Neuron, Vol. 31, 329–338, August 2, 2001, Copyright ©2001 by Cell Press

Recovering Meaning: Left Prefrontal Cortex Guides Controlled Semantic Retrieval

Anthony D. Wagner,[1,2,4] E. Juliana Paré-Blagoev,[3] Jill Clark,[2] and Russell A. Poldrack[2]
[1]Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
[2]Massachusetts General Hospital
Nuclear Magnetic Resonance Center
Charlestown, Massachusetts 02129
[3]Graduate School of Education
Harvard University
Cambridge, Massachusetts 02138

Summary

Prefrontal cortex plays a central role in mnemonic control, with left inferior prefrontal cortex (LIPC) mediating control of semantic knowledge. One prominent theory posits that LIPC does not mediate semantic retrieval per se, but rather subserves the selection of task-relevant knowledge from amidst competing knowledge. The present event-related fMRI study provides evidence for an alternative hypothesis: LIPC guides *controlled semantic retrieval* irrespective of whether retrieval requires selection against competing representations. With selection demands held constant, LIPC activation increased with semantic retrieval demands and with the level of control required during retrieval. LIPC mediates a top-down bias signal that is recruited to the extent that the recovery of meaning demands controlled retrieval. Selection may reflect a specific instantiation of this mechanism.

Introduction

Complex behavior often requires controlled access to and utilization of memory. Prefrontal cortex plays a central role in cognitive control (Norman and Shallice, 1986; Shallice, 1988; Cohen et al., 1996), and is thought to be a component of the neural circuitry underlying the control of memory (Stuss and Benson, 1984; Goldman-Rakic, 1987; Shimamura, 1995; Fuster, 1997; Miller and Cohen, 2001). The specific prefrontal regions recruited in service of mnemonic control appear to depend on the content of the memory (Goldman-Rakic, 1995; Kelley et al., 1998; Wagner et al., 1998a; but see Asaad et al., 1998), and on the particular mechanisms engaged during memory processing (Petrides, 1996; Owen, 1997; D'Esposito et al., 1998; Passingham et al., 2000). One form of mnemonic control that is central to many forms of cognition is the control of long-term semantic knowledge. Over the past decade, neuroimaging evidence has pointed to the left inferior prefrontal cortex (LIPC) as central to the control of semantic memory, including the recovery and evaluation of meaning (Petersen et al., 1988; Kapur et al., 1994; Martin et al., 1995; Gabrieli et al., 1996; Vandenberghe et al., 1996; Fiez, 1997; Thompson-Schill et al., 1997; Wagner et al., 1997; Martin and Chao, 2001). Although considerable attention has been focused on understanding LIPC function, the fundamental nature of LIPC contributions to the recovery of meaning remains unresolved and controversial.

Initial observations of increased LIPC activation during performance of semantic generation and semantic decision tasks led some theorists to posit that LIPC subserves semantic retrieval or semantic working memory processes (Petersen et al., 1988; Kapur et al., 1994; Demb et al., 1995; Fiez, 1997). From this perspective, LIPC is thought to guide access to goal-relevant semantic knowledge, and to permit "working with" or evaluation of the recovered knowledge (Kapur et al., 1994; Buckner, 1996; Gabrieli et al., 1998; Wagner, 1999). Consistent with this perspective, repeated access to task-relevant knowledge is associated with decreased engagement of LIPC control processes (Raichle et al., 1994; Gabrieli et al., 1996; Wagner et al., 2000b).

A significant challenge to the semantic retrieval hypothesis, however, recently emerged from a series of elegant functional magnetic resonance imaging (fMRI) and neuropsychological studies, with the resultant data leading to a reconceptualization of the role of LIPC in knowledge recovery (Thompson-Schill et al., 1997; Thompson-Schill et al., 1998; Thompson-Schill et al., 1999). Rather than mediating semantic retrieval per se, Thompson-Schill and colleagues have posited that LIPC specifically subserves, and is necessary for, the selection of task-relevant representations from amidst competing representations. From this perspective, LIPC mechanisms are engaged when a subset of knowledge must be recovered from amidst other competing knowledge, but are not engaged and are not necessary when semantic retrieval does not require selection.

Support for the selection hypothesis was garnered using a semantic decision paradigm that putatively varied selection demands (Thompson-Schill et al., 1997). A "high selection" condition required subjects to determine which of two targets (e.g., "tongue" and "bone") was most similar to a cue (e.g., "tooth") along a single semantic dimension or feature (such as COLOR). To assess similarity under such conditions, subjects must select from semantic memory each item's color attributes from amidst other competing semantic attributes (such as SIZE, SHAPE, etc). In contrast, a "low selection" condition required subjects to assess the global similarity between the cue (e.g., "flea") and targets (e.g., "tick" and "well"). Critically, no selection was required for this task because all semantic knowledge is relevant for computing global similarity. That is, it was argued that assessment of how globally similar two items are depends on a comparison between the items along all semantic dimensions or features. Thus, because all features are task-relevant, no subset of associated semantic knowledge would be selected against in favor of other knowledge. As clearly posited by Thompson-Schill and colleagues (p. 14792), "[c]omparisons of global similarity, which are based on all available information and thus do not require selection," should not elicit LIPC

[4]Correspondence: awagner@psyche.mit.edu

activation. Consistent with the selection hypothesis, fMRI revealed greater LIPC activation during high selection compared to low selection conditions.

Importantly with respect to the semantic retrieval hypothesis, the study by Thompson-Schill and colleagues (1997) also contained two low selection conditions that, although not requiring selection, were designed to putatively differ in semantic retrieval demands. Specifically, the conditions differed in the number of possible targets, either two or four, that were to be considered for their global similarity to the cue. The rationale was that increasing the number of targets should increase semantic retrieval demands, that is, the amount of semantic knowledge that must be accessed. However, because the global similarity task does not require selection, increasing the number of targets should not affect selection demands. To the extent that LIPC mediates semantic retrieval, then activation should increase with the number of targets even though competition or selection putatively does not vary. FMRI failed to reveal such an increase. As noted by Thompson-Schill and colleagues "[t]he absence of [LIPC] activity for this comparison counters the argument that the effects of selection can be attributed solely to variations in degree of semantic retrieval" (p. 14792). Critically, this finding stands as the most significant challenge to the hypothesis that LIPC mediates semantic retrieval even when selection demands are nonexistent due to an absence of between-feature competition for recovery.

Although these and related findings have led to acceptance of the selection account by many theorists, here, we posit an alternative account of LIPC function that is equally consistent with extant findings. Specifically, we propose that LIPC contributes to *controlled* semantic retrieval. That is, LIPC mechanisms may guide the recovery of semantic knowledge under situations where preexperimental associations or prepotent responses do not support the recovery of task-relevant knowledge through more automatic mechanisms. When a strong association exists between two elements, be they two stimuli or a stimulus and a response, the presentation of the first element may yield sufficient activation of the second element such that this associated representation may be accessed relatively automatically. That is, the second element may be recovered even in the absence of top-down facilitation or bias. Importantly, considerable evidence suggests that prefrontal regions are particularly important for cognition and behavior under conditions where strong stimulus-stimulus or stimulus-response associations are absent (Norman and Shallice, 1986; Cohen et al., 1996; Miller and Cohen, 2001). This increased role of prefrontal cortex when associations are weak may reflect the greater need for top-down bias signals to guide controlled access to or retrieval of the associate when presented the first element.

