

The Role of Frontopolar Cortex in Subgoal Processing during Working Memory

Todd S. Braver¹ and Susan R. Bongiolatti

Department of Psychology, Washington University, St. Louis, Missouri 63130

Received May 21, 2001; published online January 22, 2002

Neuroimaging studies have implicated the anterior-most or frontopolar regions of prefrontal cortex (FP-PFC, e.g., Brodmann's Area 10) as playing a central role in higher cognitive functions such as planning, problem solving, reasoning, and episodic memory retrieval. The current functional magnetic resonance imaging (fMRI) study tested the hypothesis that FP-PFC subserves processes related to the monitoring and management of subgoals, while maintaining information in working memory (WM). Subjects were scanned while performing two variants of a simple delayed response WM task. In the control WM condition, subjects monitored for the presence of a specific concrete probe word (LIME) occurring following a specific abstract cue word (FATE). In the subgoal WM condition, subjects monitored for the presence of any concrete probe word immediately following any abstract cue word. Thus, the task required semantic classification of the probe word (the subgoal task), while the cue was simultaneously maintained in WM, so that both pieces of information could be integrated into a target determination. In a second control condition, subjects performed abstract/concrete semantic classification without WM demands. A region within right FP-PFC was identified which showed significant activation during the subgoal WM condition, but no activity in either of the two control conditions. However, this FP-PFC region was not modulated by direct manipulation of active maintenance demands. In contrast, left dorsolateral PFC was affected by active maintenance demands, but the effect did not interact with the presence of a subgoal task. Finally, left ventral PFC regions showed activation in response to semantic classification, but were not affected by WM demands. These results suggest a triple dissociation of function within PFC regions, and further indicate that FP-PFC is selectively engaged by the requirement

to monitor and integrate subgoals during WM tasks.

© 2002 Elsevier Science (USA)

INTRODUCTION

The prefrontal cortex (PFC) has long been thought to play a central role in higher cognition (Fuster, 1989; Goldman-Rakic, 1987; Stuss and Benson, 1986). With the advent of high spatial resolution functional neuroimaging it has become possible to study the functional properties of specific PFC subregions. Consequently, there has been an explosion of research examining more fine-grained relationships between different PFC subregions and specific cognitive processes. For example, much attention has been focused on the role of the dorsal and ventral PFC in active maintenance and manipulation processes within working memory (WM) (Cabeza and Nyberg, 2000; D'Esposito *et al.*, 1998; Smith and Jonides, 1999). More recently, there have also been studies focused on the most anterior lateral regions of PFC, also termed frontopolar PFC (FP-PFC). In particular, it has been found that the FP-PFC is often engaged during episodic retrieval from long term memory (Buckner and Koutstaal, 1998; Tulving *et al.*, 1994). Subsequent research has been directed towards understanding the specific nature of this involvement (Konishi *et al.*, 2000; Lepage *et al.*, 2000; McDermott *et al.*, 2000; Rugg and Wilding, 2000). FP-PFC activity also appears to be reliably elicited during planning, problem solving and reasoning tasks. For example, Baker *et al.* (1996) observed FP-PFC activity in the Tower of London paradigm selectively under conditions that involved extensive planning. A recent review by Christoff and Gabrieli (2000) catalogued a number of similar studies using different task paradigms (such as the Wisconsin Card Sort Task and Raven's Progressive Matrices) that also have resulted in FP-PFC activity.

These findings suggest that FP-PFC regions are reliably activated across a number of different domains within higher cognition. However, the specific nature

¹ To whom correspondence and reprint requests should be addressed. Fax: (314) 935-4711. E-mail: tbraver@artsci.wustl.edu.

of this activation is still not well understood. Nevertheless, there have been some recent suggestions regarding the nature of FP-PFC function. Koechlin and colleagues (1999) recently proposed the term "cognitive branching" to describe the involvement of FP-PFC in situations requiring the maintenance of primary task goals while simultaneously allocating attention to subgoals. A subgoal is defined as a task that must be completed first before a higher-order goal is satisfied. Koechlin *et al.* (1999) observed selective FP-PFC activity during a fairly complex task that required integrating WM with attentional reallocation. In a recent review of the literature, Christoff and Gabrieli (2000) suggest that FP-PFC, and specifically lateral Brodmann's Area (BA) 10, may subserve the monitoring of internally (versus externally) generated information. They propose a hierarchical system in which dorsolateral PFC (DL-PFC) is involved when externally generated information is evaluated, and FP-PFC is additionally recruited when internally generated information must be evaluated, such as in reasoning tasks or in tasks requiring episodic memory. This view seems related to the one proposed by Koechlin *et al.* (1999), in that monitoring of internally generated information seems to require both maintenance in WM (i.e., creating a "mental scratchpad" to hold the information produced by internal generation during the monitoring process) and a subgoal (i.e., the monitoring operation itself).

The goal of the current study was to try to extend this previous work, by attempting to provide a clearer specification of the conditions needed to elicit activity in FP-PFC. In particular, we put forward a hypothesis similar to that espoused by Koechlin *et al.* (1999) and Christoff and Gabrieli (2000), but aimed to more specifically define the conditions under which FP-PFC activity is elicited. We suggest that a necessary minimal condition for FP-PFC activity is the requirement to perform a subgoal task while simultaneously maintaining primary goal-related information in WM. Our main goal was to test this hypothesis using a simple, and highly interpretable task design, in which each of the individual task components (i.e., active maintenance and subgoal task) are both clearly understood and well-validated from previous research. The experimental logic was to test whether each of these task subcomponents alone (WM or subgoal task) were sufficient to elicit FP-PFC activity or whether FP-PFC activity would only be present when these two components were combined. A secondary goal of the study was to further specify the nature of FP-PFC activity by testing two alternative hypotheses regarding the exact role of this region in subgoal processing during WM.

The first hypothesis is that FP-PFC activation may reflect a specialized representational code that is used to actively maintain information in tasks that require

subgoal processing. Previous research indicates that DL-PFC appears to be generally involved in tasks requiring active maintenance of goal-related or context information. However, during tasks that involve active maintenance and subgoal processing, FP-PFC may be additionally recruited to provide storage of this information in a form that is more protected from interference (due to the intervening subgoal task). A second hypothesis is that FP-PFC may be critically involved in the actual integration of the results of subgoal processing with the information that had been actively maintained prior to the subgoal task. Thus, under this second hypothesis, it is the need to integrate the two sources of information (primary task and subgoal task results) that elicits FP-PFC activity.

In order to test these hypotheses, we designed a simple paradigm based upon a WM task that we have studied extensively in previous neuroimaging and behavioral research. This task, known as the AX-CPT, is a variant of the delayed response paradigm, and requires maintenance of a cue over a delay period in order to determine the appropriate response to a subsequent probe (Barch *et al.*, 1997, 2001; Braver *et al.*, 2001; Braver and Cohen, 2001; Cohen *et al.*, 1999; Servan-Schreiber *et al.*, 1996). The current study added a subgoal component to the basic task paradigm by making stimulus categorization dependent upon elaborative semantic processing. We then compared this condition to two other closely matched conditions, one which required active maintenance in WM but had no subgoal component, and another which required performing the subgoal task (elaborative semantic processing) but had no WM component. Based on prior neuroimaging studies with the AX-CPT paradigm, we expected that without a subgoal component the WM demands of the AX-CPT would elicit activity in DL-PFC but not in FP-PFC. Likewise, tasks that require elaborative semantic processing in the absence of an explicit WM load appear to engage left ventrolateral PFC (VL-PFC) regions, but not DL-PFC or FP-PFC (Demb *et al.*, 1995; Poldrack *et al.*, 1999). Thus, we expected that only the subgoal WM condition would result in the presence of FP-PFC activity.

The task design also included a manipulation of the delay between cue and probe, while the total trial duration was held constant. Such a manipulation was expected to increase the demands on active maintenance processes (since the cue information must be maintained for a longer period), and lead to increased activity in brain regions that support active maintenance. In our previous research, we have found increased activity in DL-PFC regions (BA 46/9) resulting from such manipulations (Barch *et al.*, 1997, 2001; Braver and Cohen, 2001). Moreover, the delay manipulation is a selective one because all other components of the task remain the same. In the current study, the

delay manipulation was used to test whether FP-PFC was also specifically involved in active maintenance (under subgoal conditions). If so, we would expect to observe a significant effect of delay in this region during the subgoal WM task. In contrast, the absence of a delay effect would be more consistent with the hypothesis that FP-PFC is specifically involved in serving an integration function under these conditions.

