The Influence of Memory Load Upon Delay-Interval Activity in a Working-Memory Task: An Event-Related Functional MRI Study

Amishi P. Jha
Duke University

Gregory McCarthy
Duke University and Department of Veterans Affairs Medical Center

Abstract

■ We conducted two fMRI studies to investigate the sensitivity of delay-period activity to changes in memory load during a delayed-recognition task for faces. In Experiment 1, each trial began with the presentation of a memory array consisting of one, two, or three faces that lasted for 3 sec. A 15-sec delay period followed during which no stimuli were present. The delay interval concluded with a one-face probe to which subjects made a button press response indicating whether this face was part of the memory array. Experiment 2 was similar in design except that the delay period was lengthened to 24 sec, and the memory array consisted of only one or three faces. We hypothesized that memory maintenance processes that spanned the delay interval would be revealed by their sensitivity to memory load. Long delay intervals were employed to temporally dissociate phasic activity engendered by the memory array from sustained activity reflecting maintenance. Regions of interest (ROIs) were defined anatomically for the superior frontal gyri (SFG), middle frontal gyri (MFG), and inferior frontal gyri (IFG), intraparietal sulci (IPS), and fusiform gyri (FFG) on a subject-by-subject basis. The mean time course of activity was determined for all voxels within these regions and for that subset of voxels within each ROI that correlated significantly with an empirically determined reference waveform. In both experiments, memory load significantly influenced activation 6–9 sec following the onset of the memory array with larger amplitude responses for higher load levels. Responses were greatest within MFG, IPS, and FFG. In both experiments, however, these load-sensitive differences declined over successive time intervals and were no longer significant at the end of the delay interval. Although insensitive to our load manipulation, sustained activation was present at the conclusion of the delay interval within MFG and other prefrontal regions. IPS delay activity returned to prestimulus baseline levels prior to the end of the delay period in Experiment 2, but not in Experiment 1. Within FFG, delay activity returned to prestimulus baseline levels prior to the conclusion of the delay interval in both experiments. Thus, while phasic processes engendered by the memory array were strongly affected by memory load, no evidence for load-sensitive delay-spanning maintenance processes was obtained. ■

INTRODUCTION

Working memory is a psychological construct encompassing executive functions responsible for the manipulation of information and maintenance functions that keep information active over short intervals (Baddeley, 1996). Physiological studies in nonhuman primates have identified a network of brain regions involved in working memory that includes dorsolateral prefrontal cortex (dlPFC) (Fuster, 1997; Goldman-Rakic, 1987), the intraparietal sulcus (IPS) (Chafee & Goldman-Rakic, 2000), and posterior perceptual areas (Miller, Li, & Desimone, 1991). Neuroimaging studies in humans have also demonstrated a consistent pattern of activation in similar regions across a variety of working-memory tasks (Belger et al., 1998; Courtney, Ungerleider, Keil, & Haxby, 1997; D'Esposito et al., 1995; Haxby, Ungerleider, Horwitz, Rapoport, & Grady, 1995; Cohen et al., 1994; McCarthy et al., 1994).

Great effort has been expended to parse working memory into its component psychological processes and localize such processes within this neural architecture. This has proven to be a difficult undertaking as tasks that manipulate the core executive and maintenance functions of working memory can also introduce variation in other processes such as those necessary for encoding and retrieving items to and from memory, and for selecting and preparing responses. Thus, there remains a controversy about which neural structures

are engaged by which component psychological processes.

Our aim here is to investigate claims concerning the role of different brain regions in the maintenance of information in working memory. Maintenance occurs after mnemonic encoding of the memoranda has taken place and should remain engaged until a response concerning the maintained information has been executed. In neuroimaging studies in humans, memory maintenance has typically been investigated by systematically varying the number of items to be maintained (the memory load) and observing whether brain activation increases with increasing load. Several recent neuroimaging studies have reported load-sensitive increases in activity within prefrontal cortex (PFC) for both verbal and nonverbal stimuli (e.g., Sammer, 1999; Braver et al., 1997; Desmond, Gabrieli, Wagner, Ginier, & Glover, 1997; Jonides et al., 1997; Manoach et al., 1997; Jackson, Jackson, & Rosicky, 1995; Mellers et al., 1995; Petrides, Alivisatos, Meyer, & Evans, 1993). However, there remains considerable debate over facts and interpretations.

Part of this controversy stems from the reliance of several studies upon tasks that require the simultaneous engagement of maintenance and executive processes, such as the *n*-back task. The *n*-back task is problematic because maintenance is not isolated from other task requirements; indeed, each stimulus must be successively encoded, maintained, and compared to other stimuli in memory to determine which response should be selected. Perhaps for this reason, there has been a lack of consistency in the reported relationship between PFC activity and memory load across different studies using the *n*-back task. While some reports have suggested a linear relationship between increases in memory load and PFC activity (Braver et al., 1997), others have reported step functions (Cohen et al., 1997), or inverted-U–shaped functions indicating that, although dlPFC activation increased at moderate memory loads, it decreased at high memory loads (Callicott et al., 1999).

The delayed-recognition task engages component processes of working memory at different points in time, and, thus, is a better choice for isolating processes related to maintenance. In this task, subjects are instructed to remember stimulus items presented at the start of the trial (S1). A delay interval then ensues in which these stimuli are absent but must be maintained in memory. The trial ends with the presentation of a probe stimulus (S2). Subjects are asked to judge whether the probe was part of the initial memory set and to respond accordingly. In the typical delayed-recognition task, the delay period is uninterrupted with no concurrent processing demands. Thus, the delayed-recognition task temporally segregates memory encoding following S1, memory maintenance persisting throughout the delay interval, and retrieval and response selection following S2 presentation. Such tasks have been used extensively during recording of electrophysiological responses of single neurons within PFC and other brain regions. A subset of neurons has been discovered that fire throughout the delay interval (e.g., Goldman-Rakic, 1996; Funahashi, Bruce, & Goldman-Rakic, 1989), thus providing a plausible neuronal substrate for memory maintenance. It is notable that other subsets of neurons within this same region have different response profiles including those that respond exclusively to S1 or S2, or during multiple phases of the task (Funahashi et al., 1989).

Neuroimaging studies have also used the delayed-recognition task to investigate the neural substrates of working memory within PFC and posterior regions in humans. For example, Haxby et al. (1995) conducted a PET study in which subjects were presented with a single face to maintain in working memory. After a delay period, two probe faces were presented, and subjects indicated which of the two matched the memory stimulus. Memory load was not manipulated, but the duration of the delay period was varied, and behavioral performance decreased with increasing delay. Greater activation was observed in PFC for longer delay intervals consistent with the view that increased neuronal activity was required to maintain the information over these intervals. A subsequent event-related fMRI study from the same group (Courtney et al., 1997) reported sustained activity within the dorsolateral PFC (dlPFC) during the fixed 8-sec delay interval of a delayed-recognition task for faces. As before, these authors (Courtney et al., 1997) attributed the sustained activity to maintenance processes spanning the delay interval. However, in neither study was memory load directly manipulated; the sustained delay activity was simply interpreted as evidence for maintenance. The sustained delay-period activity could have been related to any of a number of intervening processes such as generalized response preparation, inhibition, expectancy, or mental timing. Indeed, it may be that a sustained pattern of activation is typical of prefrontal responses. Moreover, without varying the delay interval, it cannot be concluded that such activity is, in fact, sustained. The conclusions of Courtney et al. have been recently challenged by two fMRI studies that have manipulated memory load during delayed-recognition tasks (Postle, Berger, & D'Esposito, 1999; Rypma & D'Esposito, 1999). Both studies concluded that delay-period activity was not influenced by memory load within both dorsal and ventral PFC.

