

Hierarchical Organization of Cognition Reflected in Distributed Frontoparietal Activity

Ausaf A. Farooqui, Daniel Mitchell, Russell Thompson, and John Duncan

Medical Research Council Cognition and Brain Sciences Unit, Cambridge, United Kingdom CB2 7EF

Organization of behavior into a nested hierarchy of tasks and subtasks is characteristic of purposive cognition in humans. While frontoparietal regions have been shown to represent many kinds of task events, their representation of task/subtask structure has not been directly investigated. On each trial of the current study, participants carried out a sequence of four visual target detections organized by task context into subtasks of different structure (three and one or two and two). Through extended regions of frontoparietal cortex, activity elicited by target detections depended upon the hierarchical level of the episode completed. Target detections completing the entire trial elicited greatest activity, followed by targets completing a subtask, and finally targets within one subtask. Results depended on task and subtask completion, rather than the complexity of the next task stage to be established. We suggest that, through large regions of frontoparietal cortex, control representations direct each step of a behavioral program. Completion of a subtask revises control representations related just to this subtask, leaving those related to the overarching task episode intact, while completion of the entire task revises the entire assembly of representations.

Introduction

A central feature of purposive behavior is parcellation of the main goal (e.g., preparing breakfast) into smaller subgoals (preparing coffee and buttered toast). The whole sequence of behavioral events (e.g., pour water → add coffee → add cream → toast bread → spread butter) may be organized into distinct tasks and subtasks, corresponding to achievement of goals at different levels. In the current example, the event ‘add cream’ completes the subtask of preparing coffee, while the event ‘spread butter’ completes the higher task of preparing breakfast. From this perspective, the events making up a task episode can be conceptualized as having a hierarchy based on the level of episode they complete—events completing a task episode would be higher in this hierarchy than those lying within it.

Prefrontal cortex is critical to purposive behavior (Luria, 1966). In behaving monkeys, prefrontal neurons code many kinds of information relevant to the current cognitive operation, including stimuli, responses, rules, working memory contents (Procyk and Joseph, 2001; Duncan, 2010). Human neuroimaging studies, using either adaptation or multivoxel pattern analysis to examine the detailed content of task representations, show extensive coding of task-relevant information across multiple frontoparietal regions (Li et al., 2007;

Woolgar et al., 2011). It is frequently proposed that adaptive frontoparietal representations are a source of cognitive control by biasing processing in other brain regions (Desimone and Duncan, 1995; Miller and Cohen, 2001). In such a framework, completion of each task step signals that a previous control representation should be dissolved and the next one instantiated (Sigala et al., 2008).

We know little of how frontoparietal control activity is influenced by hierarchical goal–subgoal structure. While several studies have linked hierarchical task control to the organization of prefrontal cortex, especially along the anteroposterior axis, these studies have not been explicitly concerned with task–subtask structure (Koechlin et al., 2003; Badre and D’Esposito, 2007; Christoff et al., 2009). Badre and D’Esposito (2007) showed that frontal activity was more posterior for simple one-level decisions (e.g., red > left key) than for two-level decisions in which one stimulus feature indicated how the other should be processed. Koechlin et al. (2003) showed anterior frontal activity when the context of a whole task block determines how a stimulus should be interpreted. Although control in such tasks is certainly hierarchical, this sense of hierarchy is different from a comparison of task and subtask completion in complex, multistep behavior.

To investigate this issue, we devised a sequential target detection task with a hierarchical structure, such that some target detections completed subtasks and others completed a higher level task. We reasoned that when a subtask is completed, its specific content loses relevance but higher level task representations must remain in behavioral control. Achievement of the higher task, by comparison, completes a more substantial behavioral episode, causing greater revision of control representations and perhaps producing stronger and more extensive frontoparietal activity.

Received Feb 8, 2012; revised July 12, 2012; accepted July 28, 2012.

Author contributions: A.A.F., D.M., R.T., and J.D. designed research; A.A.F. and D.M., performed research; A.A.F., analyzed data; A.A.F. and J.D. wrote the paper.

This work was funded by the Medical Research Council (UK) Intramural Program MC-A060-5PQ10. A.A.F. was supported by a scholarship from the Gates Cambridge Trust, Cambridge, UK.

The authors declare no competing financial interests.

Correspondence should be addressed to Ausaf A. Farooqui, Medical Research Council and Cognition and Brain Sciences Unit, 15 Chaucer Road, Cambridge, UK CB2 7EF. E-mail: ausaf.farooqui@gmail.com.

DOI:10.1523/JNEUROSCI.0598-12.2012

Copyright © 2012 the authors 0270-6474/12/3217373-09\$15.00/0

Materials and Methods

Tasks. Three experiments were conducted. In the main experiment (Experiment 1), participants monitored a series of letters presented at a rate of 1.3 s/item on a computer screen (Fig. 1a). The task was to detect four target letters in turn. A three-letter word (e.g., ‘CAT’) presented at the start of the trial indicated the first three targets (T1, T2 T3; e.g., Fig. 1a, T1 ‘C’, T2 ‘A’, T3 ‘T’). The fourth target (T4) was always the letter X. At the end of the 52 s letter stream, the participants were cued to indicate whether all four targets had appeared in the specified order (true on 50% of trials called complete trials, with remaining incomplete trials not containing the final X).

At the start of each trial, the instruction screen specifying the first three targets (Fig. 1a, ‘CAT’) was presented for 3500 ms. The letter sequence began after a jittered gap of 1000–5000 ms, each letter being presented for 800 ms with a gap of 500 ms in between consecutive letters. The letter stream consisted of a total of 40 letter presentations. Participants began by monitoring for T1 (Fig. 1a, ‘C’), at the detection of which the search for T2 started and so on. The sequence of search was important; each target could only be searched for after the previous one had been detected and was irrelevant when it occurred before this point. For example, while Xs might appear at any point in the sequence, only an X appearing after sequential detection of T1, T2, and T3 was relevant. Within each trial, the intertarget interval varied randomly between 1.3 and 32 s. Responses were made when the letter stream terminated with a probe (‘X?’ displayed for 2000 ms) asking whether the letter X had appeared at the relevant position i.e., after the detection of T1–T3. Responses were made on a button box positioned under the participant’s right hand (index finger for ‘yes’, middle finger for ‘no’). A variable intertrial interval of 2000–7000 ms preceded onset of the cue for the next trial.

All stimuli were centered on the screen, visible from the participant’s position in the scanner via a mirror mounted within the head coil. Letters subtended a visual angle of 2° vertically. The experiment was controlled by a program written in Visual Basic. Participants learnt the task in a 10 min prescan practice session and then proceeded to a scanning session of an hour, which was divided into three separate scanning runs, each consisting of 14 trials.