Consistent with the hierarchical nature of the experimental design, our analysis procedure was geared to identify and isolate regions selectively involved with subgoal processing during WM, as well as with its component elements—the subgoal task itself (which was semantic processing in this study) and active maintenance. We were particularly concerned with making claims regarding the selectivity of processing in specific brain regions, because such claims enable the strongest inferences regarding the functional dissociability of different cognitive processes. Consequently, we used a conjunction analysis procedure (e.g., Price and Friston, 1997), that involved the application of multiple, stringent statistical tests as a means of identifying candidate brain regions.

METHODS

Subjects. Twenty-one right-handed subjects with no evidence of neurological compromise participated in this study. Subjects were 11 males and 10 females with a mean age of 23 years (age range 18–31 years). Subjects gave informed consent per guidelines set by the Washington University Medical Center Human Studies Committee and were paid \$25 for each hour of participation.

Behavioral tasks. Subjects performed two delayed-response WM tasks and a semantic classification task. The WM tasks were two variants of the AX version of the Continuous Performance Task (AX-CPT). The original AX-CPT requires subjects to make a positive target response to the probe letter “X,” but only when it follows the cue letter “A,” and to make a nontarget response to all other conditions (Servan-Schreiber *et al.*, 1996). Accurate responses to the probe letter require maintenance of information about the cue letter. The two tasks designed for the current study, which we termed the Word AX-CPT and the Semantic AX-CPT, retained the basic format of the original task but used words instead of letters. In the Word AX-CPT task, the target was the word “LIME” when it occurred following the word “FATE.” The target for the Semantic AX-CPT was any concrete word immediately following any abstract word. Thus, the Word AX-CPT is directly analogous to the original AX-CPT in that it requires subjects to actively maintain context information. The Semantic AX-CPT, however, additionally requires subjects to semantically classify each word, which pro-

duces a subgoal processing requirement. Consequently, to perform the task appropriately, subjects must classify the probe word (the subgoal task), while simultaneously keeping the context provided by the previous cue actively maintained in WM, so that the information from both cue and probe can be integrated into a target determination. A third task condition was included, Semantic Classification, which required subjects to make abstract-concrete judgments for words. This condition allowed us to isolate the subgoal component of the Semantic AX-CPT task.

For both AX-CPT tasks, words were presented in cue-probe pairs, one word at a time. All probes were underlined to differentiate them from the cue. Target and nontarget trials appeared intermixed in a pseudo-random sequence. Target trials (FATE-LIME or ABSTRACT-CONCRETE) occurred with 70% frequency, and nontarget trials occurred with 30% frequency. The frequency of nontarget trials was evenly distributed as follows: 10% “BX” trials in which an invalid cue preceded the target; 10% “AY” trials in which a valid probe was followed by a nontarget probe; and 10% “BY” trials, in which an invalid cue was followed by a nontarget probe. The frequencies of the various trial types replicates those used in many previous studies with the AX-CPT paradigm (Braver, Cohen, and Barch, *in press*; Cohen *et al.*, 1999).

The words for all three tasks were presented centrally on a visual display, in 36-point Helvetica font. Words were taken from standardized lists of abstract and concrete nouns. All words were three to seven letters in length and consisted of one or two syllables. Each word was presented for 750 ms. Responses to stimuli were made by pressing different buttons on a hand-held response box. In the AX-CPT tasks, subjects responded with their index finger for targets and their middle finger for nontargets. In the Semantic Classification task, subjects responded with their index finger if the word was abstract, and with their middle finger if the word was concrete.

A second factor, delay, was also manipulated in the study. Specifically, in the AX-CPT tasks, we hypothesized that regions involved with active maintenance of context should be influenced by the duration over which context should be maintained (i.e., the cue-probe delay period). Thus, we had subjects perform the AX-CPT under conditions of both short and long cue-probe delays. The short delay was 1 s in duration. In 12 of the subjects, the long delay was 5 seconds; in the remaining 9 subjects it was 7.5 s. Total trial duration was equated by counterbalancing the intertrial interval (ITI) with the cue-probe delay (e.g., for the short delay version, the delay was 1 s and ITI was 5 or 7.5 s). In the Semantic Classification condition, the timing of the words was also varied to control for the delay manipulation in the AX-CPT tasks. Thus, trials were arbi-

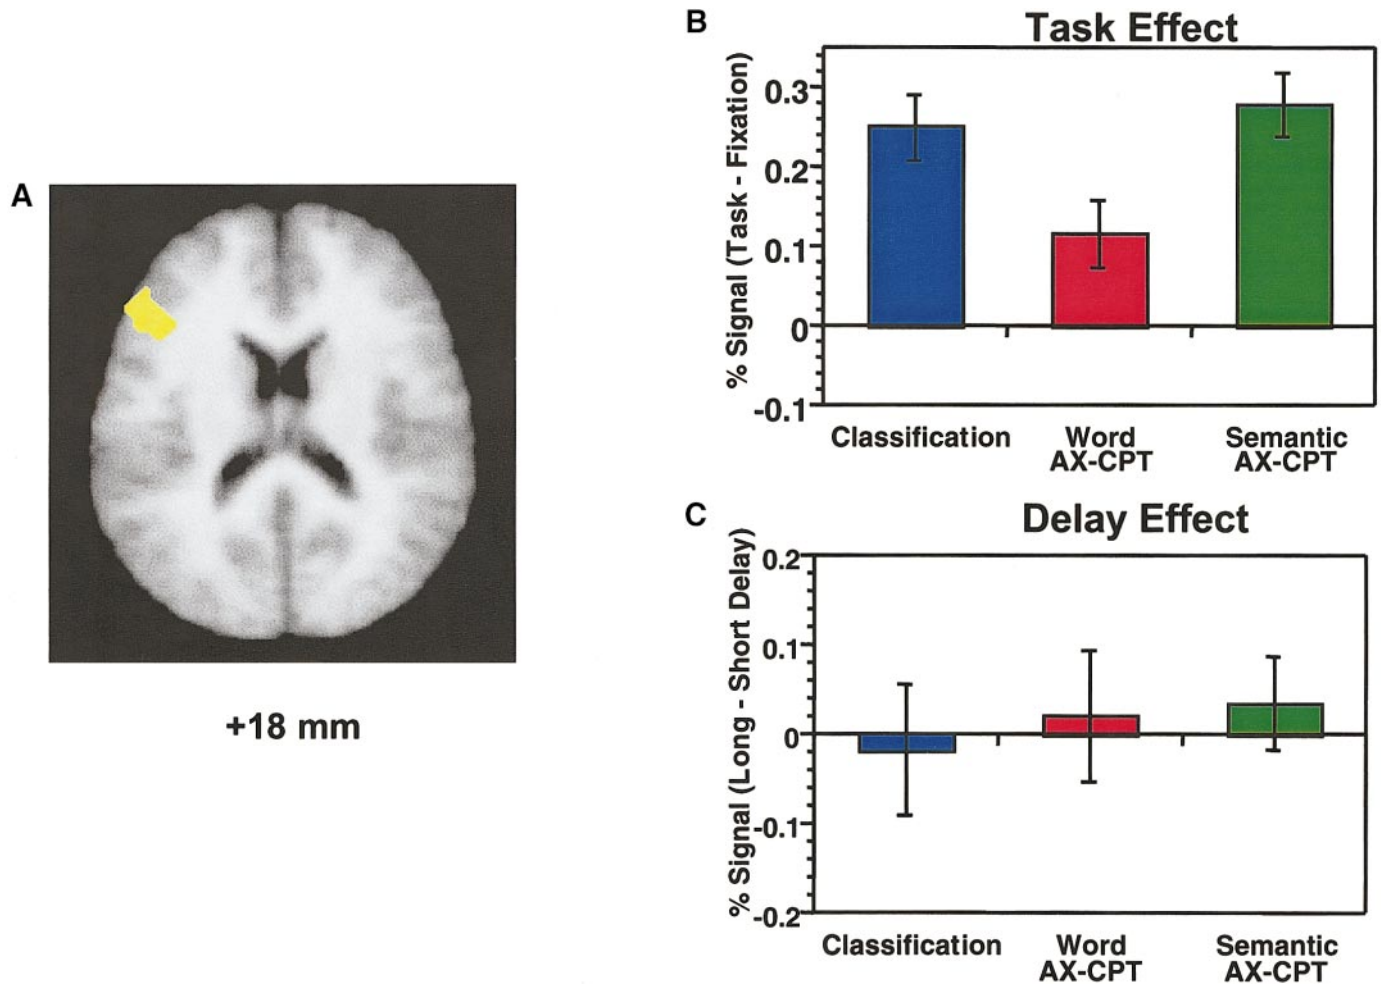


FIG. 1. Activation in a representative left VL-PFC region showing selectivity to semantic processing. In this and all subsequent figures the following conventions apply. Images are in the Talairach and Tournoux (1988) atlas space at various z coordinate locations, with activation overlaid on the corresponding anatomy image. Left of the image refers to the left side of the brain. Graph error bars represent 95% confidence interval of the mean. (A) Activation superimposed on an axial anatomic image. (B) Activation graph showing effect of task. (C) Activation graph showing effect of delay.

trarily presented as two-word pairs (but with neither word of the pair underlined), and the delay between the first and second words of the pair was either short (1 s) or long (5 or 7.5 s).

