tasks, though one of the ROIs was only marginally significant in this analysis (Fig. 3B; see also Table 2).

Testing for Behaviorally Relevant Decision Rule Representations: Possible Collapsing of Task Rule Representations During Error Trials

The above results suggest that left anterior LPFC regions 1) are sensitive to novelty, 2) represent task information, and 3) are sensitive to task accuracy—consistent with regions important for flexible cognitive control. We tested this more directly—while also testing for the possibility of a collapse in representational quality during error relative to correct trials (Rigotti et al. 2013)—by decoding rule representations separately for correct and error trials. We found that all three of the anterior LPFC ROIs could decode logical decision rules for correct trials (ROI 1: 30%, $P = 0.02$, ROI 4: 30% $P = 0.02$, ROI 8: 31% $P = 0.02$) but not for error trials (ROI 1: 24% $P = 0.36$, ROI 4: 25% $P = 0.48$, ROI 8: 26% $P = 0.41$). Critically, these classification accuracies were significantly greater for correct versus error trials ($P < 0.05$) for all of these ROIs (Fig. 4). [A secondary permutation approach that oversampled error trials (see Materials and Methods) produced results that were quite similar to the other approach, with 2 of the 3 left anterior LPFC regions being highly statistically significant. Specifically, ROI 1 was $P = 0.007$, ROI 4 was $P = 0.14$, and ROI 8 was $P = 0.001$. Undersampling of correct trials to match the number of error samples (rather than oversampling error trials to match the number of correct samples) produced similar though only marginally significant results: ROI 1 $P = 0.05$, ROI 4 $P = 0.13$, ROI 8 = 0.06. The reduction in statistical significance is likely due to throwing out a substantial portion of the (correct trial) data.]

This indicates that representational quality was likely better for correct performance trials, suggesting that rule representations in these regions contributed to behavior. In turn, this suggests the functional importance of left anterior LPFC representations to flexible cognition, as they showed all 4 signatures of interest (Table 2).




Figure 3. Activity classifications for correct versus error task performance. Accuracies for pattern classifications in each ROI for novel (A) and practiced (B) tasks are shown. P -values for the exploratory ROIs (the last 8 ROIs) were corrected for multiple comparisons using false discovery rate. Error bars indicate inter-subject standard errors.




Figure 4. Activity classifications for logical decision rules, categorized by behavioral performance. Accuracies for pattern classifications in each ROI are shown. P -values for the exploratory ROIs (the last 8 ROIs) were corrected for multiple comparisons using false discovery rate. Error bars indicate inter-subject standard errors. ROI labels with greater correct trial than error trial classification (1, 4, and 8) are in bold ($P < 0.05$).

Testing the Unity of the Anterior LPFC Regions

The apparent functional similarity and spatial proximity of the three left anterior LPFC regions suggested they might be one

large unified brain region. We tested this by considering that distinct functional brain regions have distinct whole-brain brain connectivity patterns (Cohen et al. 2008). Specifically, we assessed the similarity of functional connectivity patterns among the 3 regions based on the results from 2 previous resting-state functional connectivity fMRI studies (Power et al. 2011; Yeo et al. 2011). Importantly, recent results indicate that the network organization identified during resting state is also present across a wide variety of tasks (Cole, Bassett, et al. 2014), suggesting connectivity pattern inferences based on resting-state functional connectivity generalize to tasks.

We found that all 3 left anterior LPFC regions were in separate subnetworks in both brain network partitions (Fig. 5). Note that both studies used distinct network partition approaches. Power et al. (2011) used graph theoretical community detection, while Yeo et al. (2011) used a recently developed clustering algorithm that models data with a von Mises-Fisher distribution. Despite these differences, both approaches produced similar network partitions, both of which indicate the 3 left anterior LPFC regions identified here are in separate subnetworks. Focusing on the labels assigned by Power et al. (2011), ROI 1 is in the “salience” network, ROI 4 is in the “frontoparietal” network, and ROI 8 is in the “ventral attention” network. All 3 networks are considered to be control-related or “task positive” in nature (Power et al. 2011).