Hierarchical structure of the trial was emphasized in the initial instructions that mentioned “two tasks are to be done on each trial; search for the letters of the cued word, and then search for X.” This was reinforced by the distinction of the sources that guided these two tasks: a cue word to be kept in working memory as it changed from trial to trial, and the letter X that was constant across trials. According to such a structure, T1 and T2 completed episodes at the lowest level (component letters of the initial target word; level 1; Fig. 1a), T3 completed the episode at the next higher level (target word completed; level 2), while X completed the entire search episode (level 3). Participants were informed that correct ‘yes’ responses on complete trials increased the current score by +1, whereas a

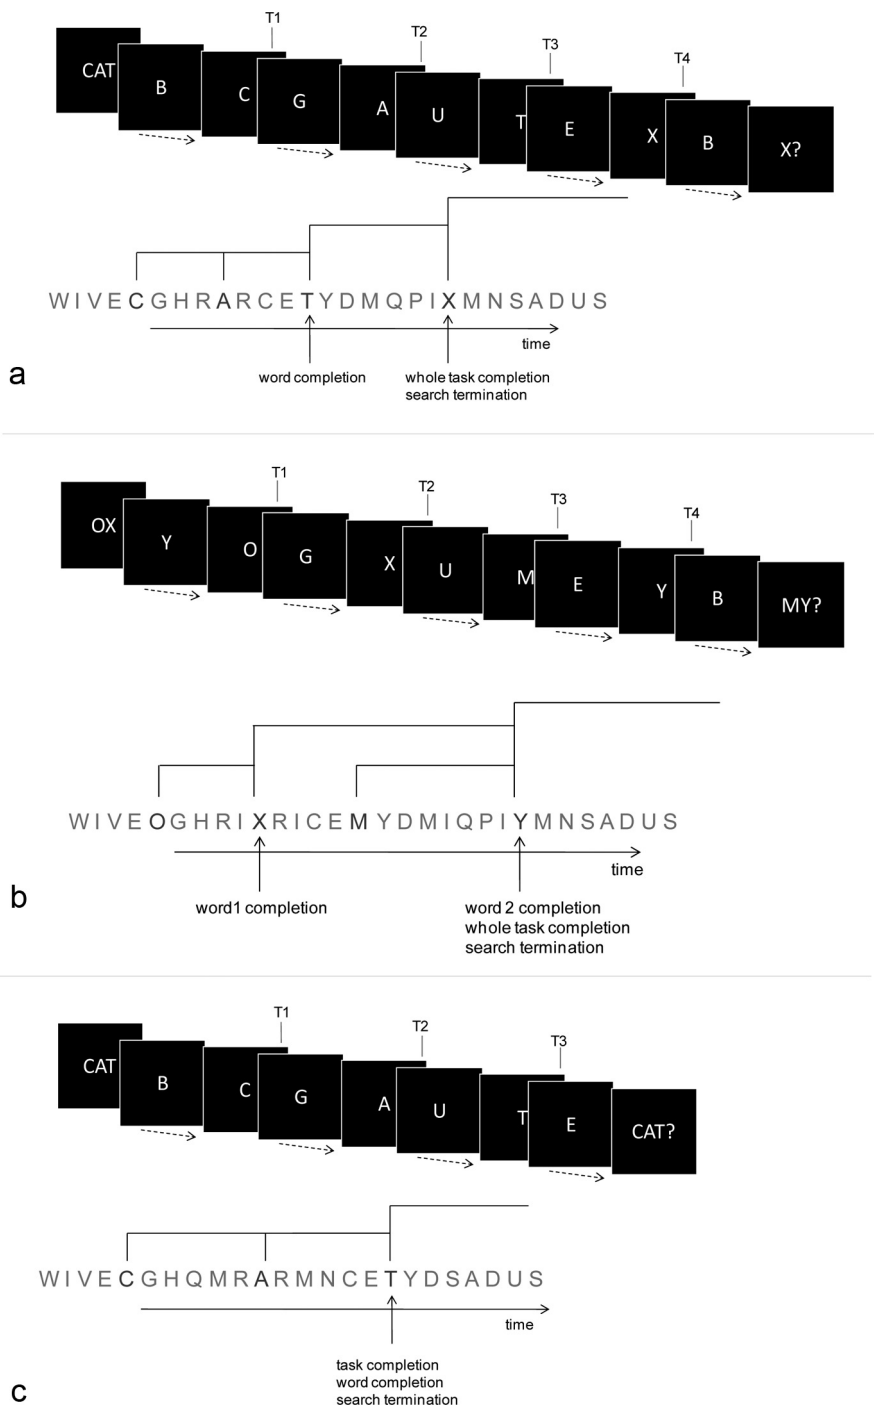


Figure 1. *a*, Structure of a typical trial in Experiment 1. Trials began with a three-letter cue word. The three letters of this word were to be covertly detected, in the correct order, in the ensuing letter stream; after all three had been detected, search for the letter X began. The complete sequence of four target letters appeared in only half of the trials. The letter stream ended with a probe asking whether the letter X had appeared in the correct sequential position, i.e., following the three letters of the target word. Thus the first two targets (T1 and T2) completed subgoals at the lowest level (component letters of the first target word; level 1); the third target, T3, completed a subgoal at the next highest level (complete target word; level 2), while the fourth target, T4 (X), completed the whole goal of the task (level 3). Dotted arrows indicate the variable number of nontarget letters between initial cue, successive targets, and final probe. *b*, Structure of a typical trial in Experiment 2. The cue consisted of a two-letter word, the detection of which was followed by detection of the letters of word ‘MY’, the latter being constant across trials. In this scheme T1 and T3 completed a level 1 episode, T2 completed a level 2 episode, and T4 completed the whole task (level 3). *c*, Structure of a typical trial in Experiment 3, similar to Experiment 1 but with only three targets, T1–T3.

correct ‘no’ response on an incomplete trials did nothing to the score; an incorrect response decreased the current score by –1.

In Experiment 1, the events completing more abstract episodes also occurred progressively later in the trial. These were dissociated in Exper-

iment 2. Now (Fig. 1*b*), a two-letter word was shown at the beginning of each trial. After these two target letters had been detected in their proper sequence, participants searched for the letters of the word 'MY'. This second word was constant across the experiment and hence was not displayed in the cue, analogous to the case with X in Experiment 1.

Other procedural details followed those of Experiment 1. The relation between T2 and T3 in Experiment 2 was thus the reverse of that in Experiment 1; T2 now completed a level 2 episode, while T3 was relegated to level 1. Comparison of activities at these two events across the two experiments allowed for a separation between the effects of the ordinal position of the event within the trial and the hierarchical level of the episode completed by it.

In Experiments 1 and 2, when each of targets T1 to T3 was detected, the previous search was to be abandoned and a new one begun. In contrast, detection of the final target T4 (X in Experiment 1 and Y in Experiment 2) was associated with termination of the search and an increase in the current score, which could be interpreted as a form of reward. To examine the significance of these features, in a follow-up experiment (Experiment 3) the task was slightly modified—T3 detection completed the entire trial (i.e., task episode) and thus there was no requirement to search for a T4 (Fig. 1*c*). Further, an explicit reward (of 10p) was given at the completion of the task. At the end of each trial participants were probed to indicate whether all the three target letters had appeared in the correct order (Fig. 1*c*, 'CAT?'). In this experiment, accordingly, activity specifically linked to search termination/reward should have been seen at T3. Note that in this experiment, target strings were a mixture of words and pronounceable nonwords.