Prior to the scanning session, subjects were given instructions regarding all tasks to be performed. Subjects were then given practice trials in which to perform each task. During practice trials, the experimenter answered any further questions, validated that instructions were understood, and ensured that the tasks were performed appropriately and with a reasonably high level of accuracy.

Functional imaging. Images were acquired on a Siemens 1.5 Tesla Vision System (Erlangen, Germany) with a standard circularly polarized head coil. A pillow and tape were used to minimize head movement. Headphones dampened scanner noise and enabled

communication with participants. Both structural and functional images were acquired at each scan. High-resolution ($1.25 \times 1 \times 1$) structural images were acquired using a sagittal MP-RAGE 3-D T1-weighted sequence (TR = 9.7 mm, TE = 4, flip = 12° , TI = 300 ms) (Mugler and Brookeman, 1990). Functional images were acquired using an asymmetric spin-echo echoplanar sequence (TR = 2500, TE = 50 ms, flip = 90°). Each image consisted of 16 contiguous, 8-mm-thick axial slices acquired parallel to the anterior-posterior commissure plane (3.75×3.75 mm in-plane), allowing complete brain coverage at a high signal-to-noise ratio (Conturo *et al.*, 1996). Subjects performed two repetitions of each of the 6 task conditions (Word AX-CPT, Semantic AX-CPT and Semantic Classification at each of Short and Long delays) in separate scanning runs (12 runs total). Each run consisted of alternating cycles

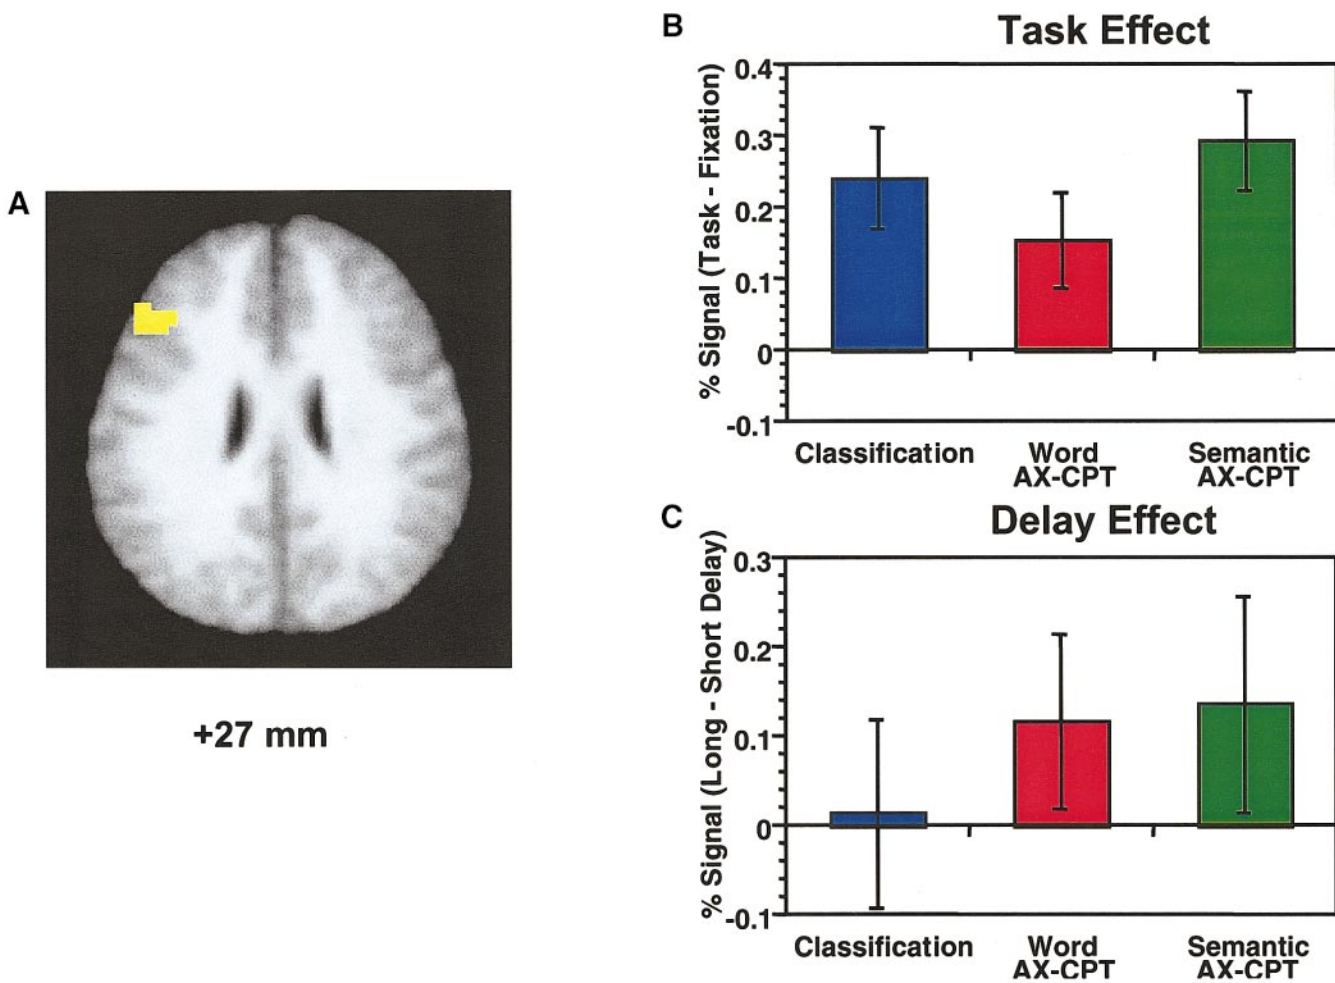


FIG. 2. Activation in a left DL-PFC activation showing selectivity to active maintenance. (A) Activation superimposed on an axial anatomic image. (B) Activation graph showing effect of task. (C) Activation graph showing effect of delay.

of task and fixation blocks. Task blocks were 10 trials in duration. Fixation blocks (denoted by a centrally presented crosshair) were either 25 or 37.5 s in duration. Finally, the first four images in each scanning run were used to allow the scanner to reach steady-state, and hence were discarded. Each run lasted approximately 5 min, and a 2-min delay occurred between runs, during which time subjects rested.

Visual stimuli were presented using PsyScope software (Cohen *et al.*, 1993) running on an Apple Power-Mac G4. Stimuli were projected to subjects with an AmPro LCD projector (model 150) onto a screen positioned at the head end of the bore. Subjects viewed the screen through a mirror attached to the head coil. A fiber-optic, light-sensitive key press interfaced with the PsyScope Button Box was used to record subjects' behavioral performance.

Data analysis. Behavioral performance data were analyzed for differential difficulty across the three task conditions by conducting ANOVAs on accuracy and RT

measures. Analyses were only conducted on probe trials because these were the only trials requiring response selection in the two AX-CPT conditions. In the Classification condition, probe trials were defined arbitrarily as the second of a two-stimulus pair.