Discussion

We found that 3 regions in left anterior LPFC (see ROIs 1, 4, and 8 in Fig. 2) met all of our criteria for behaviorally relevant task set representation. We began by looking at 2 ROIs from a previous (regional activation amplitude) study (Cole et al. 2010) using the same cognitive paradigm. We found that activity patterns within the left anterior LPFC region could distinguish novel from practiced tasks, correct from error trials, and logical decision rules from one another. Importantly, logical decision rule classification was significantly higher for correct than error trials, suggesting these representations were behaviorally relevant. Among 8 additional ROIs, only 2 regions (also in left anterior LPFC, but clearly distinct from the first ROI) showed this same pattern of results. Together these results suggest representations across 3 distinct portions of left anterior LPFC contribute to adaptive cognition and the behavior it produces.




Figure 5. The 3 left anterior LPFC regions are in separate brain network partitions. Brain network partitions were identified in 2 previous studies by clustering similar patterns of resting-state functional connectivity. Power et al. (2011) used graph theoretical community detection, while Yeo et al. (2011) used k-means clustering. Both studies used resting-state data from over 100 healthy young adults. We used the Power et al. (2011) “consensus” partition and the Yeo et al. (2011) “17 networks” partition, both from the Connectome Workbench software package. We plotted the centers of each of the 3 left anterior LPFC regions as spheres. ROI 1 is light blue, ROI 4 is red, and ROI 8 is yellow. The distinct colors under each sphere indicate that both partitions assigned these regions to 3 separate brain networks.

The cognitive paradigm used in this study was designed to investigate RITL—the ability to learn and perform tasks on the first try. RITL is one of the most highly adaptive forms of cognitive flexibility and cognitive control (Cole, Laurent et al. 2013), suggesting brain regions that contribute to RITL likely facilitate cognitive flexibility in a variety of contexts. In this paradigm, novel tasks were new to participants, while practiced tasks were performed repeatedly in a prior session. All aspects were controlled for across novel and practiced tasks except for task novelty (e.g., stimulus features, previous practice with each task rule). We found that activity patterns in both Cole et al. (2010) regions distinguished between novel and practiced tasks, suggesting fine spatial patterns in these regions contribute to producing RITL. Further, we found an additional 8 regions that distinguished novel and practiced tasks, many of which overlapped with a well-established brain system (consisting of multiple subnetworks) called the cognitive control network (Cole and Schneider 2007; Vincent et al. 2008). These results strongly suggest that the cognitive control network—especially its portions in LPFC—contribute to RITL. This is largely consistent with previous RITL studies, which primarily used regional activation amplitudes and separated instruction from task execution periods (Cole et al. 2010; Stocco et al. 2012).

The present results are also consistent with previous studies of task rule representation (Reverberi et al. 2012), which also found rule representations within LPFC along with other portions of the control network. The present study built upon previous findings to test for sensitivity of these rule representations to task performance. In particular, we hypothesized that task rules would be better represented in brain regions during correct than error trials, consistent with those rule representations contributing to behavioral performance. Such a finding would be consistent with recent results in nonhuman primate LPFC (Rigotti et al. 2013), but this time in the context of multiple complex tasks and highly abstract rules in humans. Three portions of left anterior LPFC showed the expected rule representations contingent on task performance, suggesting these brain regions are especially important for adaptive cognition and behavior.

Humans can perform RITL (and related abilities) while other animals cannot. It has been speculated that this uniquely human ability may be due to extensive expansion of anterior LPFC in human evolution (Cole 2009; Cole, Laurent, et al. 2013). Highly relevant to this evolutionary argument, a recent study directly compared the connectivity patterns of LPFC regions between humans and macaque monkeys (Neubert et al. 2014). They found very high correspondence among all regions of LPFC except for the most anterior portion of LPFC. Consistent with this, we found that the most anterior portions of LPFC were the most important for adaptive cognitive control, including RITL.

The 3 aspects of cognitive flexibility that we tested for were incompatible in some ways. For instance, a region that is involved only during novel tasks might not distinguish correct versus error trials (or decision rules) during practiced tasks. Our finding of all 3 aspects in anterior LPFC regions suggests distinct activation patterns within these regions encode for different aspects of cognitive control. Future research may be able to test whether distinct subpopulations of neurons encode these different aspects of cognitive flexibility, or whether there is overlap in the populations but orthogonal activation patterns to reduce interference.

We found that the 3 left anterior LPFC regions were part of 3 separate brain subnetworks (Fig. 5), based on having distinct functional connectivity patterns (Power et al. 2011; Yeo et al. 2011). This is strong evidence that these are distinct functional

regions, despite their spatial proximity and similar functionality in the context of this study. It will be important for future studies to determine why these 3 regions show such similar functionality despite their distinct connectivity patterns, and whether they interact with one another to support cognitive flexibility.