Participants. Eighteen participants (10 female; mean age, 22.5 ± 3.6 years) in Experiment 1, twenty-one (15 females; mean age, 24.5 ± 4.1 years) in Experiment 2, and fifteen (8 females; mean age, 28.1 ± 6.4 years) in Experiment 3 were recruited from the Medical Research Council and Cognition and Brain Sciences Unit (Cambridge, UK) volunteer panel. Participants were right handed and had normal or corrected vision. Informed consent was taken, and the participants were reimbursed for their time. The study had the approval of Hertfordshire Local Research Ethics Committee.

Acquisition. fMRI data were acquired using a Siemens 3T Tim Trio scanner with a 12 channel head coil. A sequential descending T2*-weighted echo planar imaging (EPI) acquisition sequence was used with the following parameters: acquisition time, 2000 ms; echo time, 30 ms; 32 oblique slices with slice thickness of 3 mm and a 0.75 mm interslice gap; in-plane resolution, 3.0×3.0 mm; matrix, 64×64 ; field of view, 192 mm; flip angle, 78°. T1-weighted MP RAGE structural images were also acquired for all participants (slice thickness, 1.0 mm; resolution, $1.0 \times 1.0 \times 1.5$ mm; field of view, 256 mm; 160 slices). Experimental task started after 10 "dummy" scans had been acquired. These were discarded from the general linear model to allow for T1 equilibration effects.

Analysis. The fMRI data were analyzed using SPM5 (Wellcome Department of Imaging Neuroscience, London, England; www.fil.ion.ucl.ac.uk). Before statistical analysis, all EPI volumes were slice-time corrected using the first slice as a reference and then realigned into a standard orientation using the first volume as a reference. These realigned volumes were then normalized into the Montreal Neurological Institute (MNI) space and spatially smoothed using an 8 mm full-width half-maximum Gaussian kernel. During the normalization stage, voxels were resampled to a size of $3 \times 3 \times 3$ mm. The time course of each voxel was high pass filtered with a cutoff period of 90 s.

Statistical analysis was carried out using a general linear model. To capture activity related to target detection, a 16 s epoch starting from each target onset was modeled using a finite impulse response (FIR) basis set of eight 2 s boxcar regressors. In this way the response to target detections could be modeled without making any assumptions about the shape of the hemodynamic response. Additionally, the cue and probe were modeled using epoch regressors of width equal to the duration of respective events, convolved with a basis function representing the canonical hemodynamic response. Movement parameters and block means were included as covariates of no interest. Parameter estimates for each regressor were calculated from the least-squares fit of the model to the data, and

estimates for individual participants were entered into a random effects group analysis.

Whole-brain comparisons were performed using paired *t* tests on the relevant contrast values from each participant's first-level analysis. Unless otherwise specified, all results are reported at a threshold of $p < 0.01$ and corrected for multiple comparisons using the false discovery rate. Coordinates for peak activation are reported using an MNI template.

To capture frontoparietal regions widely engaged in cognitive control, 10 regions of interest (ROIs) were created as spheres of 10 mm radius at coordinates that have been shown to be consistently active in varied tasks (Dosenbach et al., 2006; Duncan, 2006). The ROIs (in MNI space) were bilateral inferior frontal sulcus (IFS; central coordinate ± 41 23 29), bilateral intraparietal sulcus (IPS; ± 37 -56 41), bilateral anterior insula extending into frontal operculum (AI/FO; ± 35 18 3), anterior cingulate cortex (ACC; 0 31 24), and presupplementary motor area (pre-SMA; 0 18 50), all taken from Duncan (2006); bilateral anterior prefrontal cortex ROIs (APFC; 27 50 23 and -28 51 15) were taken from Dosenbach et al. (2006). ROIs were constructed using the MarsBaR toolbox for Statistical Parametric Mapping or SPM (<http://marsbar.sourceforge.net>; Brett et al., 2002). Estimated data were averaged across voxels within each ROI using the MarsBaR toolbox, and the mean values were exported for analysis using SPSS.

Results

Experiment 1

Behavior

Average response time to the probe was 723 ± 15 ms, and accuracies for responses to the probe exceeded 97% for most participants (mean 97.1 ± 1.4), suggesting that they were indeed covertly detecting target letters.

Imaging

The FIR model provided estimates of activity in eight successive 2 s windows starting from each target onset. The first analysis identified brain regions that were responsive in general to target detection. To capture the positive phase of the response, analysis of variance compared estimates of brain activity across the first five time bins. Regions showing a significant main effect of time (familywise error corrected, $p < 0.05$) are shown in Figure 2, along with the average time course for each ROI (see Materials and Methods). Coordinates of peak activity in the whole-brain analysis are shown in Table 1. Consistent with other fMRI studies showing widespread brain activity associated with target detection (Jiang et al., 2000; Downar et al., 2001; Hon et al., 2006; Hampshire et al., 2007), activity involved the occipital cortex and spread dorsally to the parietal cortex and ventrally to the temporal cortex; anteriorly the activity extended into the frontal cortex and involved all its major sectors. Beyond the cortical regions, basal nuclei and cerebellum also showed loci of activity.

The time courses in Figure 2 make it evident that in many regions of the frontoparietal cortex there was a differential response to the four successive target detections. For further analysis, we calculated an index of phasic BOLD response by subtracting the estimate of the first FIR bin from the average of the second and third. Mean values of this index for each ROI appear in Figure 3. It is evident that X, which completes the overall task episode (level 3), shows the highest activity in all ROIs. In contrast, T1 and T2 (level 1 subtask) show the least. Further, the profile of activity across the four target events differs in different ROIs. An ANOVA with ROI and target type as factors confirmed the significant effect of target type ($F_{(3, 51)} = 13.5$, $p < 0.001$) and its interaction with ROI ($F_{(27, 459)} = 2.0$, $p < 0.01$).

To look into these results further, activity indices for each ROI were examined by ANOVA, with target type (T1, T2, T3, X) and hemisphere (for lateral ROIs only) as factors. All ROIs showed a