Here, we posit that LIPC may subserve controlled semantic retrieval even when selection against competing knowledge is not required, although heightened competition from irrelevant knowledge also may increase demands on this top-down control mechanism. Importantly, the findings of Thompson-Schill and colleagues (1997) do not address this hypothesis because the cue-target associative strength was quite strong in their low selection conditions (e.g., the target response

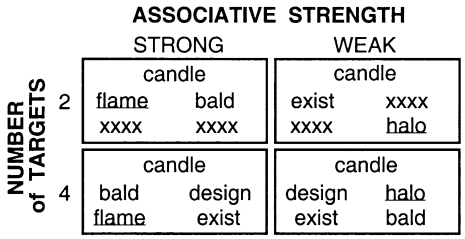

Figure 1. Example Trials Depicting the Four Experimental Conditions

On each experimental trial, a single cue word and either two or four target items were presented, and subjects indicated which target was most closely globally related to the cue by pressing a response key. The correct response (underlined for illustrative purposes) was either a strong or a weak associate of the cue. This task is analogous to the global similarity task of Thompson-Schill and colleagues (1997), and thus requires no selection as all semantic attributes are relevant to computing global relatedness.

for the cue "tick" was the strongly associated stimulus "flea"). Thus, these low selection conditions also likely constituted low controlled retrieval conditions, where performance could have been based on more automatic semantic retrieval processes not mediated by LIPC. That is, the critical data that were proffered as being inconsistent with the initial semantic retrieval hypothesis and thus were argued to support the selection account, may ultimately prove consistent with the hypothesis that LIPC guides semantic retrieval but only under situations were mnemonic control is required.

Thus, although considerable effort has been devoted to understanding the role of LIPC in the recovery of meaning, a fundamental issue that is central to understanding the nature of prefrontal contributions to mnemonic control, controversy nevertheless remains. To address this issue, the present event-related fMRI study systematically tested the selection and the controlled semantic retrieval hypotheses of LIPC function under conditions that paralleled the low selection conditions of Thompson-Schill and colleagues (1997). That is, subjects performed a semantic decision task that required access to global semantic knowledge and thus did not require selection. In this task, subjects were presented a cue and had to determine which target from a set of possible targets was most globally related to the cue (Figure 1). Controlled retrieval demands were manipulated by varying the preexperimental associative strength between the cue (e.g., "candle") and the correct target, with the target being either a strong (e.g., "flame") or weak (e.g., "halo") associate of the cue. Moreover, semantic retrieval demands were manipulated by varying the number of possible targets, either two or four, in the response set (Thompson-Schill et al., 1997). To the extent that LIPC mediates controlled semantic retrieval, varying both the associative strength and the number of targets should modulate LIPC activation: weaker associative strength and more targets should elicit greater LIPC activation. By contrast, the selection hypothesis predicts that these factors should not affect the extent of LIPC engagement (Thompson-Schill et al., 1997). Thus, this factorial design permits two direct tests of these alternative hypotheses of LIPC function.

Table 1. Accuracy and Reaction Time Associated with Task Performance

Condition	Accuracy	Reaction Time (ms)
Strong		
2-target	.95 (.01)	1704 (85.4)
4-target	.93 (.01)	1989 (86.4)
Weak		
2-target	.89 (.01)	2053 (93.8)
4-target	.79 (.02)	2513 (87.3)

Values in parentheses refer to Standard Errors.

Finally, it is worth noting that initial neuroimaging data suggest that LIPC is not functionally homogeneous. Rather, the posterior and dorsal extent of LIPC (~Brodmann's area [BA] 44) has been observed in studies of selection or interference resolution (Thompson-Schill et al., 1997; Jonides et al., 1998; Thompson-Schill et al., 1998). An anatomically similar region also has been associated with phonological control, including phonological access and maintenance (Paulesu et al., 1993; Awh et al., 1996; Fiez et al., 1996; Poldrack et al., 1999; Wagner, 1999; Davachi et al., 2001). By contrast, the anterior and ventral extent of LIPC (~BA 47/45) has been particularly associated with semantic control (Fiez, 1997; Price et al., 1997; Poldrack et al., 1999; Wagner et al., 2000a; Bokde et al., 2001). To the extent that posterior LIPC and anterior LIPC are functionally distinct, then the effects of varying controlled semantic retrieval demands were expected to be particularly robust in anterior LIPC.

Results

Behavioral Performance

Consideration of performance accuracy and response latencies revealed main effects of Associative Strength and of Number of Targets, as well as interactions (Table 1). Accuracy was lower and reaction times (RTs) were longer for Weak than for Strong associative strength trials (accuracy, $F(1,13) = 121.78$, $p < .0001$; RT, $F(1,13) = 266.00$, $p < .0001$), and for 4-target than for 2-target trials (accuracy, $F(1,13) = 28.12$, $p < .0001$; RT, $F(1,13) = 147.47$, $p < .0001$). The effects of Number of Targets were greater when Associative Strength was Weak (accuracy, $F(1,13) = 25.79$, $p < .001$; RT, $F(1,13) = 9.79$, $p < .01$). RTs did not significantly differ between the Weak/2-target and the Strong/4-target conditions (planned contrast, $F(1,13) = 2.63$, $p = .13$).

fMRI Indices of LIPC Function

An initial analysis of task performance revealed that engagement in the semantic relatedness task elicited above baseline activation in numerous neural regions, including occipital, lateral temporal, motor, and, critically, left inferior prefrontal cortices (LIPC) (Figure 2A). Of central interest is whether the response in LIPC regions was modulated by the Number of Targets and the cue-target Associative Strength; all subsequent analyses principally focus on the effects of these factors on LIPC activation, although effects in other prefrontal regions will be briefly considered.

LIPC activation during performance of each of the four conditions was initially assessed relative to baseline (Figure 2B). Analyses revealed that all four conditions elicited greater activation in the posterior and dorsal extent of the left inferior prefrontal cortex (posterior LIPC). Anatomically, this region corresponds approximately to the posterior LIPC region implicated in studies of selection or interference resolution (Thompson-Schill et al., 1997; Jonides et al., 1998; Thompson-Schill et al., 1998). Moreover, this region corresponds approximately