One concern with the correct versus error classification analysis (Fig. 3) may be the possibility of surprise/oddball activations to the infrequent error trials. Specifically, changes in activation related to infrequent events could have driven some of the differences between error and correct trials that lead to high classification accuracies. This is a general problem, as this lack of balance in correct versus error frequency is necessary for any dataset in which nonrandom (i.e., overall above chance) task performance is required. Importantly, the better rule classification performance for correct versus error trials (Fig. 4) was unlikely to have been driven by surprise/oddball effects on error trials, given that rule classification was orthogonal to behavioral accuracy. This suggests that any surprise/oddball effect that may be present for error trials is unlikely to explain the improved representational quality of task rules on correct trials within left anterior LPFC.

The present results are consistent with several theories of LPFC function. First, although several recent findings call this theory into question (Crittenden and Duncan 2014; Reynolds et al. 2012), the present findings are consistent with a hierarchical organization of LPFC, in which the most abstract rules are represented in the most anterior portions of LPFC (Fuster 2001; Badre et al. 2009). These results are also consistent with the Guided Activation theory of LPFC function (Miller and Cohen 2001), in which LPFC maintains task-relevant representations and influences remote task-relevant regions via top-down biasing of processing. Expansion of this framework based on recent advances in network science—the Flexible Hub theory (Cole, Repovs, et al. 2014)—proposes that the frontoparietal control network consists of highly connected “hub” regions that flexibly shift their functional connectivity depending on task demands. Importantly, a recent study testing this theory (Cole, Reynolds, et al. 2013) used the same cognitive paradigm as used here, and found that functional connectivity with LPFC (along with other frontoparietal regions) could be used to decode which task rules were being performed during RITL. This suggests the representations identified in the present study likely spread to other behavior-implementing regions (e.g., primary motor cortex, sensory cortices) via changes in functional connectivity. It will be important for future studies to investigate this possibility, as well as the possibility that these functional connectivity changes are sensitive to behavioral performance as found here for left anterior LPFC activation patterns.

Funding

Research reported in this publication was supported by the National Institute of Mental Health of the National Institutes of Health under Award Numbers R01 MH66078 and K99/R00 MH096801. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Notes

We thank Joset Etzel, Mattia Rigotti, Kevin Oksanen, and Maria Chushak for their help with data collection, analysis, and helpful discussions. *Conflict of Interest:* None declared.

References

- Badre D, Hoffman J, Cooney JW, D'Esposito M. 2009. Hierarchical cognitive control deficits following damage to the human frontal lobe. *Nat Neurosci.* 12:515–522.
- Braver TS, Reynolds JR, Donaldson DI. 2003. Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron.* 39:713–726.
- Cohen AL, Fair DA, Dosenbach NUF, Miezin FM, Dierker D, Van Essen DC, Schlaggar BL, Petersen SE. 2008. Defining functional areas in individual human brains using resting functional connectivity MRI. *Neuroimage.* 41:45–57.
- Cole M. 2009. The Biological Basis of Rapid Instructed Task Learning [Dissertation], University of Pittsburgh. <http://d-scholarship.pitt.edu/8386/> Last accessed March 30, 2015.
- Cole MW, Bagic A, Kass R, Schneider W. 2010. Prefrontal dynamics underlying rapid instructed task learning reverse with practice. *J Neurosci.* 30:14245–14254.
- Cole MW, Bassett DS, Power JD, Braver TS, Petersen SE. 2014. Intrinsic and task-evoked network architectures of the human brain. *Neuron.* 83:238–251.
- Cole MW, Etzel JA, Zacks JM, Schneider W, Braver TS. 2011. Rapid transfer of abstract rules to novel contexts in human lateral prefrontal cortex. *Front Hum Neurosci.* 5:142.
- Cole MW, Laurent P, Stocco A. 2013. Rapid instructed task learning: a new window into the human brain's unique capacity for flexible cognitive control. *Cogn Affect Behav Neurosci.* 13:1–22.
- Cole MW, Repovs G, Anticevic A. 2014. The frontoparietal control system: a central role in mental health. *Neuroscientist.* 20:652–664.
- Cole MW, Reynolds JR, Power JD, Repovs G, Anticevic A, Braver TS. 2013. Multi-task connectivity reveals flexible hubs for adaptive task control. *Nat Neurosci.* 16:1348–1355.
- Cole MW, Schneider W. 2007. The cognitive control network: integrated cortical regions with dissociable functions. *Neuroimage.* 37:343–360.
- Cox RW. 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res.* 29:162–173.
- Crittenden BM, Duncan J. 2014. Task difficulty manipulation reveals multiple demand activity but no Frontal Lobe Hierarchy. *Cereb Cortex.* 24(2):532–540.
- Duncan J. 2010. The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends Cogn Sci (Regul Ed).* 14:172–179.
- Etzel JA, Cole MW, Zacks JM, Kay KN, Braver TS. 2015. Reward motivation enhances task coding in frontoparietal cortex. *Cereb Cortex.* bhu327.
- Etzel JA, Zacks JM, Braver TS. 2013. Searchlight analysis: promise, pitfalls, and potential. *Neuroimage.* 78:261–269.
- Fuster J. 2001. The prefrontal cortex—an update: time is of the essence. *Neuron.* 30:319–333.
- Genovese C, Lazar N, Nichols T. 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage.* 15:870–878.
- Hanke M, Halchenko YO, Sederberg PB, Hanson SJ, Haxby JV, Pollmann S. 2009. PyMVPA: a python toolbox for multivariate pattern analysis of fMRI data. *Neuroinform.* 7:37–53.
- Mcguire JT, Botvinick MM. 2010. Prefrontal cortex, cognitive control, and the registration of decision costs. *Proc Natl Acad Sci USA.* 107:7922–7926.
- Miller E, Cohen J. 2001. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci.* 24:167–202.