Other studies have argued that different regions with in the PFC differentially engage executive and maintenance functions. Yet, a controversy remains regarding which regions are engaged by which functions. For example, it has been reported that while both dorsal and ventral PFC are comparably activated by tasks requiring maintenance of information during the delay period (D'Esposito, Postle, Ballard, & Lease, 1999; Postle et al., 1999), there is an additional engagement

of the dorsal PFC during manipulation tasks. Others have reported that maintenance tasks preferentially engage more inferior/ventral PFC regions, whereas monitoring engages dorsal PFC (Owen, Evans, & Petrides, 1996). An alternative formulation has been recently proposed by Rowe and colleagues (2000) who reported that Brodmann's area (BA) 46, which includes much of the MFG, was engaged preferentially by response selection, while BA 8, which includes much of the posterior superior frontal gyrus (SFG), was engaged preferentially by maintenance.

Just as the precise role of PFC in working memory is yet unresolved, the role of posterior perceptual systems is also contentious. Recent studies have suggested that these posterior systems remain active throughout the maintenance phase of working-memory tasks and are engaged in the active rehearsal of information (Awh et al., 1999). Single-unit recordings in monkeys, however, suggest that posterior perceptual systems remain active only until they are required for the perception of new stimuli (Desimone, 1996). For example, the pattern of unit activity in inferotemporal (IT) cortex is similar to the sustained activity observed in PFC during delay periods. However, unlike PFC, the activity within IT terminates upon the appearance of task-irrelevant stimuli during the delay period (Miller, Li, & Desimone, 1993). This pattern has also been observed in posterior parietal cortex during spatial working-memory task (Constantinidis & Steinmetz, 1996). These single-unit results raise the issue of whether sustained activity in posterior perceptual cortices is critical or necessary for memory performance.

Functional neuroimaging studies in humans have not resolved this issue. Haxby et al. (1995) found that activation within the right posterior fusiform gyrus (FFG) did not differ from a sensorimotor control tasks for delay periods beyond 11 sec and concluded that the ventral face areas did not maintain face representations over delay intervals. Similarly, Belger, et al. (1998) found that the activity of ventral extrastriate regions engaged in a delayed-recognition task did not differ from activity observed for nonmemory perceptual control tasks. In contrast, Courtney et al. (1997) found that posterior ventral extrastriate regions resulted in a small but significant activation during the delay period and concluded that face-processing regions remain active during maintenance. They conjectured that this involvement would terminate when these areas were recruited for the perception of new stimuli.

In addition to the involvement of PFC and posterior perceptual cortices in working memory, posterior parietal regions have been implicated. Single-unit studies report that dlPFC and posterior-parietal regions share reciprocal connections and produce nearly identical patterns of single unit activity during performance of memory-guided saccades (Chafee & Goldman-Rakic, 2000). In addition, a fMRI study by Belger et al. (1998) reported that the hemispheric pattern of activation covaried between dlPFC and the IPS during spatial and object delayed-response tasks. Thus, the extant data describe a network of brain regions activated during working-memory tasks that includes dorsal and ventral PFC, regions within the parietal lobe including the IPS, and perceptual regions within ventral extrastriate cortex. However, the process or processes represented by activation within these different regions is as yet unclear.

Here, we report two studies that investigated memory maintenance in a delayed-recognition task. We hypothesized that the neuronal process supporting memory maintenance would increase its activity as memory load was increased (at least until the limit of working-memory capacity was reached) and, furthermore, that this increased activity would be reflected by a load-sensitive sustained blood oxygenation level dependent (BOLD) signal that was maintained throughout the delay interval, analogous to the sustained firing of single units observed in monkey delayed-response tasks. We recognized that memory load may also influence phasic processes related to the registration and encoding of the memory array. If the delay interval was too short, these phasic responses might be mistakenly interpreted as delay-spanning maintenance processes. Indeed, we note that the 8-sec delay interval used by Courtney et al. (1997) was only slightly longer than the typical peak latency of a hemodynamic response (HDR) to a brief sensory stimulus. Encoding or other processes might increase the latency of the HDR and, thus, create a confound. As the HDR for simple visual stimuli does not return to baseline until approximately 12 sec for simple sensory stimuli (see, e.g., Huettel & McCarthy, 2000), we were concerned that potential variations in the latency of the phasic responses to the memory array would result in variations in its slow return to baseline that would masquerade as delay-spanning sustained activations in short delay intervals. We therefore used long delay intervals in an attempt to dissociate phasic and sustained activation related to memory load. As we hypothesized that maintenance processes span the delay interval, the critical test was whether there were load-sensitive differences in BOLD signal at the conclusion of the delay interval.

RESULTS

Behavioral Results

Accuracy (percent correct) and reaction time (RT) data are presented in Figure 1. Analysis of variance (ANOVA) tests revealed a significant main effect of memory load in both accuracy and RT. In Experiment 1, performance accuracy differed as a function of memory load with better performance for one- (99%, $SD = .02$) versus two-face (92%, $SD = .03$) ($F(2, 12) = 6.04, p < .05$), and one-versus three-face trials (89%, $SD = .03$) ($F(1, 6) = 32.17$,