Functional imaging data were preprocessed prior to statistical analysis according to the following procedures. All functional images were first corrected for movement using a rigid-body rotation and translation correction (Friston *et al.*, 1996; Snyder, 1996), and then registered to the subject's anatomical images (in order to correct for movement between the anatomical and function scans). The data were then scaled to achieve a whole-brain mode value (used in place of mean because of its reduced sensitivity to variation in brain margin definition) of 1000 for each scanning run (to reduce the effect of scanner drift or instability), and spatially smoothed with an 8-mm FWHM Gaussian kernel. Subjects' structural images were transformed into standardized atlas space (Talairach and Tournoux, 1988),

using a 12-dimensional affine transformation (Woods *et al.*, 1992; Woods *et al.*, 1998). The functional images were then registered to the reference brain using the alignment parameters derived for the structural scans.

Our statistical analysis procedure was designed to identify brain regions showing activation by one of three primary cognitive functions engaged by the tasks: semantic classification, active maintenance, and subgoal processing in WM. In order to make the strongest functional interpretations we further demanded that brain regions show evidence that the activation pattern was selective. Selectivity requires not only that a region is activated by a specific processing function, but also that it is significantly more activated by that function than by other comparison functions. Thus, the analysis involved a conjunction of multiple statistical tests, with each applied in a voxel-wise manner. Each statistical test was conducted with an alpha-level of $P < 0.05$, for determining statistical significance. In order for a brain region to be accepted as selective for a particular function, all voxels within the region were required to be statistically significant in all tests (described below). Because a minimum conjunction of three tests were applied for each cognitive function, the maximum voxel-wise false-positive rate was 0.0001 ($0.05 * 0.05 * 0.05$). Moreover, a region was considered significant only if it contained a cluster of 8 or more contiguous voxels. The additional cluster-size requirement ensured an overall image-wise false-positive rate of $P < 0.05$ (Forman *et al.*, 1995; McAvoy *et al.*, 2001).

The specific tests conducted were as follows. The test for selectivity to *semantic processing* involved the following conjunctions: (1) significantly increased activation (relative to fixation) in each of the two semantic conditions (Semantic Classification and Semantic AX-CPT); and (2) significantly greater activation in each of these two conditions than in the nonsemantic comparison condition (Word AX-CPT). The test for selectivity to *active maintenance* involved the following conjunctions: (1) significantly increased activation (relative to fixation) in each of the two tasks with high active maintenance demands (the long delay condition of the Semantic AX-CPT and Word AX-CPT); and (2) significantly greater activation in each of these two conditions than in the low-maintenance demand comparison conditions (the short delay condition of the Semantic AX-CPT and Word AX-CPT). Additionally, to ensure that delay-sensitive regions were not engaged because of delay effects unrelated to maintenance processes (i.e., timing, stimulus onset predictability, etc.) we further required that all voxels identified in the active maintenance conjunction test not show any activation ($P > 0.1$) in the delay contrast that did not involve maintenance (i.e., Semantic Classification long vs short delay). The test for selectivity to *subgoal processing* involved the following conjunctions: (1) signifi-

cantly increased activation (relative to fixation) in the condition requiring subgoal processing (the Semantic AX-CPT); and (2) significantly greater activation in this condition than in either of the nonsubgoal control conditions (Word AX-CPT and Semantic Classification). Additionally, for the subgoal processing test, we were concerned about false-positive identification of regions due to additive effects. In particular, a nonselective region involved in both semantic classification and active maintenance (but not necessarily subgoal processing) would be expected to show increased activity in the Semantic AX-CPT solely because both active maintenance and semantic classification are required. To guard against such false-positive identifications, we further required that voxels identified in the subgoal conjunction test not show any activation ($P > 0.1$) in the two control task contrasts (Word AX-CPT and Semantic Classification compared to fixation).

All statistical analysis were conducted using a random-effects model in which subject served as the random effect. Task-related activations were identified through a paired *t* test comparing mean signal during task blocks against mean signal during fixation blocks. For comparisons in which the delay factor was not relevant, mean signal was averaged across both short and long delay runs. To increase interpretability, only positive task-related activations were considered in all analyses. Cross-task comparisons were conducted after first subtracting mean activity during fixation blocks from mean activity during task blocks. This procedure controlled for any linear drift that may have occurred between scanning runs. Following the primary analyses, region-of-interest (ROI) analyses were conducted to quantitatively estimate the effect of experimental factors on the identified regions. For these ROI analyses (and the graphs in Figs. 1–3 displaying the results), data for the effect of task were expressed in terms of mean percentage-change in fMRI signal during task blocks relative to fixation, after first averaging across short and long delay conditions. Data for the effect of delay were expressed in terms of mean percentage-change in fMRI signal during the long delay blocks relative to the short delay blocks, after first subtracting out activation during fixation. In all graphs showing ROI effects, the bars indicating error of estimate refer to 95% confidence intervals. Thus, error bars which do not cross the zero-point on the y-axis can be considered statistically significant ($P < 0.05$).

Because of technical difficulties, behavioral data from one subject was unusable, and the scanning session had to be ended early, with only half of the scans performed (one run in each condition). We thus excluded this subject from all behavioral analyses, and imaging analyses were performed both including and excluding the subject's data. None of the reported im-

TABLE 1

	Classification	Word AX-CPT	Semantic AX-CPT
Percent errors	9.0 (2.5)	1.8 (1.6)	12.4 (5.3)
Reaction time (ms) correct trials	826.5 (50.4)	526.4 (47.7)	685.3 (65.2)
Reaction time (ms) error trials	903.8 (83.2)	642.0 (144.9)	765.8 (102.6)

Note. Data are expressed as mean values across subjects with 95% confidence intervals in parentheses.

aging results differed as a function of inclusion or exclusion.

RESULTS

Behavioral data. Accuracy and reaction time data were examined to determine whether the Semantic AX-CPT was significantly more difficult than the other two conditions. The results are summarized in Table 1. An analysis of error rates indicated a significant main effect of task ($F(2,38) = 11.52, P < 0.001$). Pair-wise comparisons indicated that performance on the Word AX-CPT task was significantly less error-prone than in either Semantic AX-CPT or Semantic Classification (Word vs Semantic AX-CPT: $t(19) = 3.78, P < 0.001$; Word vs Classification: $t(19) = 4.13, P < 0.001$). However, the Semantic AX-CPT and Classification conditions were not significantly different in error rates ($t(19) = 1.64, P > 0.1$). A similar analysis was performed on the reaction time data. Again, a highly sig-

nificant main effect of task was observed ($F(2,38) = 208.36, P < 0.001$). As with error rates, the Word AX-CPT was found to have significantly shorter RTs than Semantic AX-CPT or Classification (Word vs Semantic AX-CPT: $t(19) = 9.87, P < 0.001$; Word vs Classification: $t(19) = 26.98, P < 0.001$). Additionally, it was found that Semantic AX-CPT RTs were significantly shorter than in Classification ($t(19) = 8.66, P < 0.001$). These findings indicate that the Semantic AX-CPT was not associated with greater task difficulty, at least when compared to Classification.

However, the data also suggested a potential speed-accuracy tradeoff in the Semantic AX-CPT, since accuracy was numerically (but not significantly) worse in the Semantic AX-CPT relative to Semantic Classification, even though RTs were significantly faster. We tested for speed-accuracy tradeoffs in two ways. First, we examined RTs on correct vs error trials. If a task condition is associated with a speed-accuracy tradeoff, then one should find that error trials have faster RTs than correct trials. Yet in all conditions, error trials were *slower* than error trials (see Table 1). This effect was statistically significant for the Semantic AX-CPT (which was the only condition with enough errors to conduct a statistical test, $t(19) = 2.68, P < 0.05$). A second test for speed-accuracy tradeoff was to correlate speed against accuracy on a between-subject basis. If a speed-accuracy tradeoff were present, then there should be a significant negative correlation between speed and error rate across subjects (i.e., faster subjects should make more errors). Yet in the Semantic AX-CPT, the correlation was positive, though insignif-

TABLE 2

Brain region	Brodmann area	X	Y	Z	Size (mm ³)
Semantic processing					
Right Inferior Frontal Cortex	45	39	27	17	270
Right Frontal Operculum	47/Insula	27	20	1	540
Ventrolateral PFC	47	-42	19	-8	513
Ventrolateral PFC	44/45/47	-41	12	20	8424
Left Anterior Cingulate Cortex	32	-6	11	46	270
Left Inferotemporal Cortex	37	-46	-55	-12	1161
Active maintenance					
Left Dorsolateral PFC	46/9	-43	22	28	378
Left Thalamus	—	-14	-12	1	243
Right Inferior Parietal Cortex	40	35	-51	32	459
Right Inferior Parietal Cortex	40	39	-62	45	2025
Subgoal processing					
Right Frontopolar PFC	10	35	40	-2	432
Right Dorsolateral PFC	46/9	44	28	28	1026
Right Hippocampus	—	31	-8	-20	432
Right Superior Parietal Cortex	7	17	-72	40	243

Note. Regions are listed according to location of activation centroid. Coordinates are determined from the Talairach and Tournoux (1988) atlas.