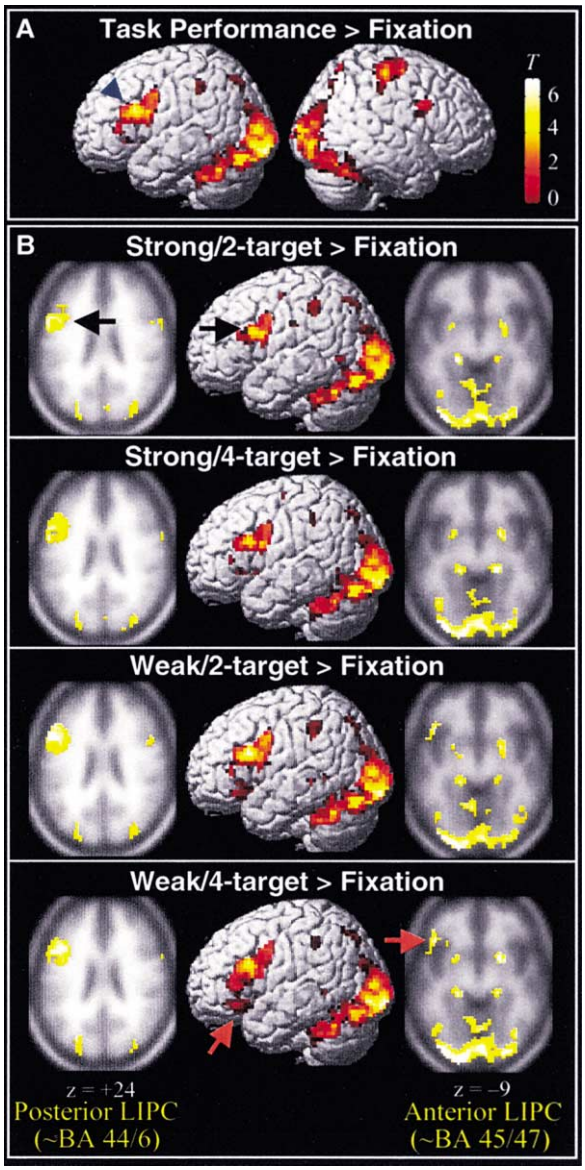

Figure 2. LIPC Regions Engaged during Performance of the Semantic Decision Task

(A) Analyses demonstrated task-related activation increments in occipital, inferior temporal, motor, and LIPC (blue arrowhead) regions relative to baseline.

(B) LIPC activation was observed during the Strong/2-target (first row), Strong/4-target (second row), Weak/2-target (third row), and Weak/4-target (fourth row) trials relative to baseline. Activation in posterior LIPC (black arrow) was elicited by all conditions, whereas activation in anterior LIPC (red arrow) was principally reliable during Weak associative strength trials.

Table 2. Prefrontal Foci Modulated by Number of Targets and Associative Strength

Contrast	Stereotaxic Coordinates			~Brodmann's Areas	Peak Z Score (No. voxels)
4-target > 2-target					
Posterior LIPC / middle frontal	−36	21	27	44, 45, 9, 46	4.40 (111)
Posterior LIPC	−39	6	24	44, 6	4.37
Anterior LIPC	−45	27	9	45, 47	4.02
Anterior cingulate	3	30	36	32, 8	3.73 (31)
Anterior cingulate	3	15	42	32, 8	3.21
Anterior LIPC	−51	21	−12	47	3.59 (15)
Weak > Strong					
Anterior LIPC	−45	27	−12	47, 45	5.33 (560)
Posterior LIPC / middle frontal	−51	18	27	44, 45, 9, 46	5.12
Anterior LIPC	−51	21	−3	47, 45	4.81
ACC/Medial superior frontal	6	18	39	32, 8	4.72 (144)
ACC	9	27	36	32, 8	4.32
Medial superior frontal	0	9	57	6	4.28
Anterior RIPC	45	21	6	45, 47	4.30 (83)
Anterior RIPC / orbital frontal	30	24	−6	47, 11	4.08
Anterior RIPC / orbital frontal	39	27	−9	47, 11	4.06
RIPC / middle frontal	54	24	27	44, 45, 9, 46	3.62 (21)
Posterior RIPC	45	9	27	44, 6, 9	3.54 (29)
Posterior RIPC	51	9	33	44, 6, 9	3.23
Weak/2-target > Strong/4-target					
RIPC / middle frontal	51	21	27	44, 45, 9, 46	4.69 (114)
Posterior RIPC	54	12	21	44, 9	4.65
Posterior RIPC / middle frontal	54	12	36	9, 8, 44	4.49
Anterior LIPC	−48	27	−12	47, 45	4.41 (50)
Anterior LIPC	−42	33	−12	47, 45	4.41
Medial superior frontal / ACC	−9	6	48	6, 32	3.86 (11)
LIPC / middle frontal	−48	30	27	45, 46	3.72 (15)
Posterior LIPC / middle frontal	−51	21	24	44, 45, 9, 46	3.63
Posterior LIPC	−57	12	30	44, 9	3.37
Medial superior frontal	6	6	54	6	3.70 (9)

ACC, anterior cingulate cortex; LIPC, left inferior prefrontal; No. voxels, number of voxels in cluster; RIPC, right inferior prefrontal.

to the posterior LIPC region previously implicated in phonological control, including phonological access and maintenance (Paulesu et al., 1993; Awh et al., 1996; Fiez et al., 1996; Poldrack et al., 1999; Wagner, 1999). By contrast, these voxel-level analyses revealed that the anterior and ventral extent of LIPC (anterior LIPC) was reliably engaged only when the cue-target associative strength was Weak (Figure 2B). Anatomically, this region corresponds approximately to the anterior LIPC region previously implicated in semantic control (Fiez, 1997; Price et al., 1997; Poldrack et al., 1999; Wagner, 1999). Collectively these results suggest that posterior LIPC and anterior LIPC were differentially sensitive to the associative strength manipulation. As detailed below, this difference was validated through region-of-interest analyses that permitted a direct test of whether these two LIPC regions demonstrated a different pattern of activation as controlled retrieval demands were varied.

In contrast to the predictions of the selection hypothesis, there was a main effect of Number of Targets on LIPC activation, with activation increasing as the number of targets increased (Table 2, Figure 3A). Number of Targets robustly affected posterior LIPC activation and yielded a modest but reliable effect in anterior LIPC (Table 2). Separate contrasts further revealed greater LIPC activation during 4-target trials than during 2-target trials both when associative strength was Weak and when it was Strong. Importantly, these latter findings reflect a direct failure to replicate the null result observed in the comparable 4-target versus 2-target contrast in the low selection conditions of Thompson-Schill et al. (1997). As detailed in the Introduction, this earlier null result constituted the primary evidence mustered against the hypothesis that LIPC activation varies with semantic retrieval demands even when selection demands are held constant.

To determine whether LIPC is sensitive to varying control demands during semantic retrieval, we next tested for the main effect of Associative Strength on LIPC activation. Critically, with respect to the controlled retrieval hypothesis, greater anterior LIPC and posterior LIPC activation was observed when cue-target associative strength was Weak relative to when it was Strong (Table 2, Figure 3B). Cortical inflation at the individual subject level further demonstrated that the effects of Associative Strength in anterior LIPC and posterior LIPC were anatomically distinct (Figure 4). Moreover, separate contrasts revealed greater LIPC activation during Weak relative to Strong trials both during 2-target trials and during 4-target trials.

A central question is whether the effect of Associative Strength differed across anterior LIPC and posterior LIPC. As detailed above, voxel-based comparisons of each condition to baseline revealed reliable activation increases in posterior LIPC during both Weak and Strong associative strength trials, whereas anterior LIPC activation was reliably above baseline only when asso-

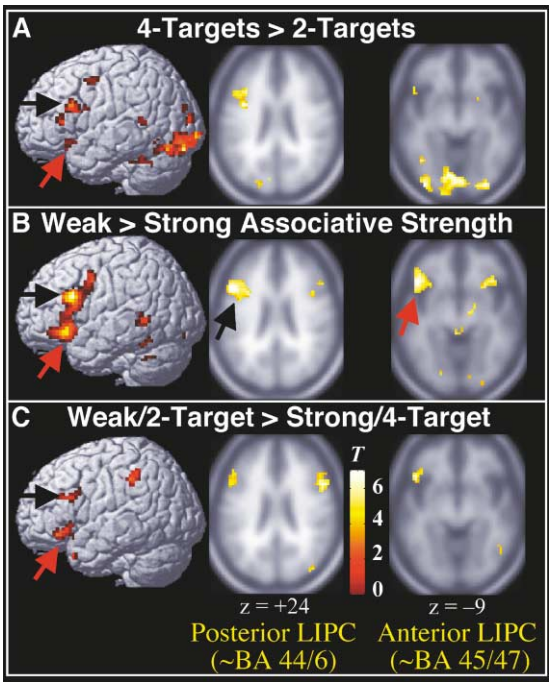

Figure 3. LIPC Activation Was Sensitive to the Number of Targets and to Associative Strength

(A) A main effect of Number of Targets was observed in posterior LIPC (black arrow) and in anterior LIPC (red arrow), with activation increasing with targets. Separate contrasts (not shown) revealed greater LIPC activation during 4-target than during 2-target trials both when Associative Strength was Weak (peak coordinates of −36, 21, 27; −45, 15, 16; −39, 9, 27; −48, 30, 0; and −45, 27, 9) and when it was Strong (coordinates of −45, 21, 24; −39, 21, 30; −39, 6, 27; and −51, 21, −12).