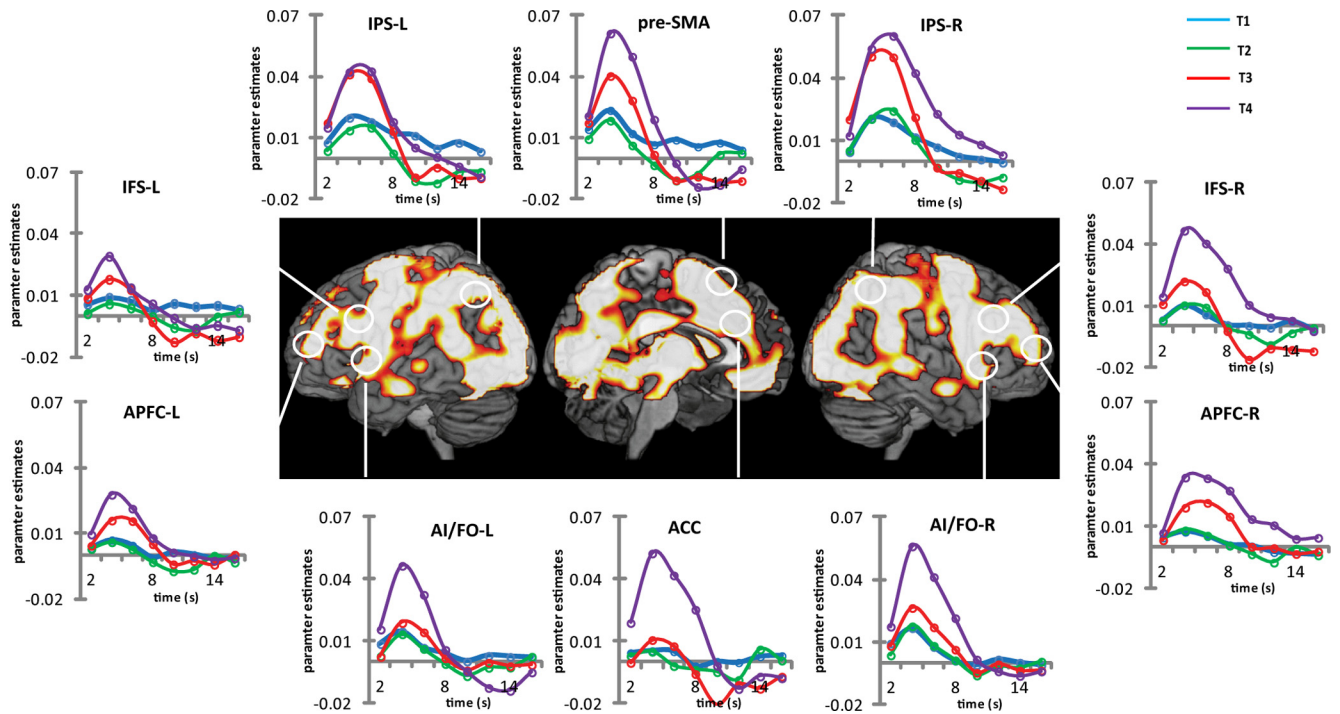


Figure 2. Experiment 1. The whole brain render shows areas where target detection significantly modulated brain activity (ANOVA comparing the first 5 FIR regressors after each target detection event; results thresholded at family-wise error corrected $p < 0.05$). For ROIs, estimates of eight FIR regressors linked to each target event were extracted and plotted to construct the time course of BOLD activity (blue, T1; green, T2; red, T3; purple, X). L, Left; R, right.

Table 1. Experiment 1: representative coordinates where activity was significantly modulated by target detection

Lobe	Region	Laterality	$F_{(4,68)}$	XYZ (mm)
Frontal	Insula	L	75.93	-28, 28, 0
		R	60.86	30, 26, -2
	Inferior frontal	L	42.04	-48, 6, 28
		R	32.37	50, 8, 28
	Middle frontal	L	45.78	-28, 0, 58
		R	46.08	28, 0, 54
	SMA	L	65.32	-6, 12, 46
		R	54.18	2, 8, 58
Cingulate	L	30.70	-10, 20, 32	
	R	52.56	6, 16, 46	
Parietal	Superior parietal	L	59.08	-22, -64, 50
		R	37.41	28, -56, 50
	Inferior Parietal	L	58.78	-30, -52, 48
		R	47.24	34, -44, 42
	Precuneus	L	42.72	-12, -68, 48
		R	35.95	14, -66, 46
Occipital	Middle occipital	L	46.06	-36, -90, -4
		R	40.22	30, -66, 34
	Calcarine	L	28.91	-4, -86, -2
		R	23.24	2, -78, 6
Temporal	Superior temporal	L	20.68	-54, -44, 14
		R	21.35	60, -44, 22
	Middle temporal	L	28.24	-44, -62, -4
		R	22.41	50, -22, -8
Subcortical	Cerebellum	L	25.02	-38, -60, -26
		R	24.55	34, 60, -26
	Pallidum	L	33.58	-12, 8, -4
		R	40.39	30, 18, 2

XYZ are the MNI coordinates along x, y, and z axes.

main effect of target type (Table 2). Among the bilateral ROIs, IFS, APFC, and IPS showed a main effect of laterality reflecting greater activity on the right; in IFS and APFC this effect of laterality interacted significantly with that of target type, showing that the differential activity across the targets was also greater on the right.

To probe the significant effect of target type further, pairwise comparisons compared T3 (level 2 episode completion) with the average of T1 and T2 (level 1). APFC, ACC, and right IPS showed significantly higher activity to T3 (Table 3). Among other target events, X was significantly higher than T1 and T2 in all ROIs except left IFS (Table 3). Across all ROIs, the tendency was also for greater activity to X than T3. This was significant, however, only in right IFS, right AI/FO, ACC, and pre-SMA. There was no significant difference between T1 and T2 in any ROI.

To check for other brain areas with differential activity to the different target events, two whole-brain analyses examined similar contrasts. In the first (Fig. 4a), we contrasted activity for T3 with mean activity for T1 and T2. Consistent with the ROI results, the only regions to show greater activity for T3 were anterior prefrontal cortex, ACC, and parts of pre-SMA along with right lateral parietal cortex and precuneus. In the second analysis, we contrasted activity to the final target X with mean activity to T1 and T2. Results of this contrast appear in Figure 4b. Again, results paralleled those from the ROI analyses. Significant activity was found in multiple frontoparietal regions and was visibly greater on the right. In the right prefrontal cortex it involved inferior frontal sulcus, frontal operculum, and the anterior insula, extending anteriorly through the middle frontal gyrus up to anterior prefrontal regions. In the left prefrontal cortex, only the anterior prefrontal regions and rolandic operculum were involved. Medially, stronger activity for X compared to T1/T2 was seen in the anterior cingulate and the pre-SMA on both