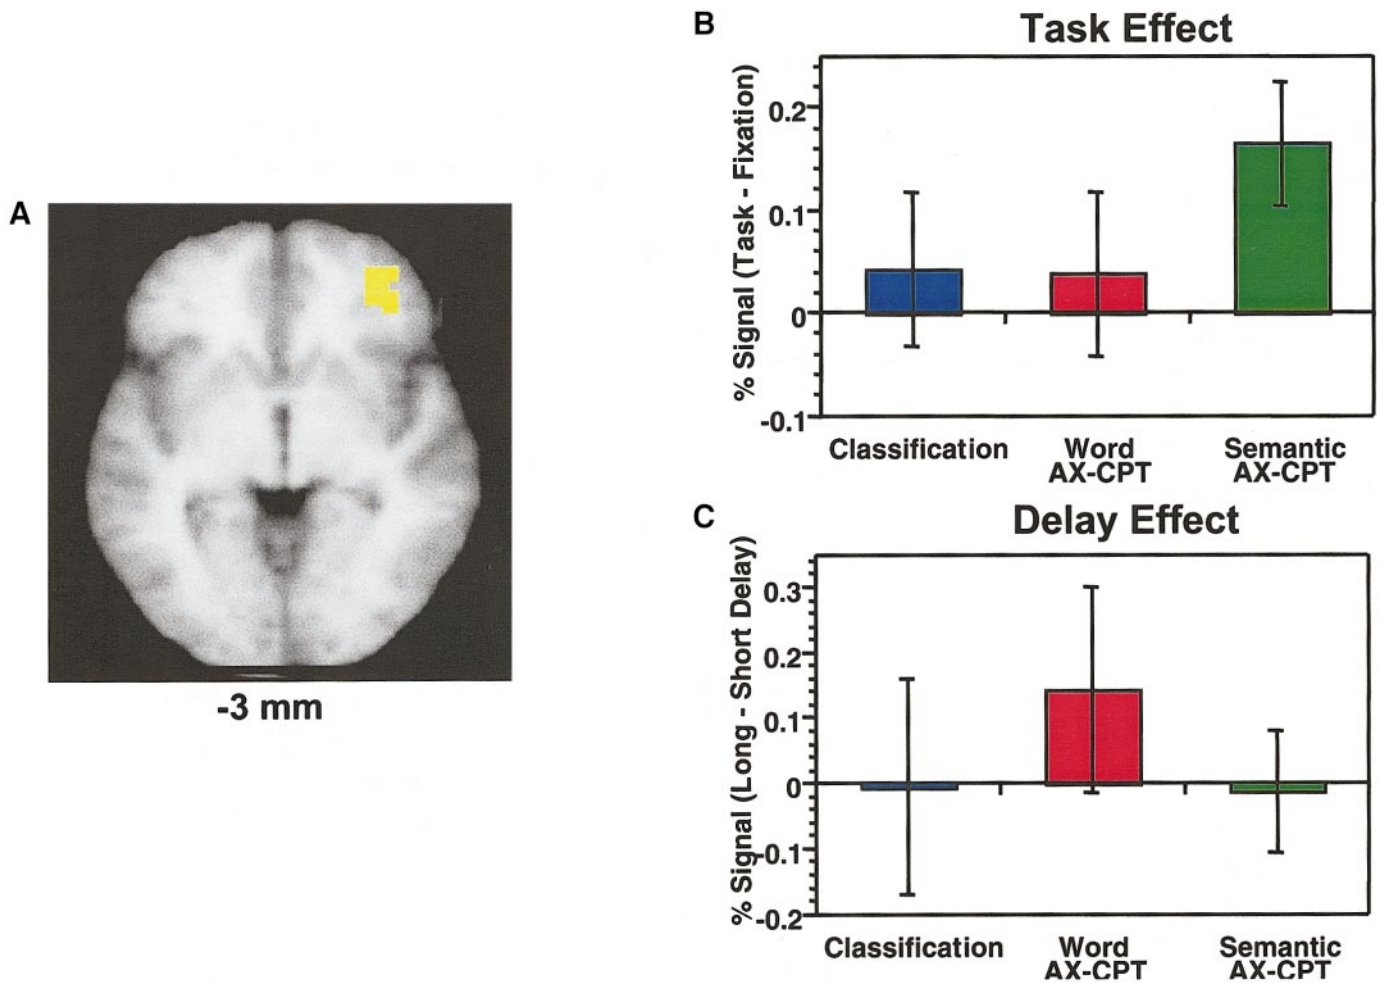


FIG. 3. Activation in a right FP-PFC region showing selectivity to subgoal processing during WM. (A) Activation superimposed on an axial anatomic image. (B) Activation graph showing effect of task. (C) Activation graph showing effect of delay.

icant ($r(18) = 0.06$). To fully rule-out the speed-accuracy relationship between the Semantic AX-CPT and Semantic Classification conditions, we also did the same correlation but on the differences in accuracy and RT across the two measures (Semantic AX-CPT minus Semantic Classification). This correlation was negative, but not even close to significant ($r(18) = -0.08$). Thus, there was no evidence for speed-accuracy tradeoffs affecting performance in the Semantic AX-CPT condition.

Imaging data. Three different analyses of the whole-brain imaging data were conducted, one testing for selectivity to semantic processing, one testing for selectivity to active maintenance, and the other testing for selectivity to subgoal processing effects. The results of these whole-brain analyses are summarized in Table 2. In the sections that follow we focus exclusively on PFC activations, since these were of primary theoretical interest.

Semantic processing. This analysis identified a primarily left-lateralized pattern of activity within VL-

PFC. Activation was observed in several large foci distributed across BA 44, 45, and 47 (see Fig. 1A and Table 2). ROI analyses of these regions indicated a high level of activity during both conditions involving semantic processing (i.e., Classification and Semantic AX-CPT), but reduced activation in the condition that did not require semantic processing (i.e., Word AX-CPT; see Fig. 1B). Moreover, in none of these regions did activity differentiate between the two semantic processing conditions, suggesting that VL-PFC regions are not sensitive to the additional processing requirements associated with the Semantic AX-CPT. There were also no effects of delay in any of these PFC regions, consistent with a selective role in semantic processing rather than active maintenance (Fig. 1C).

Active maintenance. This analysis revealed activation in left DL-PFC, within BA 46/9 (see Fig. 2A, Table 2). The location of this activity replicates the results of our previous studies (Barch *et al.*, 1997, 2001; Braver and Cohen, 2001). The effect of delay was significant in




FIG. 4. Montage showing FP-PFC regions associated with subgoal processing identified in exploratory analysis.

both AX-CPT conditions (as defined by the analysis), but moreover was equivalent in magnitude in both (Fig. 2C). However, the overall level of task-related activity was greatest in the Semantic AX-CPT condition (Fig. 2B). The ROI analysis indicated that the left DL-PFC region was also activated during Semantic Classification, although the response was not affected by delay (as defined by the analysis procedure). The finding of activity in left DL-PFC during semantic processing in the absence of maintenance demands helps explain the increased level of activity during the Semantic AX-CPT. Specifically, it suggests that the Semantic AX-CPT activity was an additive effect. Since the left DL-PFC region is engaged both by active maintenance and semantic processing, when a task involves both functions (e.g., the Semantic AX-CPT long delay condition), the activation level should be the sum of activation associated with nonsemantic active maintenance (e.g., Word AX-CPT delay effect) and with non-maintenance semantic processing (e.g., Semantic Classification). Such a hypothesis fits well with the observed pattern of left DL-PFC activity in the Semantic AX-CPT task.