(B) A main effect of Associative Strength was observed in posterior LIPC and in anterior LIPC, with activation being greater when associative strength was Weak. Separate contrasts (not shown) revealed greater LIPC activation during Weak relative to Strong trials both when there were 2-targets (coordinates of −36, 6, 30; −51, 18, 27; −33, 15, 30; −48, 27, −12; and −36, 27, 3) and when there were 4-targets (coordinates of −51, 18, 24; −45, 15, 15; and −54, 27, 3).

(C) Direct comparison of Weak/2-target to Strong/4-target trials demonstrated that posterior LIPC and anterior LIPC activation was greater during the Weak/2-target trials.

ciative strength was Weak. To further assess the effects of Associative Strength, region-of-interest (ROI) analyses were conducted on the significant foci identified as demonstrating a main effect of Number of Targets (Figure 5). Because this contrast is orthogonal to that of Associative Strength, definition of the ROIs in this manner rendered them unbiased with respect to this factor. Consistent with the observed main effect of Associative Strength in the voxel-based analysis, these ROI analyses revealed that Associative Strength influenced both anterior LIPC (BA 47; coordinates of −51, 21, −12; $F(1,13) = 33.41$, $p < .0001$) and posterior LIPC (BA 44/6; coordinates of −39, 6, 24; $F(1,13) = 19.01$, $p < .001$). Importantly, these analyses further revealed a Region X Associative Strength interaction ($F(1,13) = 7.46$, $p < .02$). That is, there was a robust response in posterior LIPC during both Weak (.63) and Strong (.40) trials, but a robust response in anterior LIPC during Weak (.66) trials and a markedly more modest response in this region during Strong (.18) trials (Figure 5). Thus, the *pattern* of responses across these two LIPC regions differed along the dimension of cue-target associative strength suggesting that these regions mediate distinct control mechanisms, with anterior LIPC being particularly sensitive to controlled semantic retrieval demands.

Both the voxel-based and ROI analyses clearly revealed that Associative Strength and Number of Targets modulated LIPC activation, although neither analysis revealed an interaction between Associative Strength and Number of Targets. Importantly, the two main effects were observed during performance of a semantic decision task that does not demand selection (Thompson-Schill et al., 1997), thus posing a significant challenge to the selection hypothesis and lending support to the controlled retrieval perspective. One concern regarding these observations, however, is that greater LIPC activation was present in those conditions that were associated with a longer duty-cycle or increased task difficulty (as indexed by RT). Although increased demands on LIPC control processes and longer duty-cycles are entirely consistent with the theoretical model being proposed, nevertheless it would be informative to know whether LIPC activation is sensitive to controlled retrieval demands even when RTs do not differ. To address this question, a direct contrast was performed between the Weak/2-target and the Strong/4-target conditions that were accompanied by comparable RTs. This contrast revealed that both anterior LIPC and posterior LIPC were more sensitive to Associative Strength than to Number of Targets, demonstrating greater activation in the Weak/2-target condition (Table 2, Figure 3C). Thus, even when duty cycle is held constant, LIPC activation is modulated by the degree of control required during semantic retrieval.

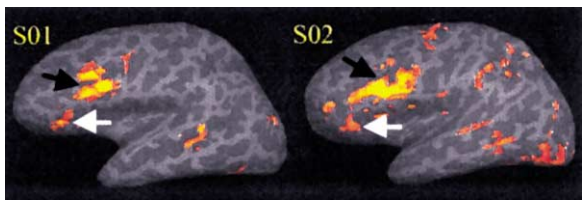

Figure 4. Distinct LIPC Regions Demonstrated Greater Activation during Weak Relative to Strong Trials

Cortical inflation illustrated the anatomical separability between the anterior LIPC (white arrow) and posterior LIPC (black arrow) regions that were sensitive to cue-target associative strength. Displayed are data from two subjects thresholded at $p < .001$; dark gray regions correspond to sulci.

Task-Related Responses in Other Regions

Although the objective of the present study was to directly test competing theories regarding the role of LIPC in the recovery of semantic knowledge, it should be noted that a few other prefrontal regions demonstrated main effects of either Number of Targets or of Associative Strength. In particular, a relatively rostral extent of the anterior cingulate cortex (ACC; Table 2, Figure 5) demonstrated a main effect of Number of Targets, whereas a slightly more caudal extent of the ACC demonstrated a main effect of Associative Strength. The

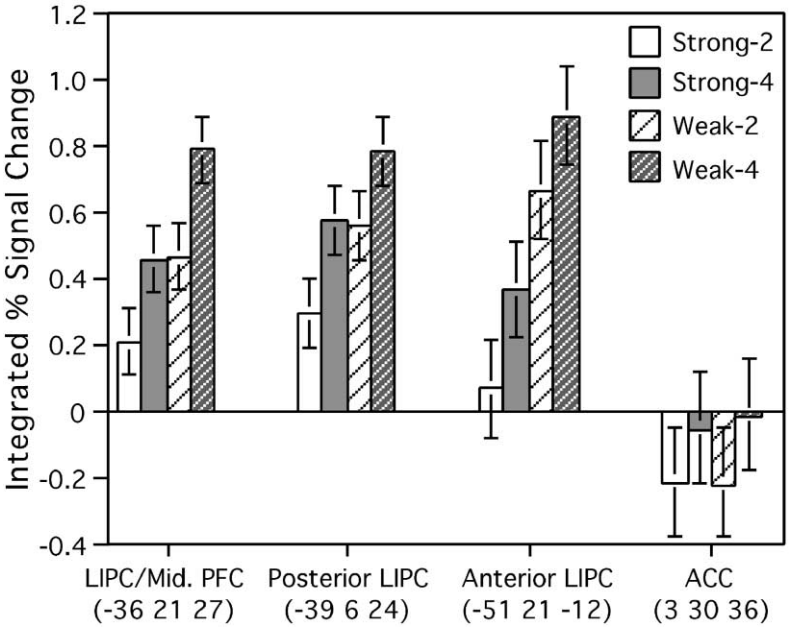

Figure 5. Region-of-Interest Analyses Revealed Differential Patterns of Activity across Prefrontal Regions

Plotted are integrated measures of percent signal change from four regions identified as demonstrating a main effect of Number of Targets; error bars reflect repeated-measures 95% confidence intervals. Although all LIPC regions demonstrated a greater response during Weak relative to Strong associative strength trials, the effect of Associative Strength differed across anterior LIPC and posterior LIPC. Specifically, posterior LIPC demonstrated a robust response during both Strong and Weak trials, whereas anterior LIPC demonstrated a robust response during Weak trials and a more modest response during Strong trials. The response in the anterior cingulate cortex (ACC) was below baseline, and did not show an effect of Associative Strength.

magnitude of activation in the former ACC region was below baseline in all four conditions, thus possibly complicating interpretation of this result. Activation in the latter ACC region was above baseline. However, in contrast to anterior LIPC and posterior LIPC, activation in this region tracked task difficulty. Thus, we note these ACC responses but refrain from speculating about their functional significance.