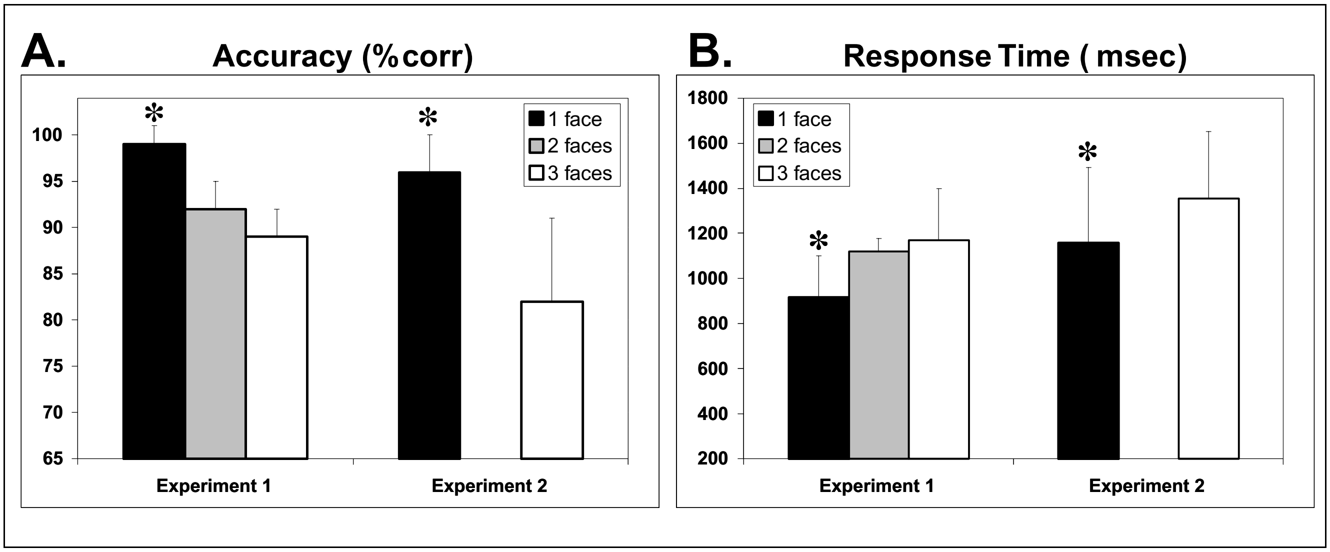

Figure 1. Behavioral results. (A) The mean percent correct scores are plotted with standard error bars for both experiments. (B) The mean response times are plotted with standard error bars for both experiments. Asterisks depict conditions that were significantly different ($p \leq .05$) from the three-face trials.

$p < .004$). As predicted, performance for the two-face load was intermediate to one- and three-face; however, accuracy for the two-face condition did not differ significantly from three-face condition. Mean RTs were significantly faster for one-face (915 msec, SD = 187 msec) versus two-face (1120 msec, SD = 56 msec) ($F(1, 6) = 18.79$, $p < .005$), and one-face versus three-face trials (1171 msec, SD = 226 msec) ($F(1, 6) = 40.2$, $p < .005$). In Experiment 2, performance accuracy differed as a function of memory load, with better performance for one-face (96%, SD = .04) versus three-face trials (82%, SD = .09) ($F(1, 8) = 13.84$, $p < .0.006$). Mean RTs were significantly faster in one-face (1158 msec, SD = 333 msec) versus three-face trials (1353 msec, SD = 298 msec) ($F(1, 8) = 10.36$, $p < .01$).

In comparing Experiments 1 and 2, a main effect of memory load was obtained for both performance accuracy and RT ($p < .05$). A main effect of the experiments (1 vs. 2) was found for accuracy ($p < .02$) with accuracy higher in Experiment 1. Thus, memory load affected behavioral performance in both experiments.

fMRI Results

Average Activation Time Courses for All Voxels Within Each ROI

Mean activation time courses for the voxels comprising each anatomically defined ROI (see Methods and Figure 2B) are illustrated in Figure 3A. These waveforms represent the contribution of all voxels within the ROI and were thus not influenced by any statistical criterion or correlational analysis. Significant S1-evoked HDRs were observed ~6–9 sec following S1 onset, and ~6 sec following S2 onset. A reduced level of activity occurred in the S1–S2 delay interval that varied considerably across the different ROIs. Although all of our ROIs produced the same general time course response, the mean signal within posterior ROIs (FFG and IPS) was greater in amplitude than ROIs comprising the PFC. Within the PFC, MFG waveforms were the largest in amplitude for all the waveshape features.

We examined activity within each ROI on a slice-by-slice basis (across all load conditions) to determine if there was an anterior to posterior gradient to the pattern of activity. The mean activation time courses for all voxels within each slice (in mm relative to the anterior commissure [AC]) are shown for the middle-frontal gyrus (MFG) and the FFG in Figure 3B. Within the MFG, the time-activation waveforms were largest anteriorly (41.25 to 26.25 mm anterior to the AC). Peak activation occurred more posteriorly along this anterior-posterior axis for the other frontal ROIs: IFG: (3.75 to 22.5 mm anterior to the AC), SFG (3.75 to 11.25 mm anterior to the AC), and cingulate gyrus (CG: 15 to 3.75 mm anterior to AC). IPS and FFG activity was greatest in more posterior slices (56 to 68 mm posterior to the AC).

A comparison of the slice-by-slice activation waveforms for the MFG and FFG (Figure 3B) is revealing. The MFG waveform shows strong activity evoked by S1 and S2 with significant above-zero activity spanning the entire delay. These three waveform features were highly correlated across slices; i.e., the slice with the largest amplitude response to S1 also had the largest amplitude response to S2 and the greatest sustained activity. For example, the correlation between the peak response to S1 and S2 across slices, hemispheres, and subjects, was r = .81, while the correlation between S2 and sustained activity was r = .74. Thus, the main waveform features did not dissociate across slices. This strong correlation among waveshape features was also true for the SFG,