Subgoal processing. Our primary analysis of interest was to identify regions activated during the Semantic AX-CPT condition, but not the other two tasks. Such a pattern of activity would suggest that the region was

selectively engaged by conditions requiring the conjunction of active maintenance and subgoal processing, rather than by individual components themselves. This analysis identified two regions in PFC showing subgoal-selective activity. Both regions were in the right hemisphere, one in DL-PFC (BA 46/9), and most importantly, the other in FP-PFC (BA10; see Fig. 3A). As defined by the analysis procedure, activation in this area during the Word AX-CPT and Semantic Classification tasks was not significantly different from fixation. Moreover, the level of activation in these two control tasks was equivalent (see Fig. 3B). This pattern argues against a difficulty-related explanation of this activity, since the behavioral data indicated that Semantic Classification task was more difficult than the Word AX-CPT. A further examination of the delay effect provided a test of whether maintenance demands influenced activity in the FP-PFC region. There was no effect of delay observed in FP-PFC activity during the Semantic AX-CPT ($F < 1$; Fig. 3C). Interestingly, there was a trend for a delay effect in the Word AX-CPT condition, although this trend did not reach significance.

We were concerned that our results regarding FP-PFC activity might be biased by the stringent criteria employed in our analysis procedure. In particular, the region of FP-PFC identified in our analysis was fairly

small in size. It is possible that this small region is not truly representative of the pattern of activity found within the larger expanse of FP-PFC. To examine this hypothesis, we conducted a more exploratory analysis, in which we relaxed the identification criteria. Specifically, we identified voxels showing significant activity in the Semantic AX-CPT and numerically greater (but not necessarily significantly greater) activity levels in that task than in either the Word AX-CPT or Semantic Classification comparison conditions. No other constraints were placed on activation levels on the two comparison conditions. When we employed this procedure, we identified a large FP-PFC region in the right hemisphere that included the original ROI, but was much greater in volume (175 contiguous voxels vs. 16 contiguous voxels), and which extended more anteriorly and superiorly (see Fig. 4). An analysis of the activation pattern in this ROI indicated that the pattern was highly similar to the original region, with significantly greater activity in the Semantic AX-CPT than in either of the other two conditions. The similarity of the larger and smaller right FP-PFC ROIs was formally tested by an ANOVA including ROI (small, large) as a factor of interest when comparing all task conditions. The ANOVA revealed no significant main effect of ROI or ROI \times condition interaction ($P > 0.1$). The exploratory analysis also identified a somewhat smaller (58 contiguous voxels) left hemisphere FP-PFC region located homotopically to the right hemisphere ROI. The left FP-PFC activation pattern was also similar to the original ROI, albeit with a less strong task effect (i.e., the activation differences between the Semantic AX-CPT and the comparison tasks were not significant). Nevertheless, an ROI-based ANOVA comparing the left and right regions showed no significant main effects or interactions involving ROI ($P > 0.1$). Thus, the exploratory analyses suggest that the FP-PFC region identified in the original analysis is in fact representative of the general pattern of right FP-PFC activity, and also to a lesser extent of left FP-PFC activity.

Finally, to more conclusively rule out the role of active maintenance in contributing to FP-PFC during the Semantic AX-CPT, we conducted a second analysis of the delay effect using the two large exploratory FP-PFC ROIs. We tested whether there were any voxels within these FP-PFC ROIs showing even a trend level ($P < 0.1$ uncorrected) delay effect in the Semantic AX-CPT task. We failed to identify any voxels that showed such a pattern.

DISCUSSION

The goal of the current study was to look at potential functional dissociations between different subregions of PFC. In particular, we were interested in determin-

ing the task conditions that elicit activity in FP-PFC as compared to the conditions that engage DL-PFC and VL-PFC regions. To examine this question, we utilized cognitive tasks that when performed in isolation reliably activate either VL-PFC regions (semantic classification) or DL-PFC regions (working memory), but do not engage FP-PFC. We then employed an experimental design in which these tasks were combined in a particular fashion that we hypothesized would lead to frontopolar activity. Specifically, the critical experimental condition required performing semantic classification as a subgoal task while concurrently maintaining information in WM, and then using the results of both sources of information (WM contents and classification judgment) to determine the appropriate response. The results confirmed our hypotheses while being consistent with prior research. Replicating prior findings, semantic processing engaged primarily left-lateralized VL-PFC regions (Demb *et al.*, 1995; Poldrack *et al.*, 1999), while the active maintenance of information in WM engaged left DLPFC (Barch *et al.*, 1997, 2001; Braver and Cohen, 2001). Moreover, performance of these cognitive functions by themselves was not associated with FP-PFC activity. However, in the critical experimental condition, FP-PFC was found to be significantly activated. These results suggest a specialized role for FP-PFC in enabling the processing of subgoal tasks during WM.

Our findings confirm and extend the recent work of Koechlin and colleagues (1999), who also observed FP-PFC activity related to subgoal processing in WM. The interpretation of Koechlin *et al.* was that FP-PFC is engaged under conditions that involve "cognitive branching." However, the particular functions subserved by FP-PFC during cognitive branching were not clearly specified in the Koechlin *et al.* study. For example, the task design utilized by Koechlin *et al.* was both novel and complex; consequently, it was not clear whether the phenomenon of branching would generalize across different task domains. Moreover, from the results presented by Koechlin *et al.* it was not clear whether FP-PFC is engaged by the unique maintenance requirements associated with cognitive branching, or by the demands associated with resuming a primary task after completing a subgoal task. In the current study, we observed FP-PFC activity that was directly related to the presence of a subgoal processing requirement during performance of a simple WM task. This finding demonstrates that the subgoal processing effect, or "cognitive branching" in the terminology of Koechlin *et al.*, is a phenomenon that generalizes across task domains. More importantly, our results also further specify the nature of FP-PFC involvement in subgoal processing. Specifically, we observed that FP-PFC activity was not increased in relationship to increased maintenance demands occurring under sub-

goal conditions (i.e., the null delay effect). This finding suggests that in our task, FP-PFC was not responsible for maintaining cue information over the delay period.

The lack of delay effect in FP-PFC favors alternative hypotheses of FP-PFC function that relate activity in this region specifically to the integration of WM and subgoal task processing. One such hypothesis is that FP-PFC is engaged to resolve potential interference occurring specifically during the period in which the subgoal task is being completed. In our task, this would correspond to the period immediately following probe presentation, when the probe word must be semantically classified, while at the same time, active maintenance of the cue information must continue. Under this hypothesis, it is the simultaneous demands on WM and subgoal processing that elicits FP-PFC activity, especially under conditions when these two demands involve the same neural substrates for processing. Moreover, under these conditions, it is the duration of subgoal processing rather than the duration of active maintenance in WM that influences FP-PFC activity. A source of support for this hypothesis is the finding in the current study that the left DL-PFC appears to be activated both by semantic processing requirements and by active maintenance of the cue. When these two tasks were combined in the Semantic AX-CPT condition, FP-PFC may have been recruited to help reduce the interference that would otherwise occur in DL-PFC.

A second hypothesis is that FP-PFC is primarily engaged under conditions in which the results of subgoal processing need to be integrated with the information stored in WM, in order to determine appropriate responding. This hypothesis suggests that the primary factor that influences FP-PFC is not subgoal processing during WM *per se*, but rather the requirement to integrate these two sources of information. Thus, a task that requires the integration of subgoal processing with WM contents would be predicted to elicit greater activation than an equivalent task in which the two sources of information (i.e., WM contents and subgoal results) remain unintegrated and unrelated. A third hypothesis is that FP-PFC activity reflects more of a “state” effect—a sustained representation of higher-level task goals across a series of trials in task conditions requiring goal-subgoal coordination and integration. Thus, the lack of the delay effect may not be indicative of an absence of maintenance-related activity in FP-PFC, but rather that the maintenance demands (of the appropriate task-set) are equivalent in both long and short delay Semantic AX-CPT conditions. All of these hypotheses are consistent with the current results, and accordingly, warrant more direct investigation in future studies.

One conclusion to be drawn from the current results is that FP-PFC may play a very specific functional role

in cognition that nevertheless appears across a wide variety of behavioral domains. The coordination and management of subgoal tasks is a computation that is present in many behaviors, such as planning, problem solving, reasoning, and other complex cognitive activities. For example, the classic Tower of Hanoi problem is one that requires subjects to decompose a higher-order goal (moving the disks to a specific peg) into a hierarchical sequence of subgoals (e.g., to move the second disk, first move the top disk to a spare peg). Each of these subgoals must be appropriately chained together, while at the same time keeping them integrated with the higher-order goal. Our findings begin to suggest the specific cognitive operations that require the engagement of FP-PFC during performance of these types of behaviors. Indeed, the findings of the current study are consistent with our previous computational analyses, in which we have suggested that anterior PFC regions represent more abstract information, such as higher-level goals, and will be selectively engaged under conditions when this information must be integrated with information being maintained and updated in more posterior PFC regions (O'Reilly *et al.*, in press).