The Weak associative strength trials also elicited greater activation in right inferior prefrontal cortex (RIPC) relative to Strong associative strength trials (Table 2, Figure 3). As with the LIPC regions, RIPC responses were above baseline (with the exception of 45, 21, 6) and demonstrated greater activation when comparing the Weak/2-target trials to the Strong/4-target trials (and thus did not track task difficulty; Table 2). Interestingly, RIPC activation failed to demonstrate a main effect of Number of Targets, both in voxel-based analyses and in ROI analyses, suggesting that this region is selectively modulated by controlled retrieval demands. Although RIPC activation occasionally has been observed in prior studies of semantic retrieval (e.g., Wagner et al., 1998b; Poldrack et al., 1999), at present the factors that determine when RIPC activation will cooccur with LIPC activation remain undetermined. While speculative, the present results raise the possibility that when controlled retrieval demands are particularly high, RIPC processes may be additionally recruited to guide the recovery of semantic knowledge.

Finally, the two hypotheses being tested, controlled retrieval and selection, both posit that LIPC plays a role in accessing semantic knowledge that is assumed to be represented elsewhere in posterior cortex (e.g., Thompson-Schill et al., 1999). Prior neuropsychological and neuroimaging evidence would suggest that such long-term conceptual knowledge is, at least partially, represented in lateral temporal cortices (Petersen et al., 1988; Martin and Chao, 2001). Consistent with this possibility, the present results revealed two left temporal responses that were modulated by the Number of Targets and by Associative Strength. Specifically, left middle temporal gyrus (~BA 21/22) demonstrated greater activation in the 4-target, relative to the 2-target, condition (coordinates of −63, −51, 6), and when cue-target associative strength was Weak relative to when it was Strong (coordinates of −57, −51, 0). In addition, left fusiform cortex (~BA 37) demonstrated a greater response when associative strength was Weak relative to when it was Strong (coordinates of −42, −51, −15). Both regions have been observed in prior studies of semantic processing (see Kirchhoff et al., 2000; Koutstaal et al., 2001 for discussion), and thus these effects are consistent with the possibility that LIPC control mechanisms provide a top-down bias that modulates left temporal processing in the course of semantic retrieval.

Discussion

The present fMRI study sought to address a central and, as yet, unanswered question regarding prefrontal contributions to the control of memory: What is the functional role of LIPC in the recovery of meaning? The results revealed three important outcomes. First, consistent with the hypothesis that LIPC subserves the controlled retrieval of long-term semantic knowledge, the level of control required during semantic retrieval modulated LIPC activation. Importantly, these effects were observed within the context of a semantic decision task that did not require selection between competing semantic representations. Second, two anatomically distinct regions within LIPC were observed to be sensitive to controlled retrieval demands, but the effects of varying mnemonic control on the pattern of activation in these LIPC regions were distinct. Specifically, anterior LIPC (BA 47/45) demonstrated greater sensitivity to cue-target associative strength than did posterior LIPC (BA 44/6). Third, directly failing to replicate the central previ-

ous observation proffered against a semantic retrieval interpretation of LIPC function, the results revealed greater LIPC activation when semantic retrieval demands increased but selection was held constant. Collectively, these data are inconsistent with the hypothesis that LIPC function is restricted to the selection of task-relevant representations from amidst competing knowledge. Rather, these data provide the strongest evidence to date indicating that LIPC mediates controlled semantic retrieval.

LIPC Activation Is Modulated by Controlled Retrieval Demands

Compelling support for the controlled semantic retrieval hypothesis derives from the present observation that LIPC activation increased as the cue-target associative strength weakened. An extensive literature points to the centrality of prefrontal function in the control of cognition and behavior (Norman and Shallice, 1986; Cohen et al., 1996; Miller and Cohen, 2001). When goal-directed behavior cannot be performed based on relatively automatic knowledge retrieval due to prepotent stimulus-stimulus and stimulus-response associations, successful behavior becomes increasingly dependent on prefrontal function. The observation that preexperimental cue-target associative strength markedly modulates anterior LIPC and posterior LIPC activation provides striking evidence for a role of LIPC in the controlled recovery of meaning (Wagner, 2001).

Prior neuroimaging studies suggest that anterior LIPC may play a more central role in the retrieval and evaluation of semantic knowledge relative to posterior LIPC (Poldrack et al., 1999). For example, direct comparisons of tasks that differentially depend on semantic and phonological control have consistently implicated anterior LIPC in the former and posterior LIPC in the latter (Fiez, 1997; Price et al., 1997; Poldrack et al., 1999; Kirchhoff et al., 2000). Moreover, recent data indicate that anterior LIPC activation is selectively influenced by prior semantic processing, whereas posterior LIPC activation is modulated by earlier semantic and nonsemantic processing (Wagner et al., 2000a). That is, experience-dependent reductions in engagement of anterior LIPC mechanisms appear to be specifically associated with increased automaticity of semantic access.

The present findings build on these prior observations as they indicate that, even in the absence of selection demands, greater controlled semantic retrieval requirements elicit increased activation in both anterior LIPC and posterior LIPC. Importantly, however, the present data further provide strong evidence for a differential role of these two LIPC regions in controlled semantic retrieval. Specifically, the region (anterior LIPC/posterior LIPC) X associative strength (Strong/Weak) interaction indicates that anterior LIPC is particularly associated with the controlled retrieval of semantic knowledge.

LIPC Activation Is Modulated by Semantic Retrieval Demands

The selection hypothesis posits that LIPC does not guide semantic retrieval per se, but rather supports the selection of task-relevant knowledge from amidst competing irrelevant knowledge (Thompson-Schill et al., 1997; Thompson-Schill et al., 1998; Thompson-Schill et al., 1999). The primary source of evidence against the semantic retrieval hypothesis was the failure to observe increased activation under low selection conditions when semantic retrieval demands putatively varied (i.e., when comparing a 4-target to a 2-target condition, Thompson-Schill et al., 1997). In contrast to this earlier reported null result, however, the present experiment provides clear evidence that an increase in the number of targets, and thus putative semantic retrieval demands, elicits greater LIPC activation. This effect was observed in both posterior LIPC and anterior LIPC.

One account for the discrepancy between the present findings and those of Thompson-Schill and colleagues is that there was a strong association between the cue and the correct target in all "low selection" trials in that earlier study (Thompson-Schill et al., 1997). Thus, performance of the semantic decision task might have been based on more automatic semantic retrieval processes that yield recovery of sufficient knowledge for identifying the highly associated stimulus without relying on LIPC control processes. Moreover, given the blocked experimental design implemented in that earlier study, subjects might have been able to adopt a performance strategy based on automatic semantic access without encountering a significant performance cost. That is, because all trials consisted of one strongly associated target and a set of unassociated targets, subjects might have been able to depend solely on the semantic knowledge accessed in a more automatic manner in order to determine the correct target. The present experiment, in contrast, implemented an event-related design that intermixed trials in which the target response was a strong associate of the cue and trials in which the target response was a weak associate of the cue. Given this design, reliance on more automatic retrieval processes might not have been afforded as such an approach would have led to failure in the Weak trials. Thus, the present event-related design likely encouraged strategic engagement of controlled semantic retrieval mechanisms during all trials. Consistent with this perspective, the present data revealed that LIPC activation was sensitive to the number of targets when the cue-target associative strength was weak as well as when associative strength was strong.