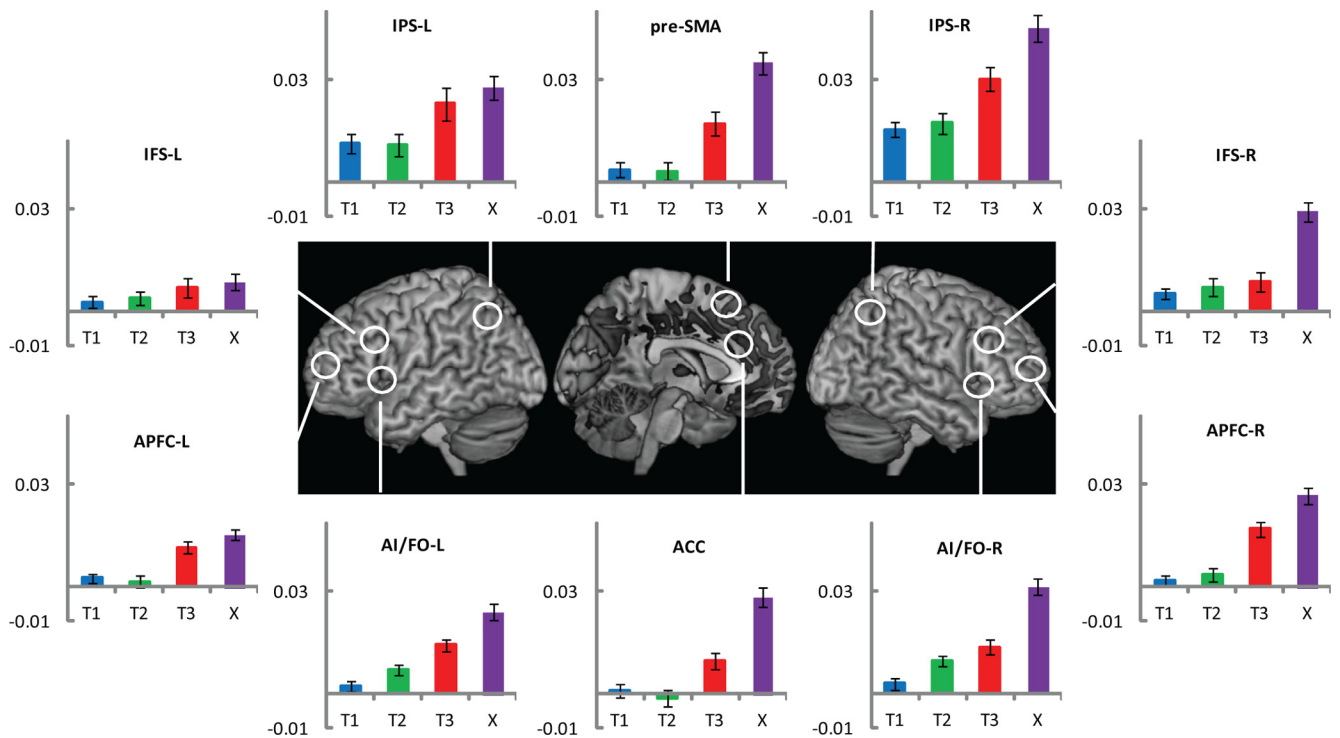


Figure 3. Experiment 1. Comparison of phasic activity in response to various target events. Plots show activity index for each target. This was derived for each target event by subtracting the estimate of the first FIR regressor from the average of the second and third. The error bars represent one standard error of the mean.

Table 2. Experiment 1: effects of target position, hemisphere, and their interaction in each ROI

ROI	Target position $F_{(3,51)}$	Hemisphere $F_{(1,17)}$	Target \times hemisphere $F_{(3,51)}$
IFS	4.9	11.4	14.2
IPS	6.3	9.9	2.7
AI/FO	7.1	2.6	1.4
APFC	9.1	4.9	3.4
ACC	9.0	—	—
Pre-SMA	11.2	—	—

Values in bold were significant at $p < 0.05$.

Table 3. Experiment 1: $t_{(17)}$ values from the pairwise comparisons of different target detections in each ROI

ROI	Laterality	T1 and T2 versus T3	T1 and T2 versus X	T3 versus X
IFS	L	1.0	1.2	0.3
	R	0.7	4.9	3.8
IPS	L	1.9	2.6	0.6
	R	2.5	4.4	1.8
AI/FO	L	2.0	2.9	1.6
	R	1.2	3.7	3.2
APFC	L	2.8	3.7	0.8
	R	3.3	4.8	1.7
ACC	n.a.	2.6	4.4	3.1
Pre-SMA	n.a.	2.2	4.8	3.0

Values in bold were significant at a Holm–Bonferroni corrected threshold of 0.05 for multiple comparisons within each ROI.

sides. Posteriorly, significant activity was found in the intraparietal sulcus extending into the temporoparietal junctions, posterior cingulate, precuneus, cuneus, and parts of the calcarine and lingual regions.

To summarize, activity was strongest and most widespread in response to the completion of the highest level episode (level 3). Selected regions, especially APFC, medial PFC (ACC and pre-SMA), and right parietal regions, also showed greater activity for episode completion at level 2 compared to that at level 1. The effects were generally stronger in right-sided regions.

Experiment 2

Experiment 1 confounds position in the trial with the level of episode completed in that T3 had a higher ordinal position than T1 and T2, apart from completing a higher level (level 2) episode. In Experiment 2 we dissociated these factors. Materials were constructed such that T2 rather than T3 completed a level 2 episode.

Behavior

Average response time was 727 ± 35 ms, while average accuracy was 95.2% (± 1.1).

Imaging

Figure 5 shows indices of activation for the four target events. It is evident that, across most regions, the activity at T3 was not higher than that at T2; in several regions indeed, activity was greater for T2. To compare results from Experiments 1 and 2, we used the index: T3 activity – T2 activity. A comparison of this index across the two experiments is mathematically identical to the main effect of chunk completion (T3 of Experiment 1 + T2 of Experiment 2 vs T2 of Experiment 1 + T3 of Experiment 2) without the confound of serial position. This index was significantly more positive in Experiment 1 than in Experiment 2 for right IPS ($t_{(31)} = 2.6, p < 0.01$), pre-SMA ($t_{(31)} = 2.2, p < 0.02$), left IPS ($t_{(31)} = 1.9, p < 0.05$), and right APFC ($t_{(31)} = 1.7, p < 0.05$), with all tests one-tailed. A trend was also seen in the left AI/FO ($t_{(31)} = 1.4, p < 0.1$). The results confirm that a distributed set of frontoparietal regions is sensitive not to the ordinal position within the trial, but to the hierarchical structure of task episodes.




Figure 4. Experiment 1. *a*, Whole-brain render showing areas that had higher activation index for T3 compared to T1 and T2. *b*, Areas with greater activity for X than T1 and T2.