Interestingly, another domain that may require the management and integration of subgoals is intentional retrieval of information from episodic memory. For example, episodic retrieval involves carrying out the subgoal task of searching LTM based upon a retrieval cue while actively maintaining the overall task goal (the episodic context to be recovered), and then comparing the outputs from LTM with that task goal to determine whether the information actually satisfies the goal. This conceptualization of episodic retrieval in terms of monitoring subgoal processes in WM has important implications for studies of FP-PFC function (Christoff and Gabrieli, 2000; Koechlin *et al.*, 1999). Specifically, hypotheses regarding FP-PFC function might be better characterized in terms of general computational operations related to subgoal processing, rather than those specifically associated with the cognitive domain of episodic memory. An important test of this reconceptualization will be to directly compare whether the FP-PFC regions engaged under conditions of episodic retrieval are also engaged under situations that more clearly involve subgoal processing in WM, but which do not seem to involve episodic retrieval. With regard to this issue, it is worth noting that FP-PFC activation in episodic retrieval studies has most commonly been right-lateralized, although left hemisphere activity has also been observed (Cabeza and Nyberg, 2000). A similar pattern was observed in the current study, with the primary analysis revealing only right-hemisphere FP-PFC activation, but with weaker left-hemisphere activity detected in the exploratory analysis. A recent meta-analysis of retrieval studies has suggested that

the centroid of right hemisphere FP-PFC activity associated with episodic retrieval is 30, 46, 8 (Lepage *et al.*, 2000). Although this location is more anterior and superior than the region observed in our primary analysis, it is highly consistent with the larger ROI detected in the exploratory analysis (centroid = -34, 45, 6). Thus, based on a comparison of activation, it seems that the activity found in the current study related to subgoal processing is consistent in anatomical location with that observed in retrieval studies. Nevertheless, it will be important to conduct direct within-subjects comparisons to determine whether or not the retrieval and subgoal processing areas are truly anatomically overlapping.

In addition to relating our findings anatomically to the neuroimaging literature on episodic retrieval, it also important to consider whether the current results can be alternatively interpreted as reflecting retrieval-related rather subgoal processing. In particular, a retrieval interpretation would suggest that subjects engage episodic retrieval processes to perform the Semantic AX-CPT and use these processes more reliably in the Semantic AX-CPT than in the control conditions. One problem with this interpretation of the results is that the Semantic AX-CPT seems very different from most episodic memory tasks. Specifically, subjects need to store only one piece of information (the cue) and hold it over a short period of time (less than 10 s). Moreover, given that subjects were not explicitly instructed to use an episodic retrieval strategy for any of the tasks, the question remains as to why subjects may have covertly adopted such a strategy. Moreover, why would this strategy be preferentially adopted in the Semantic AX-CPT condition than in the Word AX-CPT condition, given that they are both so similar in task demands? It is conceivable that subjects may have spontaneously adopted a strategy of rapidly storing each cue in LTM and then retrieving it following completion of the subgoal task. It is further possible that such a retrieval strategy may have been adopted preferentially in the Semantic AX-CPT because of the high load imposed by the addition of the subgoal task. Specifically, the subgoal task may have made it more difficult for subjects to have used WM to maintain cue information, which may have led to the deployment of other processing strategies. Despite the logical plausibility of a retrieval-based account of the data, it nevertheless seems like a hypothesis that would be difficult to falsify. Independent of the presence of FP-PFC activity, it is not clear how one could ascertain whether a task involves the use of covert episodic retrieval strategies on the part of subjects.

In addition to the retrieval hypothesis, there are other possible alternative interpretations of the current results. One such interpretation is that the selectively increased FP-PFC activity in the Semantic AX-

CPT condition reflects non-specific factors associated with the increased task difficulty of this condition, rather than more specific computations associated with subgoal processing during WM. This interpretation is unlikely to be correct in that the behavioral data indicated that the Semantic AX-CPT was not significantly more difficult than the Semantic Classification control task. In contrast, the behavioral data did indicate significant differences in difficulty between the Semantic Classification and Word AX-CPT tasks, yet there was no difference in FP-PFC activity across these tasks. It is worth noting that the apparently increased difficulty of the Semantic Classification condition relative to the Semantic AX-CPT (i.e., significantly slower RTs) is fully consistent with the AX-CPT experimental design. Specifically, the manipulation of trial type frequencies that occurs in the Semantic AX-CPT (e.g., 70% target trials), leads to the ability to partially predict the semantic category of the probe word based on the category of the cue. For example, following an abstract cue, subjects can predict that the probe word will likely be concrete (given that this will occur 87.5% of the time). This increased predictability can serve to prime the appropriate response prior to probe onset, thus leading to a significant facilitation in RT (and a decrease in difficulty if such contextual information is actively maintained and utilized). In the Classification condition, no equivalent facilitation occurred because there was no predictive relationship between the first and second items in a pair.

Another possible interpretation of the current results is that the Semantic AX-CPT condition may represent a dual-task situation, whereas the Classification condition and Word AX-CPT involve only single-task processing. Specifically, the Semantic AX-CPT requires the processes associated with both the Word AX-CPT and Classification conditions to be completed concurrently. Under a dual-task hypothesis, it is this requirement to complete two concurrent tasks that elicits activity in FP-PFC during the Semantic AX-CPT. Although it is true that the Semantic AX-CPT does contain a dual-task component, it does not seem likely that this component, per se, is the cause of FP-PFC activity. Previous neuroimaging studies investigating processes associated with dual-task performance have not identified FP-PFC as a brain region that is critically involved in these situations (Bunge *et al.*, 2000; D'Esposito *et al.*, 1995; Klingberg, 1998). In contrast, DL-PFC regions are most likely to be engaged under standard dual-task conditions. More importantly, Koechlin *et al.* (1999) directly compared dual-task with subgoal processing conditions and found significantly greater FP-PFC activity in the latter condition. Thus, it appears that subgoal processing in WM appears to be an important special case of a dual-task condition that preferentially engages functions

suberved by FP-PFC regions. Nevertheless, it could be important to more completely test this hypothesis in future work.

Thus, to summarize, the results of this study are consistent with the hypothesis that FP-PFC subserves cognitive functions related to the coordination, monitoring, and integration of subgoal processes within WM. Using a simple and highly decomposable experimental paradigm, we observed that neither the subgoal task itself nor the WM demand were sufficient to engage FP-PFC, but when the subgoal task was required in conjunction with WM, FP-PFC was reliably activated. However, manipulation of the duration of WM maintenance had no effect on FP-PFC activity, suggesting that this region was not specifically subserving maintenance functions within the task. In contrast, DL-PFC regions were affected by manipulations of active maintenance requirements, but showed no interaction with the presence of a subgoal task. Finally, we observed that VL-PFC regions were sensitive to the subgoal task of semantic classification, but were not further modulated by the addition of WM demands. Taken together, these findings represent a triple dissociation of function within ventrolateral, dorsolateral, and frontopolar PFC regions. We suggest that the hypothesized function of FP-PFC in subgoal processing during WM represents an important fundamental cognitive computation that may be present across a number of domains, including episodic memory. Such a reconceptualization of FP-PFC function may help to drive further research within cognitive neuroscience regarding the central mechanisms underlying higher cognition.

ACKNOWLEDGMENTS

The authors thank Deanna Barch and Kathleen McDermott for their thoughtful comments and helpful suggestions, the members of the Cognitive Control and Psychopathology Lab for productive discussions regarding these issues, Tom Conturo, Abraham Synder, and Erbil Akbudak for their technical assistance, and Marcus Raichle for his support in helping make this work possible. This study was supported by a grant from the National Institute of Mental Health (RO3 MH61615) to the first author.