The effect of number of targets on LIPC activation is consistent with the perspective that LIPC is sensitive to increased semantic retrieval demands. However, it should be noted that moving from two to four targets also increases demands on phonological and lexical processes. Notably, increasing the number of targets in the present study yielded increased activation in posterior LIPC, a region that has been previously implicated in the control of phonological representations (Awh et al., 1996; Fiez et al., 1996). Posterior LIPC, but not anterior LIPC, also has been observed during performance of "nonsemantic" processing tasks that require or permit access to phonological codes (Poldrack et al., 1999; Wagner et al., 2000a). Moreover, neuropsychological evidence indicates that posterior LIPC lesions can impair performance on lexical or phonological processing tasks (Fiez and Petersen, 1998; Swick, 1998). Thus, the presently observed increase in posterior LIPC activation may derive from differential demands on phonological

control processes, rather than demands specifically related to semantic retrieval. By contrast, the pattern of activation in anterior LIPC suggests that engagement of this region is modulated by controlled semantic retrieval requirements.

LIPC Contributions to the Recovery of Meaning

In contrast to the selection hypothesis, the controlled retrieval perspective can readily account for the effects of associative strength and number of targets on LIPC activation. The controlled retrieval hypothesis is also capable of handling earlier observations that have been posited to specifically support the selection account. Notably, the cue-target associative strength appears to have covaried with selection demands in most prior studies supporting selection (Thompson-Schill et al., 1997; Thompson-Schill et al., 1998), with associative strength being weaker in the "high-selection" than in the "low-selection" conditions. Thus, such effects of "selection demands" may actually reflect the influence of differing demands on controlled semantic retrieval. Moreover, in contrast to a recent null report that preexperimental associative strength does not modulate left prefrontal activation, a finding that was interpreted as supporting the selection hypothesis (Barch et al., 2000; but see, Fletcher et al., 2000), again, the present data demonstrate that LIPC activation is sensitive to cue-target associative strength.

The present findings are consistent with a unitary perspective on anterior LIPC function that posits that this LIPC region subserves a top-down bias signal that facilitates the controlled recovery of task-relevant knowledge. Importantly, this bias mechanism may be differentially engaged under retrieval conditions where the target knowledge is not recovered through more automatic access processes. Automatic access may be insufficient or may fail either because the cue-target associative strength is not strong enough to result in automatic knowledge recovery, or because task-irrelevant representations compete with, and interfere with the automatic recovery of, task-relevant knowledge. From this unitary perspective, controlled semantic retrieval may constitute the most basic level functional outcome of top-down facilitative bias of semantic memory. Moreover, in the process of guiding the controlled recovery of semantic knowledge, this bias mechanism may tune or "sculpt" semantic space (Fletcher et al., 2000; Wagner et al., 2000b) such that the retrieved knowledge is readily more accessible in the future (Raichle et al., 1994; Demb et al., 1995; Thompson-Schill et al., 1999). Finally, consistent with the bias competition model of prefrontal function (Desimone and Duncan, 1995; Miller and Cohen, 2001), this controlled bias mechanism may also contribute to the recovery of meaning when task-relevant knowledge must be favored or selected from amidst competitive knowledge (Thompson-Schill et al., 1997; Jonides et al., 1998; Thompson-Schill et al., 1998; Thompson-Schill et al., 1999; Wagner et al., 2000a). That is, from a unitary perspective, selection may constitute a specific instantiation of top-down controlled retrieval.

Although a unitary account of LIPC function might be the most satisfactory model, a challenge to this perspective remains. Specifically, although the controlled semantic retrieval and the selection hypotheses can be integrated at the mechanistic level, it is important to emphasize that extant data suggest that the neurobiology may differ. The LIPC region implicated in selection (at or near BA 44) appears to fall well posterior and dorsal to the anterior LIPC region that is clearly implicated in controlled semantic retrieval in this and related studies (Gabrieli et al., 1998; Poldrack et al., 1999; Wagner et al., 2000a). Rather, the LIPC region implicated in studies of selection appears to correspond to the posterior LIPC region observed in the present study and in studies of phonological working memory and control. It remains possible that the LIPC subregions that support controlled semantic retrieval may be distinct from, or may be only partially overlapping with, those that guide selection. The present design did not manipulate selection demands, and thus our data do not indicate whether the anterior LIPC region observed to be differentially sensitive to controlled retrieval demands is also modulated by selection demands when cue-target associative strength is held constant. Additional studies that cross selection demands with controlled retrieval demands should serve to adjudicate between the two possibilities. Nevertheless, the present results provide some of strongest evidence that, when selection demands are held constant, the left prefrontal cortex supports the controlled retrieval of long-term semantic knowledge. In so doing, LIPC plays a fundamental role in goal-directed behavior that depends on the recovery of meaning.

Experimental Procedures

Participants

Fourteen right-handed, native speakers of English (five male; ages 18–40 yr) were remunerated $50 for their participation. Data from one additional participant were excluded due to poor task performance (<65% correct). Informed consent was obtained in a manner approved by the Human Studies Committee of the Massachusetts General Hospital.

Stimuli and Paradigm

On each of 288 trials, a single cue word and either two or four target items were presented. Subjects decided which of the targets was most closely globally related to the cue (by pressing one of four response keys). The correct response was either a strong or a weak associate of the cue, as detailed below. This task is analogous to the global similarity task of Thompson-Schill and colleagues (1997), and thus requires no selection, as all semantic attributes are relevant to computing global relatedness. The words were presented for 3 s, with a 1 s fixation following stimuli offset (participants had all 4 s to respond). Periods of fixation lasting either 2 or 4 s were interspersed between experimental trials as determined by an optimization algorithm (Dale, 1999). Each subject participated in four fMRI runs, with the order of trials determined by optimizing the efficiency of the design matrix (Dale, 1999).

Items for the experimental trials were chosen from associative norms requiring generation of a single associate to a cue (Postman and Keppel, 1970). For each cue, one strongly and one weakly associated target item were chosen. The mean normative probability that the item was generated as the associate of the cue differed across Strong (.22) and Weak (.01) targets. Distractors were chosen randomly (Kucera and Francis, 1967), with the constraint that they were unrelated to their assigned cues. Weak target, Strong target, and distractor lists were constructed matched on word frequency and word length. Eight lists of 72 distractors and four lists each of cues and their weakly and strongly associated targets were created. Trials were created based upon these lists such that there were 72

trials in each of four conditions. Across subjects, lists were counterbalanced such that each cue and distractor list appeared in each condition (2 versus 4 targets, Weak versus Strong associative strength). Target lists were counterbalanced across number of targets (2 versus 4), and median word frequency (Kucera and Francis, 1967) and word length were equated across weakly and strongly associated target lists. The experimental paradigm, implemented using PsyScope (Cohen et al., 1993), is illustrated in Figure 1.