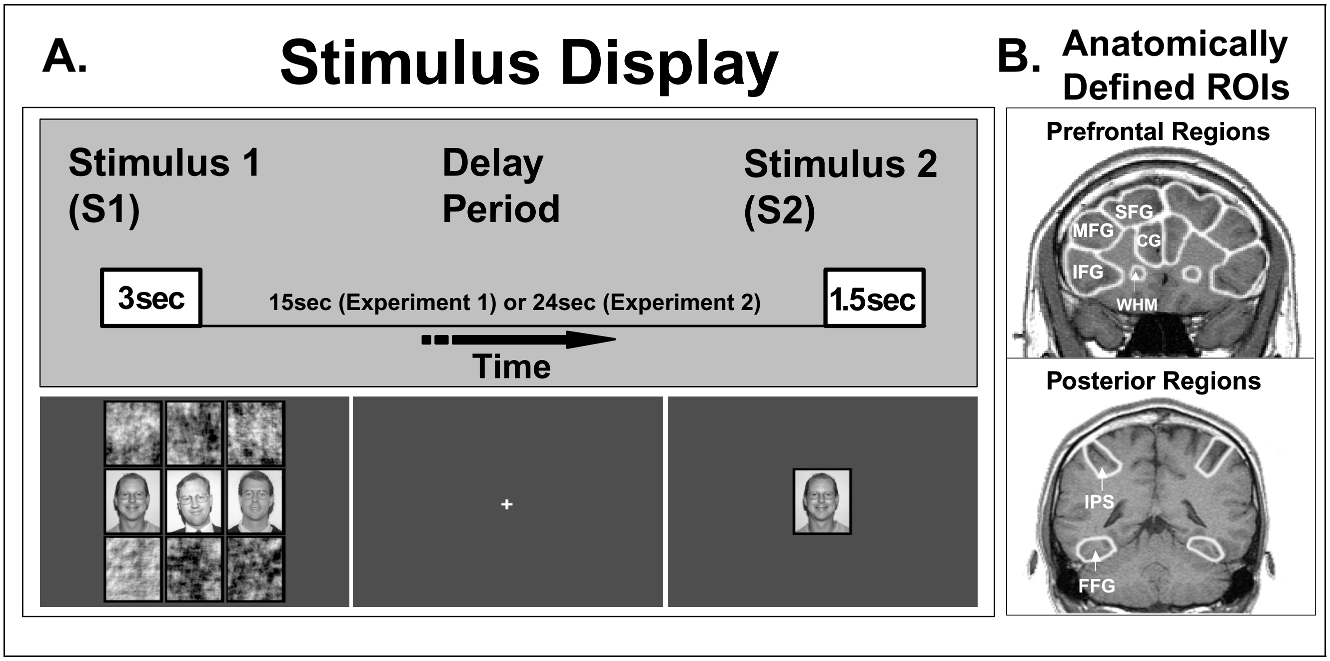

Figure 2. (A) Each trial consisted of a memory array (S1) presented for 3 sec, a delay interval free of all stimuli, and a single probe stimulus (S2) presented for 1.5 sec that terminated the delay interval. On half of the trials, the probe stimulus matched one of the items in the memory array while on the remaining half the probe was a new stimulus. The subjects made a choice-button response with the index finger of their left or right hand to indicate whether the probe did, or did not, match one of the memory items. (B). Anatomical regions of interest (ROIs) were drawn on each subject's high-resolution anatomical images by a trained observer who was blind to other statistical analyses of the data. These regions were superior (SFG), middle (MFG), and inferior (IFG) frontal gyri, cingulate gyrus (CG), intraparietal sulcus (IPS), and fusiform gyrus (FFG). In addition, white matter control regions were determined (WHM).

IFG, CG, and IPS. In marked contrast to MFG, there was no evidence for delay-spanning activity in the FFG, where the BOLD signal returned to prestimulus baseline levels by 15 sec.

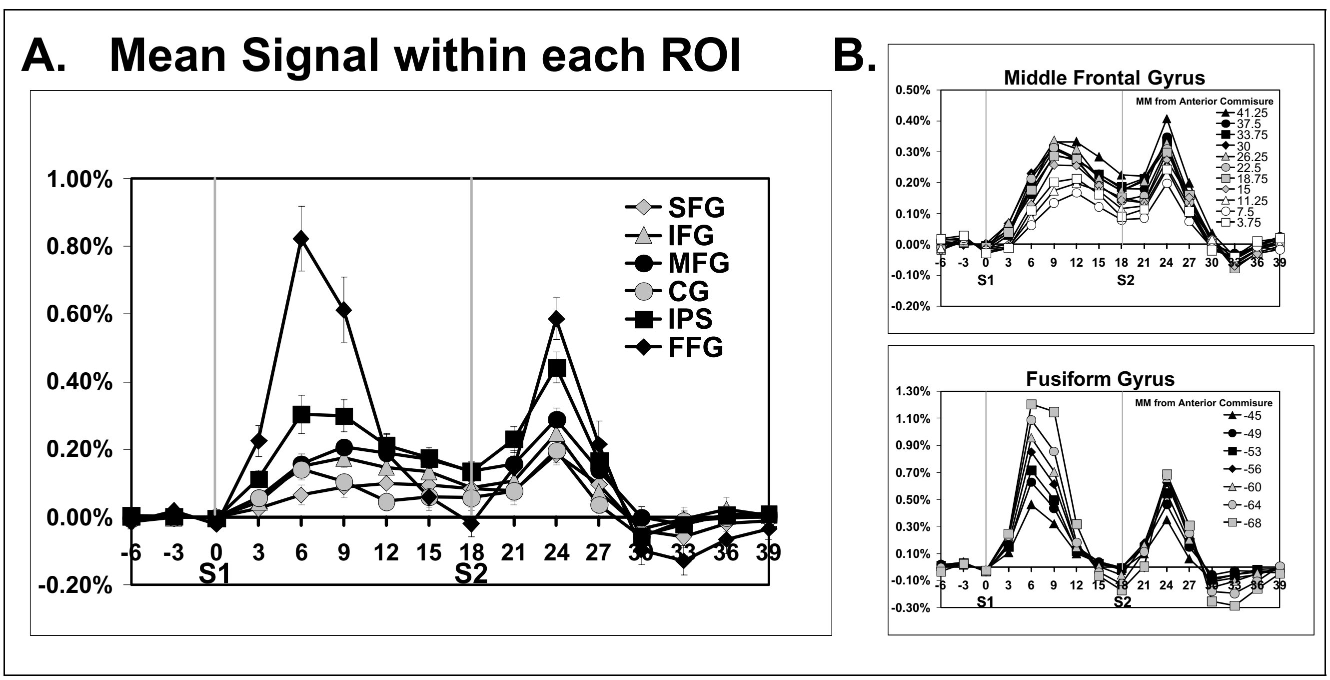

Figure 3. Mean signal within each ROI. (A) These waveforms represent the contribution of all voxels within each ROI and were thus not influenced by any statistical criterion or correlational analysis. In Experiment 1, significant stimulus-evoked hemodynamic responses were observed following S1 and S2 onset. A reduced level of activity occurred in the S1–S2 delay interval that varied considerably across the different ROIs. (B) Activity within MFG and FFG on a slice-by-slice basis (across all load conditions) is plotted to demonstrate that three waveform features were highly correlated across slices. The slice with the largest amplitude response to S1 also had the largest amplitude response to S2 and the greatest sustained activity.

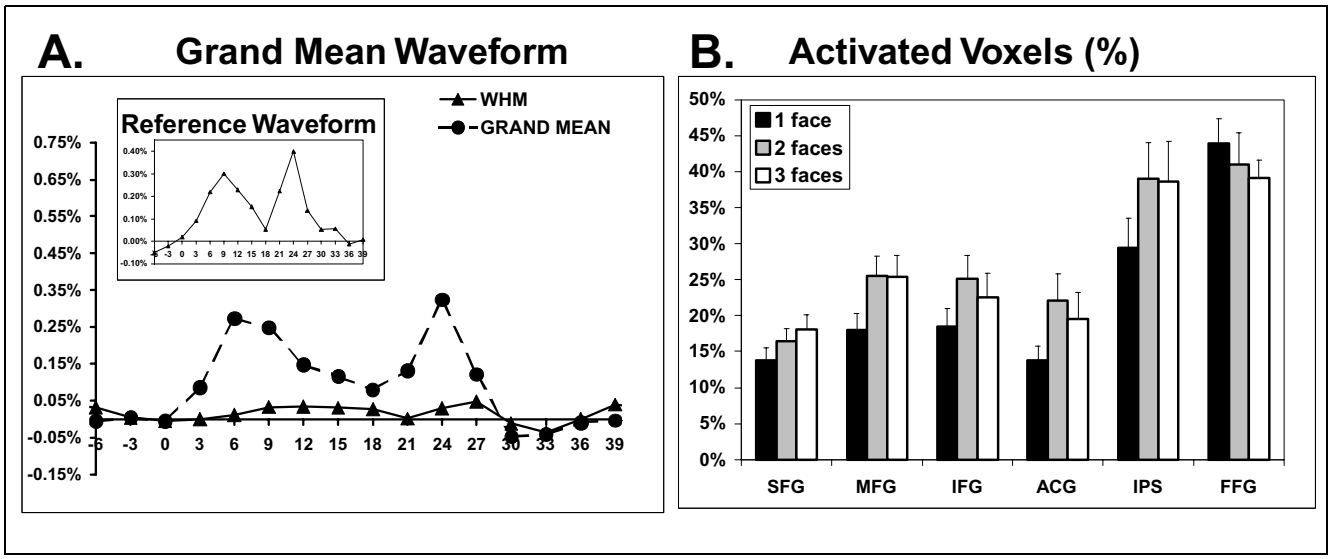

Figure 4. (A) Depicts the grand mean waveform from all voxels within all anatomical ROIs (dotted line) and within the white matter control regions (WHM — solid line) from Experiment 1. The inset depicts the reference waveform used for the correlation analysis. (B) Depicts the number of voxels whose correlation with the reference waveform was significant ($t > 1.96$, $p < .05$, uncorrected) at each load condition.