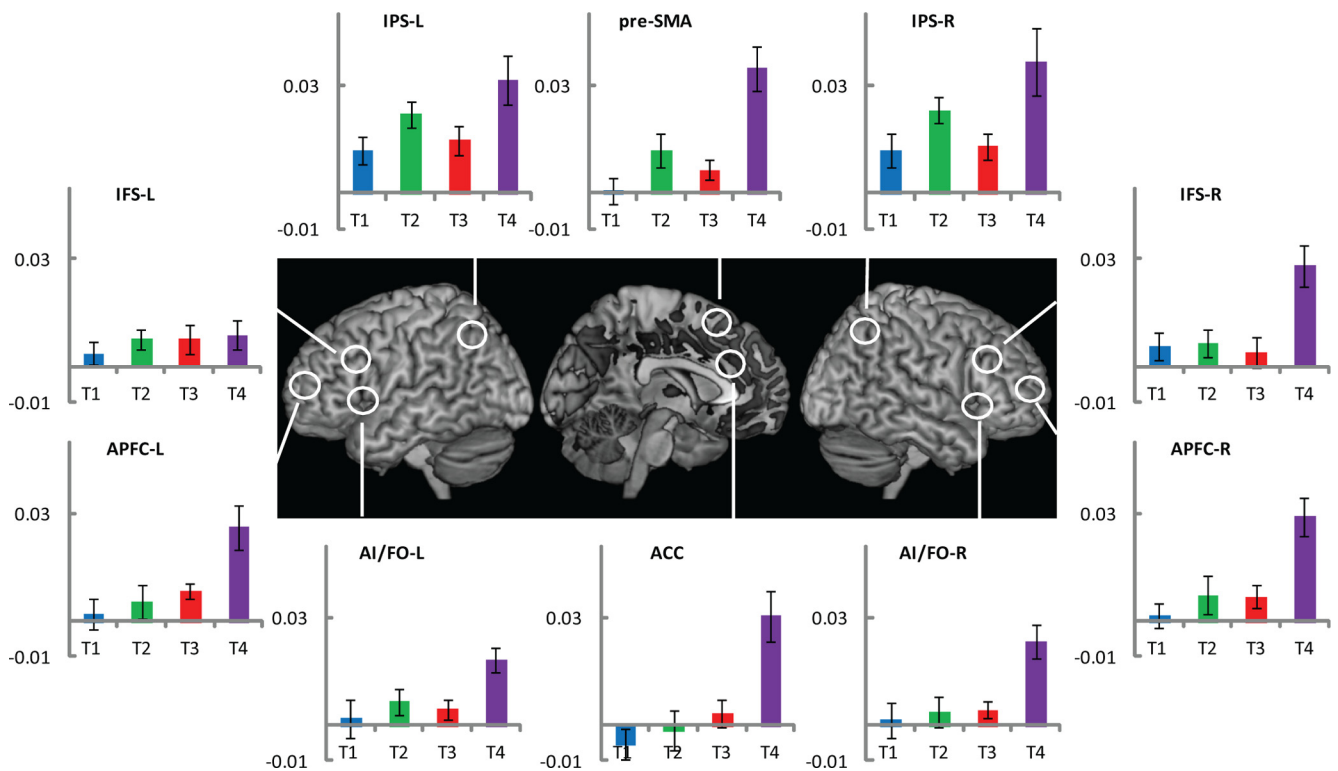


Figure 5. Experiment 2. Comparison of activity index in response to various target events. Across multiple ROIs, compared to Figure 3, note relative increase in T2 activity and decrease in T3 activity.

We also confirmed that the response to target events at the same level of hierarchy was the same in Experiments 1 and 2. A comparison of the difference of activity between T1 and T4 (level 1 vs level 3) across the two experiments revealed no significant

difference in any ROI ($p > 0.2$ in all comparisons). Similarly, a comparison between T3 vs T1 and T2 in Experiment 1 and T2 vs T1 and T3 in Experiment 2 was insignificant ($p > 0.13$ in all comparisons).

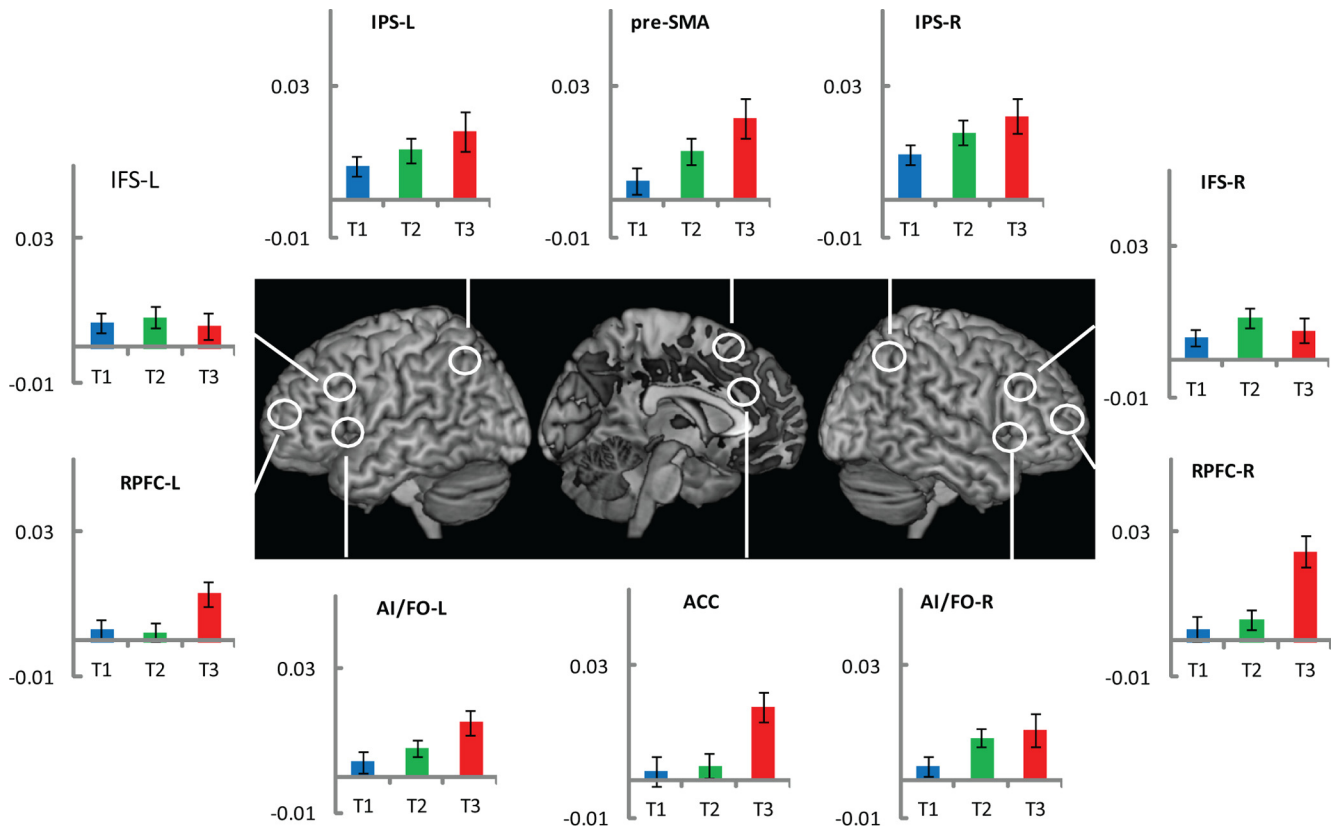


Figure 6. Experiment 3. Comparison of activity index in response to the three target events. Results were closely similar to those of Experiment 1.

Experiment 3

In Experiment 3, only the three letters of the target letter string (T1 to T3) were searched for on each trial. By associating T3 with task termination and reward—which in Experiment 1 were linked to the final target X (T4)—we planned to investigate the importance of these factors. The structure of the search episode in the current experiment is identical to that of the search episode for the first three target detections in Experiment 1. For brain regions where activity is sensitive only to task episode structure, the contrast of T3 vs T1/T2 should show similar results in the two experiments. Any influence of task termination and reward should be shown by additional T3 activity in Experiment 2, corresponding to the strong activity seen in Experiment 1 for the final X (T4).

Behavior

Average response time was 770 ± 18 ms, while average accuracy was 96.1% (± 2.7).

Imaging

Figure 6 shows indices of activation for the three target events. To directly compare Experiments 1 and 3, we carried out an ANOVA with factors target type (T1, T2, and T3), hemisphere (for lateral ROIs only), and experiment. No ROI showed a significant interaction, either between experiment and target type ($p > 0.3$ in all comparisons) or between experiment, target type, and hemisphere ($p > 0.08$ in all comparisons).