Functional Imaging

Scanning was performed on a 1.5T Siemens Sonata MRI system using a whole-head coil. Functional data were acquired using a gradient-echo echo-planar pulse sequence (TR = 2 s, TE = 40 ms, 21 axial slices, 3.125 × 3.125 × 5 mm, 1 mm inter-slice gap, 180 volume acquisitions per run). High-resolution T1-weighted (MP-RAGE) anatomical images were collected for anatomical visualization. Head motion was restricted using a pillow that surrounded the head. Visual stimuli were projected via a collimating lens onto a screen, which was viewed through a mirror attached to the head coil.

Data were preprocessed using SPM99 (Wellcome Dept. of Cognitive Neurology, London). Images were corrected for differences in slice acquisition timing by resampling all slices in time to match the first slice, followed by motion correction across all runs (using sinc interpolation). Structural and functional data were spatially normalized to an EPI template based on the MNI305 stereotactic space (Cocosco et al., 1997), an approximation of canonical space (Talairach and Tournoux, 1988), using a 12-parameter affine transformation along with a nonlinear transformation using cosine basis functions. Images were resampled into 3 mm cubic voxels and then spatially smoothed with an 8 mm FWHM isotropic Gaussian kernel.

Statistical analysis was performed using the general linear model in SPM99. Correct trials from each condition were modeled using a canonical hemodynamic response, with error trials from all conditions modeled separately. Effects were estimated using a subject-specific fixed-effects model, with session-specific effects and low-frequency signal components treated as confounds. Linear contrasts were used to obtain subject-specific estimates for each effect. These estimates were entered into a second-level analysis treating subjects as a random effect, using a one-sample t test against a contrast value of zero at each voxel. Given a priori expectations that LIPC activation would accompany performance of the relatedness task, regions were considered reliable to the extent that they consisted of at least five contiguous voxels that exceeded an uncorrected threshold of $p < .001$.

The group-level voxel-based contrasts were supplemented with region-of-interest (ROI) analyses that further characterized the effects of number of targets and of cue-target associative strength in functionally defined regions. For the ROI analyses, spherical regions of interest were identified by choosing all significant voxels within a 5 mm radius of the chosen maximum identified in the group statistical map. Signal within each ROI was then calculated for each individual subject by selectively averaging the data with respect to peristimulus time for trials in each condition. Integrated percent signal change was calculated by summing the hemodynamic response over the 0–14 s peristimulus window. The resultant data were then subjected to mixed-effect analysis of variance (ANOVA) that treated number of targets and associative strength as repeated measures, and subjects as a random effect.

Finally, to further visualize the anatomical separability of anterior LIPC and posterior LIPC regions and the location of these activations with respect to inferior frontal sulcus, inferior and premotor gyri, and anterior insula, the cortical surface was reconstructed from high-resolution MR scans of each subject and inflated using automated techniques described previously (Dale et al., 1999; Fischl et al., 1999). Statistical analyses using SPM99 were performed on individual subject data in native image space following smoothing with an 8 mm Gaussian kernel. These data were coregistered with the high-resolution anatomical image using SPM99, and the coregistered statistical maps were sampled onto the surface reconstruction and rendered on the inflated surface for visualization.

Acknowledgments

Supported by the NIDCD (DC04466), Ellison Medical Foundation, Surdna Foundation, P. E. Newton, and the Alafi Family Foundation. We thank the Learning & Memory Lab at MIT for insightful comments and discussion, and B. Fischl, D. Greve, R. Desikan, D. Salat, and A. Dale for assistance with cortical reconstruction.

Received February 21, 2001; revised May 17, 2001.

References

Asaad, W.F., Rainer, G., and Miller, E.K. (1998). Neural activity in the primate prefrontal cortex during associative learning. Neuron *21*, 1399–1407.

Awh, E., Jonides, J., Smith, E.E., Schumacher, E.H., Koeppe, R.A., and Katz, S. (1996). Dissociation of storage and rehearsal in verbal working memory: evidence from PET. Psychol. Sci. *7*, 25–31.

Barch, D.M., Braver, T.S., Sabb, F.W., and Noll, D.C. (2000). Anterior cingulate and the monitoring of response conflict: evidence from an fMRI study of overt verb generation. J. Cogn. Neurosci. *12*, 298–309.

Bokde, A.L.W., Tagaments, M.-A., Friedman, R.B., and Horwitz, B. (2001). Functional interactions of the inferior frontal cortex during the processing of words and word-like stimuli. Neuron, *30*, 609–617.

Buckner, R.L. (1996). Beyond HERA: contributions of specific prefrontal brain areas to long-term memory retrieval. Psychonom. Bull. Rev. *3*, 149–158.

Cocosco, C.A., Kollokian, V., Kwan, R.K.S., and Evans, A.C. (1997). Brainweb: online interface to a 3D MRI simulated brain database. NeuroImage *5*, 425.

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 Instrum. Comput. *25*, 257–271.

Cohen, J.D., Braver, T.S., and O'Reilly, R.C. (1996). A computational approach to prefrontal cortex, cognitive control and schizophrenia: recent developments and current challenges. Philos. Trans. R. Soc. Lond. B. Biol. Sci. *351*, 1515–1527.

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.

Dale, A.M. (1999). Optimal experimental design for event-related fMRI. Hum. Brain Mapp. *8*, 109–114.

Dale, A.M., Fischl, B., and Sereno, M.I. (1999). Cortical surface-based analysis. I. Segmentation and surface reconstruction. NeuroImage *9*, 179–194.

Davachi, L., Maril, A., and Wagner, A.D. (2001) When keeping in mind supports later bringing to mind: neural markers of phonological rehearsal predict subsequent remembering. J. Cogn. Neurosci., in press.

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.

Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. Annu. Rev. Neurosci. *18*, 193–222.

Fiez, J.A. (1997). Phonology, semantics, and the role of the left inferior prefrontal cortex. Hum. Brain Mapp. *5*, 79–83.

Fiez, J.A., and Petersen, S.E. (1998). Neuroimaging studies of word reading. Proc. Natl. Acad. Sci. USA *95*, 914–921.

Fiez, J.A., Raife, E.A., Balota, D.A., Schwarz, J.P., Raichle, M.E., and Petersen, S.E. (1996). A positron emission tomography study of the short-term maintenance of verbal information. J. Neurosci. *16*, 808–822.

Fischl, B., Sereno, M.I., and Dale, A.M. (1999). Cortical surface-based analysis. II: inflation, flattening, and a surface-based coordinate system. NeuroImage *9*, 195–207.

Fletcher, P.C., Shallice, T., and Dolan, R.J. (2000). "Sculpting the response space"–An account of left prefrontal activation at encoding. NeuroImage *12*, 404–417.

Fuster, J.M. (1997). The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe (Philadelphia: Lippincott-Raven).

Gabrieli, J.D.E., Desmond, J.E., Demb, J.B., Wagner, A.D., Stone, M.V., Vaidya, C.J., and Glover, G.H. (1996). Functional magnetic resonance imaging of semantic memory processes in the frontal lobes. Psychol. Sci. *7*, 278–283.

Gabrieli, J.D., Poldrack, R.A., and Desmond, J.E. (1998). The role of left prefrontal cortex in language and memory. Proc. Natl. Acad. Sci. USA *95*, 906–913.

Goldman-Rakic, P.S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In Handbook of Physiology, Section 1: The Nervous System, Vol. V. Higher Functions of the Brain, Part 1., F. Plum and V. Mountcastle, eds. (Bethesda, MD: American Physiological Society), pp. 373–417.