Despite the presence of significant delay-spanning activity within the prefrontal kois, ANOVA revealed no significant load-related differences in late delay activity (defined as 15–18 sec post-S1 onset, where the probe was delivered at 18 sec) as a function of memory load for any ROIs. Within MFG, for example, both time points 15 sec ($F(2, 9) = 1.7$, $p = .21$) and 18 sec ($F(2, 9) = .44, p = .65$) post-S1 onset revealed no significant differences as a function of load.

Correlation Analysis with Reference Waveform

As stated in the Methods, the average time courses shown in Figure 3 represent the contribution of all voxels within the ROI, many of which may not have been activated by the task. Thus, small load-dependent differences in the late delay interval in the subset of active voxels may have been obscured. As our anatomical ROI analyses revealed little variation in mean waveshape across slice and hemisphere (and, indeed, across ROIs), and all waveform features were highly correlated across slices and hemispheres, we conducted a correlation analysis with an empirically derived reference waveform that, while obtained from a separate pilot study, closely resembled the grand mean waveform of the present study (see Figure 4A and Methods). Activated voxels within each ROI were identified by their correlation with the reference waveform.

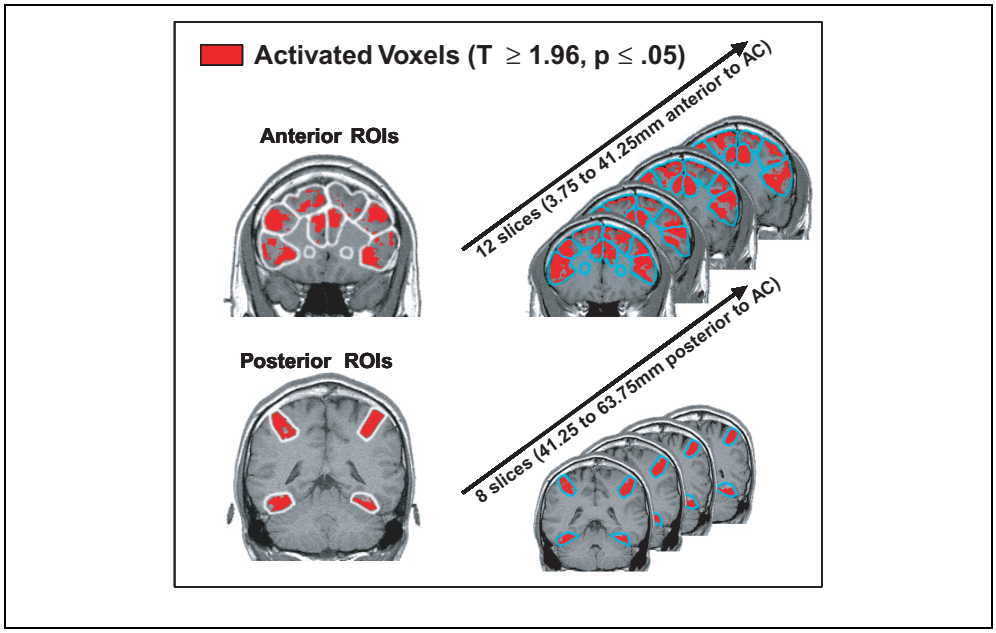

Figure 5. Activated voxels defined as those whose correlation with the reference waveform was significant ($t > 1.96$, $p < .05$, uncorrected) are shown in red for a single subject (from Experiment 1) superimposed upon that subject's high-resolution structural images. The outline of the anatomical ROIs can also be seen superimposed upon the anatomical images. These ROIs were used to group activated voxels by region.

By using a low statistical criterion for the significance of the correlation ($t > 1.96$, $p < .05$, uncorrected), we purposively defined as activated those voxels whose actual waveshapes could vary considerably from the reference waveshape.

Figure 4B presents the mean percentage of activated voxels (relative to the total number of voxels) in each ROI as a function of memory load. The overall percentage of activated voxels was surprisingly high varying from a low of approximately 16% for the SFG and a high of approximately 40% in the FFG. We examined the spatial distribution of the activated voxels within each ROI by superimposing them upon anatomical images. Activated voxels were not randomly distributed but rather clustered within each ROI as can be seen in the representative subject shown in Figure 5.

A greater number of activated voxels were obtained during three-face and two-face trials compared to one-face trials within MFG ($F(1, 9) = 10.92$, $p < .01$), IFG ($F(1, 9) = 7.10$, $p < .01$), CG ($F(1, 9) = 7.29$, $p < .03$), and IPS ($F(1, 9) = 20.42$, $p < .003$). No significant effect of memory load condition was obtained for the voxel counts within SFG and FFG. While instructive, these voxel count analyses do not reveal which waveshape components were sensitive to memory load manipulations. This issue is taken up below for each experiment.

Experiment 1

To determine which activation waveshape components were sensitive to the memory load manipulation, we performed a second activation time course analysis—but only included the contribution of activated voxels (defined above) within each ROI. The resulting waveshapes were similar to those calculated for unselected voxels; however, the amplitudes were approximately twice as large. Load sensitivity was determined for S1-evoked responses, delay-period activity, and S2-evoked activity.

The mean activation time courses for all ROIs are depicted in Figures 6 and 7. Load sensitivity was assessed by testing the influence of load upon the mean ampli-