- Neubert F-X, Mars RB, Thomas AG, Sallet J, Rushworth MFS. 2014. Comparison of human ventral frontal cortex areas for cognitive control and language with areas in monkey frontal cortex. *Neuron*. 81:700–713.
- Norman K, Polyn S, Detre G, Haxby J. 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci (Regul Ed)*. 10:424–430.
- Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL, et al. 2011. Functional network organization of the human brain. *Neuron*. 72:665–678.
- Reverberi C, G6rgen K, Haynes J-D. 2012. Compositionality of rule representations in human prefrontal cortex. *Cereb Cortex*. 22:1237–1246.
- Reynolds JR, O'Reilly RC, Cohen JD, Braver TS. 2012. The function and organization of lateral prefrontal cortex: a test of competing hypotheses. *PLoS One*. 7:e30284.
- Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S. 2013. The importance of mixed selectivity in complex cognitive tasks. *Nature*. 497(7451):585–590.
- Ruge H, Jamadar S, Zimmermann U, Karayanidis F. 2013. The many faces of preparatory control in task switching: reviewing a decade of fMRI research. *Hum Brain Mapp*. 34:12–35.
- Stocco A, Lebiere C, O'Reilly RC, Anderson JR. 2012. Distinct contributions of the caudate nucleus, rostral prefrontal cortex, and parietal cortex to the execution of instructed tasks. *Cogn Affect Behav Neurosci*. 12(4):611–628.
- Todd MT, Nystrom LE, Cohen JD. 2013. Confounds in multivariate pattern analysis: theory and rule representation case study. *Neuroimage*. 77:157–165.
- Vincent JL, Kahn I, Snyder AZ, Raichle ME, Buckner RL. 2008. Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. *J Neurophysiol*. 100:3328–3342.
- Waskom ML, Kumaran D, Gordon AM, Rissman J, Wagner AD. 2014. Frontoparietal representations of task context support the flexible control of goal-directed cognition. *J Neurosci*. 34:10743–10755.
- Woolgar A, Golland P, Bode S. 2014. Coping with confounds in multivoxel pattern analysis: what should we do about reaction time differences? A comment on Todd, Nystrom & Cohen 2013. *Neuroimage*. 98:506–512.
- Woolgar A, Hampshire A, Thompson R, Duncan J. 2011. Adaptive coding of task-relevant information in human frontoparietal cortex. *J Neurosci*. 31:14592–14599.
- Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Z6llei L, Polimeni JR, et al. 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol*. 106:1125–1165.
- Zhang J, Kriegeskorte N, Carlin JD, Rowe JB. 2013. Choosing the rules: distinct and overlapping frontoparietal representations of task rules for perceptual decisions. *J Neurosci*. 33:11852–11862.