REFERENCES

- Baker, S. C., Rogers, R. D., Owen, A. M., Frith, C. D., Dolan, R. J., Frackowiak, R. S. J., and Robbins, T. W. 1996. Neural systems engaged by planning: A PET study of the Tower of London Task. *Neuropsychologia* **34**: 515–526.
- Barch, D. M., Braver, T. S., Nystrom, L. E., Forman, S. D., Noll, D. C., and Cohen, J. D. 1997. Dissociating working memory from task difficulty in human prefrontal cortex. *Neuropsychologia* **35**(10): 1373–1380.
- Barch, D. M., Carter, C. S., Braver, T. S., McDonald, A., Sabb, F. W., Noll, D. C., and Cohen, J. D. 2001. Selective deficits in prefrontal cortex regions in medication naive schizophrenia patients. *Arch. Gen. Psychiatry* **50**: 280–288.
- Braver, T. S., Barch, D. M., Keys, B. A., Carter, C. S., Cohen, J. D., Kaye, J. A., Janowsky, J. S., Taylor, S. F., Yesavage, J. A., Mummenthaler, M. S., Jagust, W. J., and Reed, B. R. Context processing in older adults: Evidence for a theory relating cognitive control to neurobiology in healthy aging. *J. Exp. Psychol. Gen.* **130**: 746–763.
- Braver, T. S., and Cohen, J. D. 2001. Working memory, cognitive control, and the prefrontal cortex: Computational and empirical studies. *Cogn. Process.* **2**: 25–55.
- Braver, T. S., Cohen, J. D., and Barch, D. M. The role of the prefrontal cortex in normal and disordered cognitive control: A cognitive neuroscience perspective. In D. T. Stuss and R. T. Knight (Eds.). *Principles of Frontal Lobe Function*. Oxford University Press, Cambridge.
- Buckner, R. L., and Koutstaal, W. 1998. Functional neuroimaging studies of encoding, priming and explicit memory retrieval. *Proc. Natl. Acad. Sci. USA* **95**: 891–898.
- Bunge, S. A., Klingberg, T., Jacobsen, R. B., and Gabrieli, J. D. E. 2000. A resource model of the neural basis of executive working memory. *Proc. Natl. Acad. Sci. USA* **97**(7):3573–3578.
- Cabeza, R., and Nyberg, L. 2000. Imaging cognition II: An empirical review of 275 PET and fMRI studies. *J. Cogn. Neurosci.* **12**: 1–47.
- Christoff, K., and Gabrieli, J. D. E. 2000. The frontopolar cortex and human cognition: Evidence for a rostrocaudal hierarchical organization within the human prefrontal cortex. *Psychobiology* **28**: 168–186.
- Cohen, J. D., Barch, D. M., Carter, C., and Servan-Schreiber, D. 1999. Context-processing deficits in schizophrenia: Converging evidence from three theoretically motivated cognitive tasks. *J. Abnormal Psychol.* **108**: 120–133.
- Cohen, J. D., MacWhinney, B., Flatt, M. R., and Provost, J. 1993. PsyScope: A new graphic interactive environment for designing psychology experiments. *Behav. Res. Methods Instr. Comput.* **25**(2):257–271.
- Conturo, T. E., McKinstry, R. C., Akbudak, E., Snyder, A. Z., Yang, T., and Raichle, M. E. 1996. *Sensitivity Optimization and Experimental Design in Functional Magnetic Resonance Imaging*. Paper presented at the Society for Neuroscience, Washington D.C.
- D'Esposito, M., Aguirre, G. K., Zarahn, E., Ballard, D., Shin, R. K., and Lease, J. 1998. Functional MRI studies of spatial and nonspatial working memory. *Cogn. Brain Res.* **7**: 1–13.
- D'Esposito, M., Detre, J. A., Alsop, D. C., Shin, R. K., Atlas, S., and Grossman, M. 1995. The neural basis of the central executive system of working memory. *Nature* **378**(November 16): 279–281.
- Demb, J. B., Desmond, J. E., Wagner, A. D., Vaidya, C. J., Glover, G. H., and Gabrieli, J. D. E. 1995. Semantic encoding and retrieval in the left inferior prefrontal cortex: A functional MRI study of task difficulty and process specificity. *J. Neurosci.* **15**: 5870–5878.
- Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., and Noll, D. C. 1995. Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): Use of a cluster-size threshold. *Magn. Reson. Med.* **33**: 636–647.
- Friston, K. J., Williams, S., Howard, R., Frackowiak, R. S. J., and Turner, R. 1996. Movement-related effects in fMRI time-series. *Magn. Reson. Med.* **35**: 346–355.
- Fuster, J. M. 1989. A theory of prefrontal functions: The prefrontal cortex and the temporal organization of behavior, *The Prefrontal Cortex: Anatomy, Physiology and Neuropsychology of the Frontal Lobe*, pp. 157–192. Raven Press, New York.
- Goldman-Rakic, P. S. 1987. Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In *Handbook of Physiology—The Nervous System V* (F. Plum and V. Mountcastle, Eds.), Vol. 5, pp. 373–417. Am. Physiol. Soc., Bethesda, MD.

- Klingberg, T. 1998. Concurrent performance of two working memory tasks: Potential mechanisms of interference. *Cerebral Cortex* **8**: 593–601.
- Koechlin, E., Basso, G., Pietrini, P., Panzer, S., and Grafman, J. 1999. The role of the anterior prefrontal cortex in human cognition. *Nature* **399**: 148–151.
- Konishi, S., Wheeler, M. E., Donaldson, D. I., and Buckner, R. L. 2000. Neural correlates of episodic retrieval success. *NeuroImage* **12**(3): 276–286.
- Lepage, M., Ghaffar, O., Nyberg, L., and Tulving, E. 2000. Prefrontal cortex and episodic retrieval mode. *Proc. Natl. Acad. Sci.* **97**: 506–511.
- McAvoy, M. P., Ollinger, J. M., and Buckner, R. L. 2001. Cluster size thresholds for assessment of significant activation in fMRI. *NeuroImage* **13**: S198.
- McDermott, K. B., Jones, T. C., Petersen, S. E., Lageman, S. K., and Roediger, H. L., III. 2000. Retrieval success is accompanied by enhanced activation in anterior prefrontal cortex during recognition memory: An event related MRI study. *J. Cogn. Neurosci.* **12**(6): 965–976.
- Mugler, J. P. I., and Brookeman, J. R. 1990. Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP-RAGE). *Magn. Reson. Med.* **15**: 152–157.
- O'Reilly, R. C., Noelle, D. C., Braver, T. S., and Cohen, J. D. Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. *Cerebral Cortex*, in press.
- Poldrack, R. A., Wagner, A. D., Prull, M. W., Desmond, J. E., Glover, G. H., and Gabrieli, J. D. 1999. Functional specialization for semantic and phonological processing in the left inferior prefrontal cortex. *NeuroImage* **10**: 15–35.
- Price, C. J., and Friston, K. J. 1997. Cognitive conjunction: A new approach to brain activation experiments. *Neuroimage* **5**: 261–270.
- Rugg, M. D., and Wilding, E. L. 2000. Retrieval processing and episodic memory. *Trends Cogn. Sci.* **4**: 108–115.
- Servan-Schreiber, D., Cohen, J. D., and Steingard, S. 1996. Schizophrenic deficits in the processing of context: A test of a theoretical model. *Arch. Gen. Psychiatry* **53**: 1105–1113.
- Smith, E. E., and Jonides, J. 1999. Storage and executive processes in the frontal lobes. *Science* **283**: 1657–1661.
- Snyder, A. Z. 1996. Difference image versus ratio image error function forms in PET–PET realignment. In *Quantification of Brain Function Using PET* (D. Bailer and T. Jones, Eds.) Academic Press, San Diego.
- Stuss, D. T., and Benson, D. F. 1986. *The Frontal Lobes*. Raven Press, New York.
- Talairach, J., and Tournoux, P. 1988. *Co-planar Stereotaxic Atlas of the Human Brain*. Thieme, New York.
- Tulving, E., Kapur, S., Craik, F. I., Moscovitch, M., and Houle, S. 1994. Hemispheric encoding/retrieval asymmetry in episodic memory: Positron emission tomography findings. *Proc. Natl. Acad. Sci. USA* **91**: 2016–2020.
- Woods, R. P., Cherry, S. R., and Mazziotta, J. C. 1992. Rapid automated algorithm for aligning and reslicing PET images. *J. Comput. Assist. Tomogr.* **16**: 620–633.
- Woods, R. P., Grafton, S. T., Holmes, C. J., Cherry, S. R., and Mazziotta, J. C. 1998. Automated image registration: I. general methods and intrasubject, intramodality validation. *J. Comput. Assist. Tomogr.* **22**(1): 139–152.