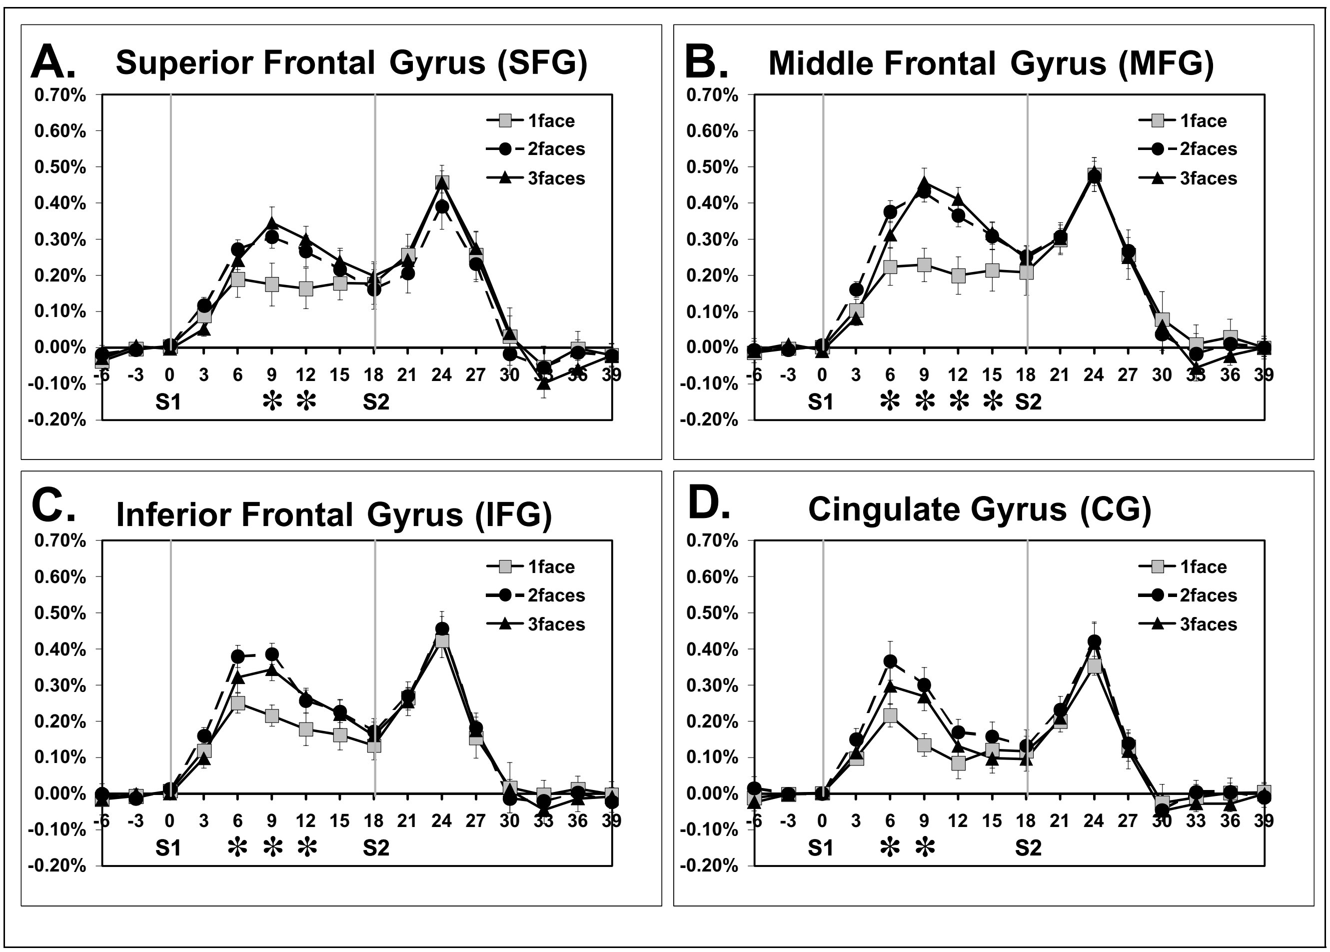

Figure 6. Mean activity within activated voxels is plotted for (A) superior frontal gyrus (SFG), (B) middle frontal gyrus (MFG), (C) inferior frontal gyrus (IFG), and (D) cingulate gyrus (CG) as a function of memory load for Experiment 1. Each time point at which mean signal values were significantly greater for three- and two-face compared to one-face trials is marked with an asterisk.

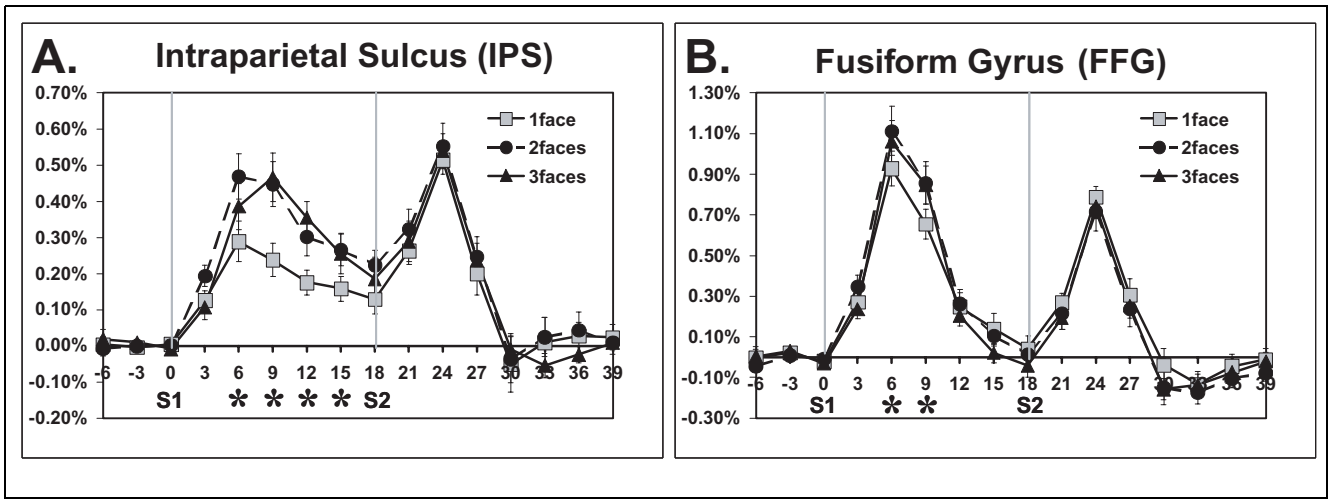

Figure 7. Mean activity within activated voxels is plotted for (A) intraparietal sulcus (IPS) and (B) fusiform gyrus (FFG) as a function of memory load for Experiment 1. Each time point at which mean signal values were significantly greater for three- and two-face compared to one-face trials is marked with an asterisk.

tude values at each of the nine time points following S1 onset (i.e., 6, 9, 12, 15, 18, 21, 24, 27, and 30 sec). The significance levels (F and p values) are summarized in Table 1. Amplitude differences were classified as significant for p values $\leq .05$. At 6 sec, mean signal values were significantly greater for three-face and two-face compared to one-face trials within MFG, IFG, CG, IPS, and FFG. At 9 sec, mean signal values were significantly greater for three-face and two-face compared to one-face trials within SFG, MFG, IFG, CG, IPS, and FFG. At 12 sec, mean signal values were significantly greater for three-face and two-face compared to one-face trials within SFG, MFG, IFG, and IPS. At 15 sec, mean signal values were significantly greater for three-face and two-face compared to one-face trials for MFG and IPS. At 18 sec, the end of the delay interval, there was no significant load effect for any ROI. To determine if responses to S2 were load sensitive, time points 21, 24, 27, and 30

Table 1. F and p Values for the 6-, 9-, 12-, 15-, and 18-sec Time Points Following S1 Onset for ROIs That Were Significantly Active ($p \leq .05$; Experiment 1)