We then compared the increase in activity at the final targets (X in Experiment 1 and T3 in Experiment 3) over the initial targets (mean of T1 and T2), across Experiments 1 and 3. The effect of hierarchy on activity of these regions would predict that this increase would be greater in Experiment 1, since the comparison here is between level 3 over level 1 completion (X vs T1 and

T2), while in Experiment 3 the comparison involves level 2 versus level 1 (T3 over T1 and T2). This was indeed the case in right IFS, IPS, AI/FO, and pre-SMA ($F_{(2,74)} > 3.2, p < 0.05$ in all comparisons). The effect was stronger in right sided ROIs ($F_{(2,74)} > 2.8, p < 0.03$).

The results suggest that associating T3 with task termination and reward has little effect on frontoparietal activity. Instead it is the hierarchical level of the task episode completed that is the major determinant of activity in these regions.

Discussion

These experiments show that activity in various regions of the frontoparietal cortices is elicited at the completion of behavioral episodes, the magnitude of which depends upon the hierarchical level of the completed episode. Completion of the lowest level subtasks (level 1) led to the weakest frontoparietal activity. At the completion of next higher level (level 2) episodes, stronger activity was seen especially in anterior and medial prefrontal regions and IPS. Completion of the highest level episode (level 3) resulted in greatest and most widespread activity.

The first question raised by these results is why the completion of discrete task episodes should elicit frontoparietal activity. Plausibly, a task episode would be characterized by a unique neurocognitive configuration of multiple neural and mental processes. Such a configuration would create the relevant cognitive focus by coordinating activity in multiple mental domains to enable task-relevant processing across all of these domains. While much is unknown, representations of purpose/goal (Miller and Cohen, 2001; Duncan, 2006) and task rules (Sakai, 2008; Dumontheil et al., 2011) could play important roles in this process. At the end of task episodes, this neurocognitive configuration would be revised or completely dismantled, depending upon the

nature of the future behavioral episode. A plausible interpretation of our results is that this restructuring at the end of an episode leads to widespread frontoparietal activity.

Further, greater change in cognitive focus at the completion of higher level episodes would correspond to greater frontoparietal activity. As lower level episodes are completed, only their corresponding neurocognitive configurations lose their relevance, while the configurations corresponding to the higher level subsuming episode would remain intact; on the other hand, completion of the highest level episode changes this structure in its entirety. The corresponding neural activity can be expected to parallel this; hence in our experiments the lowest level completions elicited minimal activity, while the final completions elicited maximal activity. This is in line with earlier findings that it is more effortful to switch across hierarchically higher level task episodes (Schneider and Logan, 2006), and that errors on task events at the junction of higher level episodes cause greater behavioral disorganization than lower level errors (Schwartz et al., 1991).

Highest level completions in the above experiments were distinct from others in terms of two additional factors—termination of the on-going task (visual search) and a feeling of reward associated with the successful completion of the trial. These, however, appeared to be less significant. T3 of Experiment 3 was similar to T4 of Experiment 1 in terms of these factors, since visual search terminated successfully at both. T3 of Experiment 3 additionally had a monetary reward attached to it. However, in terms of the hierarchical level of the episode completed, this T3 was identical to the T3 of Experiment 1 (Fig. 1*a,c*). As would be expected if the hierarchy of the episode completed was the major determinant of frontoparietal activity, T3 in the two experiments elicited very similar patterns of frontoparietal activity (Figs. 3, 6). Hierarchical changes in the representations organizing the behavior, therefore, seem to be the most plausible explanation of our results. Note that our finding does not necessarily contradict studies showing effects of reward on frontoparietal activity (Beck et al., 2010), since our results only show that task episode completion activity is not additionally modulated by the knowledge of future reward attainment; they do not imply that reward does not per se affect frontoparietal activity.

Another possibility raised by the results of Experiment 1 was that steps closer to the final goal elicit increasingly higher activity, providing an alternative explanation for greater activity at T3 over T1 and T2. Experiment 2 dissociated ordinal position in the trial from the level of episode completed, with T2 now completing a higher level episode than T3. Several frontoparietal regions—including IPS, APFC and pre-SMA—showed a significantly different pattern of activity compared to Experiment 1, with accompanying trends elsewhere. The data confirm that widespread regions of frontoparietal cortex are sensitive to the hierarchical organization of the search episode.

Earlier investigations into changes in task-related representations have approached it with the concept of task set (Sakai, 2008), understood as a cognitive structure linking the stimulus to the relevant response as per the task rules. It has frequently been proposed that restructuring the task set in response to changes in task rules requires cognitive control (Monsell, 2003), which results in widespread frontoparietal activity (Brass et al., 2005; Savine and Braver, 2010; Kim et al., 2011). While our results are broadly consistent with this idea, they have novel features. Restructuring of task set is often conceived as a process of loading the control structure for the next task episode; but in our data, the complexity of the following episode appeared not to be critical. In

our tasks the last target event involved only target detection, followed by a wait until the end of the letter sequence so that a response could be made. This appears less complex than the requirements triggered by earlier targets (target detection followed by establishment of the next search set), yet no brain region showed the lowest activity for the final target. Similarly, T3 activity was closely similar in Experiments 1 and 3, while the next trial phase was very different in the two experiments. Our results suggest that frontoparietal activity is determined by the hierarchical position of the episode just completed and not merely by the complexity of the next control episode to be established.

While the activity pattern in all frontoparietal regions (except left IFS) showed some effect of hierarchical organization of task episode, regional differences were evident. First, the effect was greater in right hemisphere regions (Table 2). Second, results depended on hierarchical level; while strong activity for level 3 decisions was very widespread, the contrast between level 2 and level 1 was most striking in APFC, IPS, and dorsomedial frontal regions and weak to absent in the IFS. The results suggest increasing the spread of frontoparietal activity with completion of successively higher-level episodes. More work is needed to show whether, beyond this quantitative effect, there may also be qualitative differences between IFS and other regions.

Other studies have shown similar frontoparietal activity at transitions across mental episodes generated while perceiving temporally extended sensory events. Zacks et al. (2001) had participants passively view movies of some episode of purposive behavior like gardening, etc.; they found increased activity in posterior prefrontal cortex, IPS, and precuneus along with the superior temporal sulcus and lateral occipital cortices at points that were a posteriori identified by the same subjects as transition between coherent segments (e.g., digging the ground, watering the plants, etc.). Similarly, Sridharan et al. (2008) had participants listen to musical pieces that had transitions between movements of symphony works. They also found activity in frontoparietal regions at such transitions. In line with our results, prefrontal cortical neurons show phasic peaks of activity at action sequence boundaries (Fujii and Graybiel, 2003).