Goldman-Rakic, P.S. (1995). Architecture of the prefrontal cortex and the central executive. In Structure and functions of the Human Prefrontal Cortex, J. Grafman, F. Boller, and K. Holyoak, eds. (New York: New York Academy of Sciences), pp. 71–83.

Jonides, J., Smith, E.E., Marshuetz, C., and Koeppe, R.A. (1998). Inhibition in verbal working memory revealed by brain activation. Proc. Natl. Acad. Sci. USA *95*, 8410–8413.

Kapur, S., Rose, R., Liddle, P.F., Zipursky, R.B., Brown, G.M., Stuss, D., Houle, S., and Tulving, E. (1994). The role of the left prefrontal cortex in verbal processing: semantic processing or willed action? NeuroReport *5*, 2193–2196.

Kelley, W.M., Miezin, F.M., McDermott, K.B., Buckner, R.L., Raichle, M.E., Cohen, N.J., Ollinger, J.M., Akbudak, E., Conturo, T.E., Snyder, A.Z., and Petersen, S.E. (1998). Hemispheric specialization in human dorsal frontal cortex and medial temporal lobe for verbal and nonverbal memory encoding. Neuron *20*, 927–936.

Kirchhoff, B.A., Wagner, A.D., Maril, A., and Stern, C.E. (2000). Prefrontal-temporal circuitry for novelty encoding and subsequent memory. J. Neurosci. *20*, 6173–6180.

Koutstaal, W., Wagner, A.D., Rotte, M., Maril, A., Buckner, R.L., and Schacter, D.L. (2001). Perceptual specificity in visual object priming: functional magnetic resonance imaging evidence for a laterality difference in fusiform cortex. Neuropsychologia, *39*, 184–199.

Kucera, H., and Francis, W.N. (1967). Computational Analysis of Present-day English (Providence, RI: Brown University Press).

Martin, A., and Chao, L.L. (2001). Semantic memory and the brain: structure and processes. Curr. Opin. Neurobiol., in press.

Martin, A., Haxby, J.V., Lalonde, F.M., Wiggs, C.L., and Ungerleider, L.G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. Science *270*, 102–105.

Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci., in press.

Norman, D.A., and Shallice, T. (1986). Attention to action: willed and automatic control of behavior. In Consciousness and Self-Regulation, R.J. Davidson, G.E. Schwartz, and D. Shapiro, eds. (New York: Plenum).

Owen, A.M. (1997). The functional organization of working memory processes within human lateral frontal cortex: the contribution of functional neuroimaging. Eur. J. Neurosci. *9*, 1329–1339.

Passingham, R.E., Toni, I., and Rushworth, M.F. (2000). Specialisation within the prefrontal cortex: the ventral prefrontal cortex and associative learning. Exp. Brain Res. *133*, 103–113.

Paulesu, E., Frith, C.D., and Frackowiak, R.S. (1993). The neural correlates of the verbal component of working memory. Nature *362*, 342–345.

Petersen, S.E., Fox, P.T., Posner, M.I., Mintun, M., and Raichle, M.E. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. Nature *331*, 585–589.

Petrides, M. (1996). Specialized systems for the processing of mnemonic information within the primate frontal cortex. Philos. Trans. R. Soc. Lond. B. Biol. Sci. *351*, 1455–1462.

Poldrack, R.A., Wagner, A.D., Prull, M.W., Desmond, J.E., Glover, G.H., and Gabrieli, J.D.E. (1999). Functional specialization for semantic and phonological processing in the left inferior frontal cortex. NeuroImage *10*, 15–35.

Postman, L., and Keppel, G. (1970). Norms of Word Association (New York: Academic Press).

Price, C.J., Moore, C.J., Humphreys, G.W., and Wise, R.S.J. (1997). Segregating semantic from phonological processes during reading. J. Cogn. Neurosci. *9*, 727–733.

Raichle, M.E., Fiez, J.A., Videen, T.O., Macleod, A.M.K., Pardo, J.V., Fox, P.T., and Petersen, S.E. (1994). Practice-related changes in human brain functional anatomy during nonmotor learning. Cereb. Cortex *4*, 8–26.

Shallice, T. (1988). From Neuropsychology to Mental Structure (New York: Cambridge University Press).

Shimamura, A.P. (1995). Memory and frontal lobe function. In The Cognitive Neurosciences, M.S. Gazzaniga MS, ed. (Cambridge, MA: MIT Press), pp 803–813.

Stuss, D.T., and Benson, D.F. (1984). Neuropsychological studies of the frontal lobes. Psychol. Bull. *95*, 3–28.

Swick, D. (1998). Effects of prefrontal lesions on lexical processing and repetition priming: an ERP study. Cogn. Brain Res. *7*, 143–157.

Talairach, J., and Tournoux, P. (1988). A Co-Planar Stereotactic Atlas of the Human Brain (Stuttgart: Thieme).

Thompson-Schill, S.L., D'Esposito, M., Aguirre, G.K., and Farah, M.J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. Proc. Natl. Acad. Sci. USA *94*, 14792–14797.

Thompson-Schill, S.L., Swick, D., Farah, M.J., D'Esposito, M., Kan, I.P., and Knight, R.T. (1998). Verb generation in patients with focal frontal lesions: a neuropsychological test of neuroimaging findings. Proc. Natl. Acad. Sci. USA *95*, 15855–15860.

Thompson-Schill, S.L., D'Esposito, M., and Kan, I.P. (1999). Effects of repetition and competition on activity in left prefrontal cortex during word generation. Neuron *23*, 513–522.

Vandenberghe, R., Price, C., Wise, R., Josephs, O., and Frackowiak, R.S. (1996). Functional anatomy of a common semantic system for words and pictures. Nature *383*, 254–256.

Wagner, A.D. (1999). Working memory contributions to human learning and remembering. Neuron *22*, 19–22.

Wagner, A.D. (2001). Cognitive control and episodic memory: contributions from prefrontal cortex. In Neuropsychology of Memory, 3rd ed. L.R. Squire and D.L. Schacter, eds. (New York: Guilford Press).

Wagner, A.D., Desmond, J.E., Demb, J.B., Glover, G.H., and Gabrieli, J.D.E. (1997). Semantic repetition priming for verbal and pictorial knowledge: a functional MRI study of left inferior prefrontal cortex. J. Cogn. Neurosci. *9*, 714–726.

Wagner, A.D., Poldrack, R.A., Eldridge, L.L., Desmond, J.E., Glover, G.H., and Gabrieli, J.D. (1998a). Material-specific lateralization of prefrontal activation during episodic encoding and retrieval. NeuroReport *9*, 3711–3717.

Wagner, A.D., Schacter, D.L., Rotte, M., Koutstaal, W., Maril, A., Dale, A.M., Rosen, B.R., and Buckner, R.L. (1998b). Building memories: remembering and forgetting of verbal experiences as predicted by brain activity. Science *281*, 1188–1191.

Wagner, A.D., Koutstaal, W., Maril, A., Schacter, D.L., and Buckner, R.L. (2000a). Task-specific repetition priming in left inferior prefrontal cortex. Cereb. Cortex *10*, 1176–1184.

Wagner, A.D., Maril, A., and Schacter, D.L. (2000b). Interactions between forms of memory: when priming hinders new episodic learning. J. Cogn. Neurosci. *12:S2*, 52–60.