Experiment 1											
		6		9		12		15		18	
Test	*ROI*	*F*	*p*	*F*	*p*	*F*	*p*	*F*	*p*	*F*	*p*
Load (*df*=2, 18)	CG	7.41	.005	10.67	.001						
	FFG	8.13	.003	12.06	.001						
	IFG	6.50	.008	19.29	.000	5.15	.017				
	IPS	12.52	.000	13.67	.000	8.99	.002	3.63	.047		
	MFG	4.77	.022	16.75	.000	11.24	.001				
	SFG			6.79	.006	3.67	.046				
Load by Hemisphere (*df*=2, 18)	IPS					4.05	.035				
	MFG			4.00	.037	5.50	.014				
	SFG			5.62	.013						
Load 1 versus 3 (*df*=1, 9)	CG			12.37	.007						
	FFG	10.78	.010	31.64	.000						
	IFG			17.04	.003	6.96	.027				
	IPS			19.25	.002	15.35	.004				
	MFG			19.21	.002	16.51	.003	6.78	.029		
	SFG			7.72	.021						
Load 2 versus 3 (*df*=1, 9)	IPS	6.61	.030								
Hemisphere (*df*=1, 9)	MFG			6.15	.035	10.76	.010				

Those ROIs that were not active for a given test at any time point are not displayed.

sec following S1 onset were analyzed. Although a large transient activation was observed following S2 for all ROIs, this activation was not influenced by load.

The strength of the memory load manipulation was most apparent early in the delay interval (6 to 12 sec following S1 onset) and declined thereafter. Figure 8 depicts this pattern. Here, the activation time course for the one-face condition has been subtracted from the three-face and two-face time courses. The peak difference occurs at approximately 9 sec after the onset of S1 and declines to near-zero by the end of the delay interval.

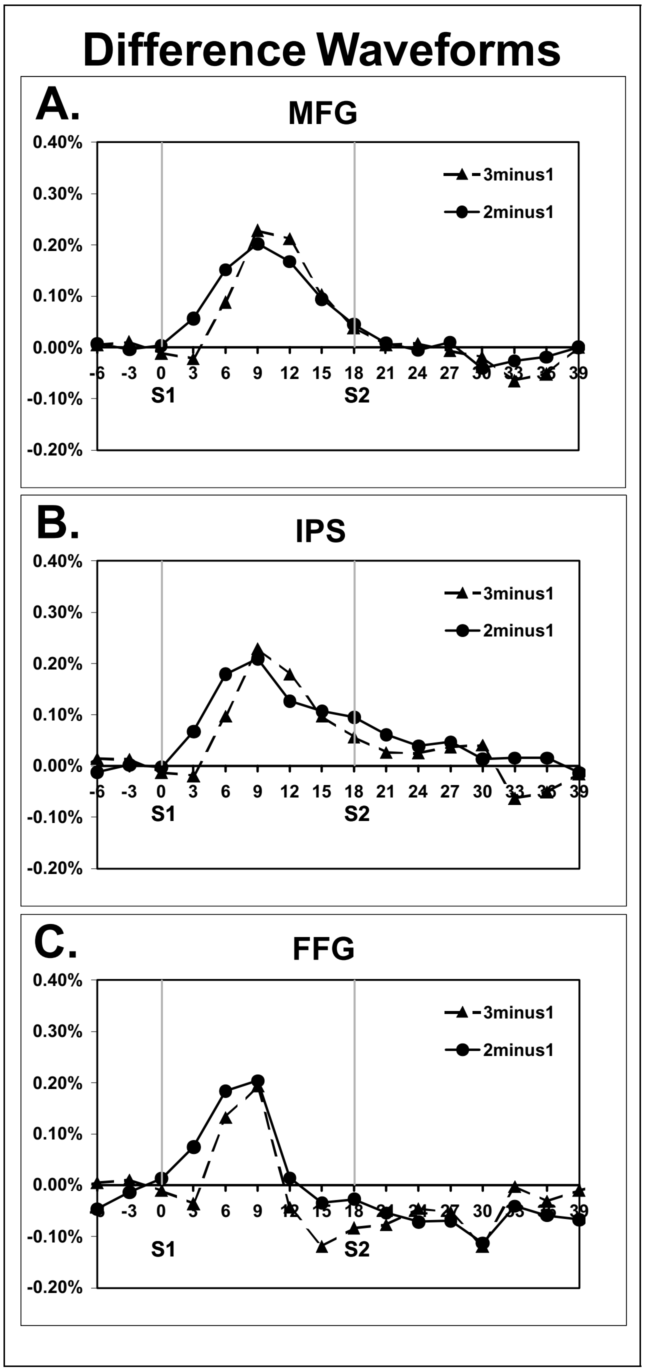

Figure 8. Activation time courses for the one-face condition subtracted from the three-face and two-face time courses are depicted for MFG (A), IPS (B), and FFG (C) from Experiment 1. The peak difference occurred ~6–9 sec after S1 onset and declined to near-zero by the end of the delay interval.

In addition to analyses performed on mean amplitude measures, ANOVAs were performed on peak amplitude and latency measures to determine if the latency of peak responses shifted with memory load. Within MFG and IPS, higher loads resulted in longer latencies to peak amplitude within the HDR following S1. Within MFG, this was revealed as a main effect of load on the latency of peak amplitude responses ($F(2, 9) > 6.00, p < .01$). S1 peak amplitudes for stimuli comprised of one-face occurred earlier (8 sec) than two-face arrays (9 sec), which were earlier than three-face arrays (10 sec). Within IPS, the latency of peak amplitude responses to S1 occurred earlier in the two-face (8 sec) relative to three-face arrays (9 sec) ($F(1, 9) = 22.70, p < .001$). Peak amplitude effects in the S1 time range were similar to mean amplitude effects of load. No latency or amplitude load effects were observed for S2 responses.

Experiment 2

Although there were no significant load-sensitive differences in delay activity at the conclusion of the delay interval (i.e., 18 sec post-S1) in Experiment 1, Figure 8 shows that the difference waveforms had not quite reached zero at 18 sec for either MFG or IPS. To determine whether this nonzero activity might reflect a small load-sensitive maintenance effect missed due to a Type II error, we replicated the experiment. The shape of the difference waveforms suggested a rise and decay of a phasic response evoked by the memory load at S1. To better separate sustained activity from the decaying phasic response, we increased the delay interval to 24 sec (measured from S1 offset to S2 onset). Only the one-face and three-face conditions were run. Figure 9 shows the time-activation waveforms for the active voxels within the MFG and the IPS. These waveforms were computed in the same manner as those shown in Figures 6 and 7, with the exception that the reference waveform used in the correlation analysis was appropriate for the longer delay interval.