A number of proposals link the hierarchical control of behavior to anteroposterior gradient in the lateral prefrontal cortex, with higher levels of the hierarchy linked to more anterior control (Koechlin et al., 2003; Badre and D'Esposito, 2007; Botvinick, 2008, for a review; Reynolds et al., 2012 for an alternative view). These have emphasized rather different issues from ours, e.g., decisions based only on current stimulus input versus those requiring memory of previous events (Koechlin et al., 2003), or decisions in which one stimulus feature instructs processing of another (Badre and D'Esposito, 2007). While it remains to be seen how these different senses of hierarchy relate to the current case of approach to a final goal through achievement of a series of component subgoals, our data give little evidence linking the anteroposterior axis of the prefrontal cortex (PFC) to task/subtask hierarchy. In Experiments 1 and 2, strong activity for level 3 completions (T4) was extremely widespread in both frontal and parietal cortex. Level 2 completions also gave stronger activity than level 1 completions across multiple frontoparietal ROIs, with the exception of IFS. Such results are not well described as a progressive posterior to anterior spread of activity for successively higher-level goal completions.

An important implication of the current results is that the design of the trial and the mental plan used by participants need to be taken into account in interpreting neural activity related to a cognitive event of interest, since the role of the cognitive event

in the overall plan of behavior is an important determinant of its elicited frontoparietal activity. For example, in light of the above experiments, what regions should be considered to be involved in visual target detection? Clearly, answering this requires taking into account the role of target detection within an overall task organization. More generally, extensive regions of frontoparietal cortex are involved in the cognitive restructuring that takes place when a current cognitive episode is completed, with the degree of activity reflecting the position of the completed episode in the hierarchically organized cognitive plan.

References

- Badre D, D'Esposito M (2007) Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *J Cogn Neurosci* 19:2082–2099. [CrossRef Medline](#)
- Beck SM, Locke HS, Savine AC, Jimura K, Braver TS (2010) Primary and secondary rewards differentially modulate neural activity dynamics during working memory. *PLoS ONE* 5:e9251. [CrossRef Medline](#)
- Botvinick MM (2008) Hierarchical models of behavior and prefrontal function. *Trends Cogn Sci* 12:201–208. [CrossRef Medline](#)
- Brass M, Derrfuss J, Forstmann B, von Cramon DY (2005) The role of the inferior frontal junction area in cognitive control. *Trends Cogn Sci* 9:314–316. [CrossRef Medline](#)
- Brett M, Johnsrude IS, Owen AM (2002) The problem of functional localization in the human brain. *Nat Rev Neurosci* 3:243–249. [CrossRef Medline](#)
- Christoff K, Keramiatian K, Gordon AM, Smith R, Mädlar B (2009) Prefrontal organization of cognitive control according to levels of abstraction. *Brain Res* 1286:94–105. [CrossRef Medline](#)
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222. [CrossRef Medline](#)
- Dosenbach NU, Visscher KM, Palmer ED, Miezin FM, Wenger KK, Kang HC, Burgund ED, Grimes AL, Schlaggar BL, Petersen SE (2006) A core system for the implementation of task sets. *Neuron* 50:799–812. [CrossRef Medline](#)
- Downar J, Crawley AP, Mikulis DJ, Davis KD (2001) The effect of task relevance on the cortical response to changes in visual and auditory stimuli: an event-related fMRI study. *Neuroimage* 14:1256–1267. [CrossRef Medline](#)
- Dumontheil I, Thompson R, Duncan J (2011) Assembly and use of new task rules in frontoparietal cortex. *J Cogn Neurosci* 23:168–182. [CrossRef Medline](#)
- Duncan J (2006) EPS mid-career award 2004: brain mechanisms of attention. *Q J Exp Psychol (Hove)* 59:2–27. [CrossRef](#)
- Duncan J (2010) The multiple-demand (MD) system of the primate brain: mental programs for intelligent behavior. *Trends Cogn Sci* 14:172–179. [CrossRef Medline](#)
- Fujii N, Graybiel AM (2003) Representation of action sequence boundaries by macaque prefrontal cortical neurons. *Science* 301:1246–1249. [CrossRef Medline](#)
- Hampshire A, Duncan J, Owen AM (2007) Selective tuning of the blood oxygenation level-dependent response during simple target detection dissociates human frontoparietal subregions. *J Neurosci* 27:6219–6223. [CrossRef Medline](#)
- Hon N, Epstein RA, Owen AM, Duncan J (2006) Frontoparietal activity with minimal decision and control. *J Neurosci* 26:9805–9809. [CrossRef Medline](#)
- Jiang Y, Haxby JV, Martin A, Ungerleider LG, Parasuraman R (2000) Complementary neural mechanisms for tracking items in human working memory. *Science* 287:643–646. [CrossRef Medline](#)
- Kim C, Johnson NF, Cilles SE, Gold BT (2011) Common and distinct mechanisms of cognitive flexibility in prefrontal cortex. *J Neurosci* 31:4771–4779. [CrossRef Medline](#)
- Koechlin E, Ody C, Kouneiher F (2003) The architecture of cognitive control in the human prefrontal cortex. *Science* 302:1181–1185. [CrossRef Medline](#)
- Li S, Ostwald D, Giese M, Kourtzi Z (2007) Flexible coding for categorical decisions in the human brain. *J Neurosci* 27:12321–12330. [CrossRef Medline](#)
- Luria AR (1966) Higher cortical functions in man. Oxford: Basic Books.
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202. [CrossRef Medline](#)
- Monsell S (2003) Task switching. *Trends Cogn Sci* 7:134–140. [CrossRef Medline](#)
- Procyk E, Joseph JP (2001) Characterization of serial order encoding in the monkey anterior cingulate sulcus. *Eur J Neurosci* 14:1041–1046. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Sakai K (2008) Task set and prefrontal vortex. *Annu Rev Neurosci* 31:219–245. [CrossRef Medline](#)
- Savine AC, Braver TS (2010) Motivated cognitive control: reward incentives modulate preparatory neural activity during task-switching. *J Neurosci* 30:10294–10305. [CrossRef Medline](#)
- Schneider DW, Logan GD (2006) Hierarchical control of cognitive processes: switching tasks in sequences. *J Exp Psychol Gen* 135:623–640. [CrossRef Medline](#)
- Schwartz MF, Reed ES, Montgomery MW, Palmer C, Mayer NH (1991) The quantitative description of action disorganization after brain damage: a case study. *Cogn Neuropsychol* 8:381–414. [CrossRef](#)
- Sigala N, Kusunoki M, Nimmo-Smith I, Gaffan D, Duncan J (2008) Hierarchical coding for sequential task events in the monkey prefrontal cortex. *Proc Natl Acad Sci U S A* 105:11969–11974. [CrossRef Medline](#)
- Sridharan D, Levitin DJ, Menon V (2008) A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proc Natl Acad Sci U S A* 105:12569–12574. [CrossRef Medline](#)
- Woolgar A, Thompson R, Bor D, Duncan J (2011) Multi-voxel coding of stimuli, rules, and responses in human frontoparietal cortex. *Neuroimage* 56:744–752. [CrossRef Medline](#)
- Zacks JM, Braver TS, Sheridan MA, Donaldson DI, Snyder AZ, Ollinger JM, Buckner RL, Raichle ME (2001) Human brain activity time-locked to perceptual event boundaries. *Nat Neurosci* 4:651–655. [CrossRef Medline](#)