As before, the memory load manipulation produced a strong effect upon activity early in the delay interval with the peak difference occurring at ~9 sec post-S1 onset. For MFG, no significant effect of the load manipulation occurred at 15 sec post-S1 onset or at any later time point. It is notable that a plateau of sustained activity was reached at about 21 sec post-S1, but this activity was essentially equivalent for both memory load conditions. A similar result was obtained for IPS; again, there were no significant differences due to load in the late delay interval, and here, the waveforms returned to prestimulus baseline levels by the end of the interval. In this experiment, as in the preceding, there was no load-

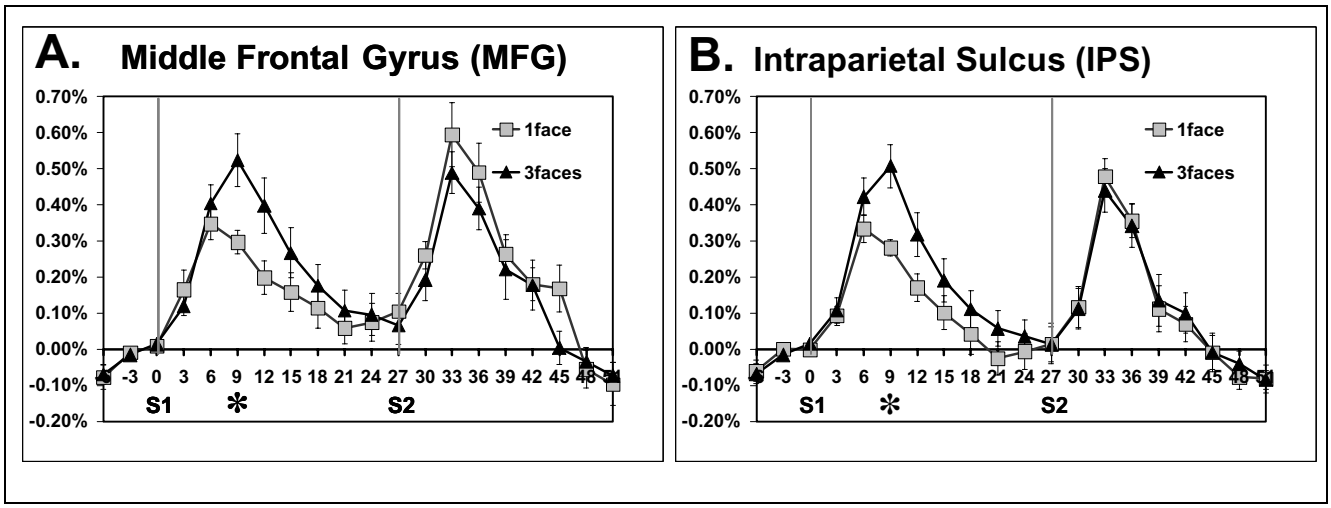

Figure 9. Mean activity within activated voxels is plotted for (A) MFG and (B) IPS as a function of memory load for Experiment 2. Each time point at which mean signal values were significantly greater for three-face compared to one-face trials is marked with an asterisk.

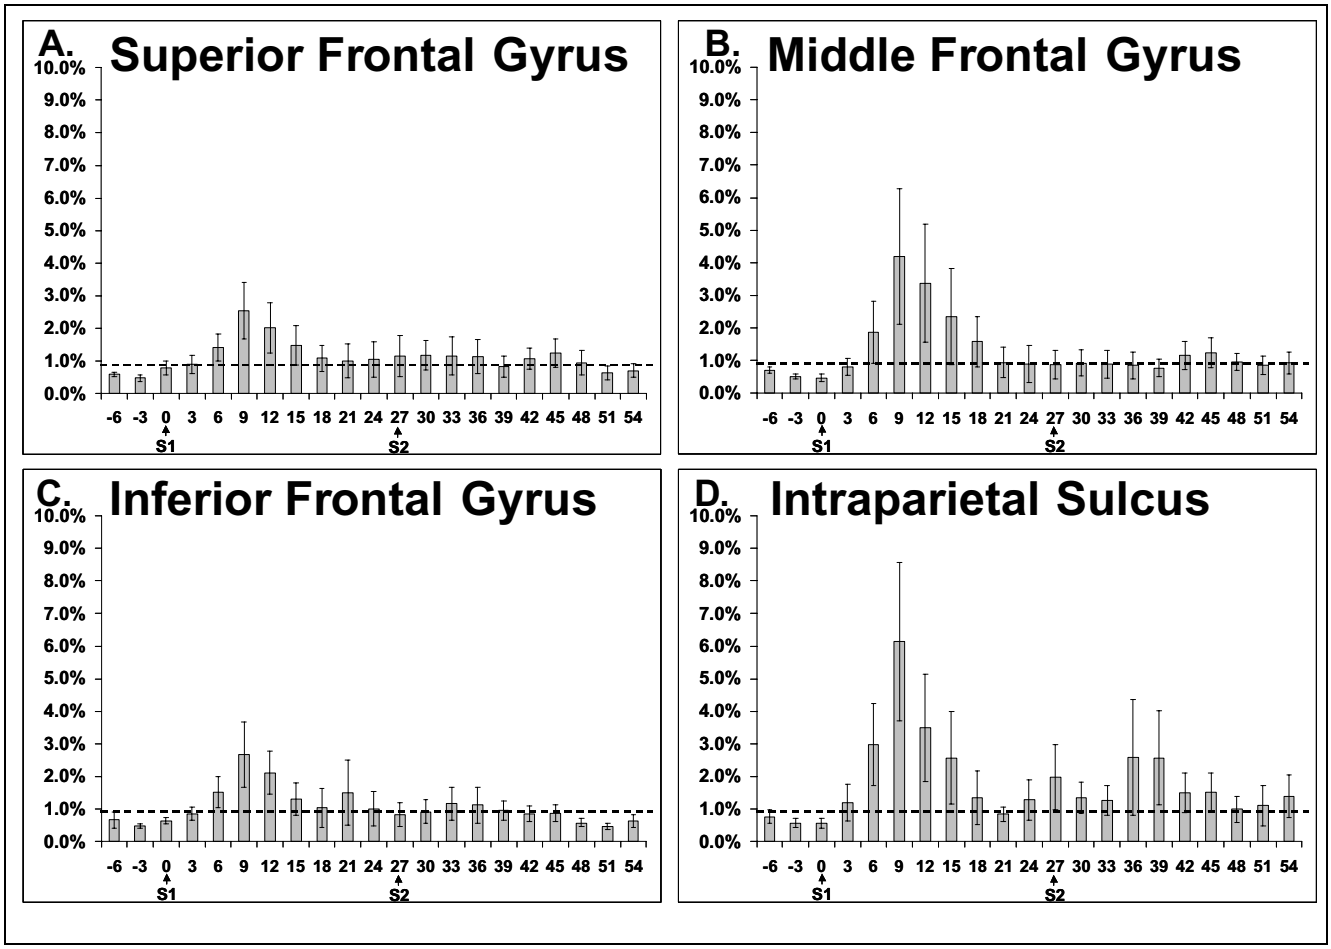

Figure 10. A *t* statistic was computed for each voxel, comparing the one- and three-face load conditions from Experiment 2. This analysis was performed for each time point and computed across the single trial epochs for each subject. The number of voxels at each time point whose *t* value exceeded 1.96 ($p < .05$, uncorrected) was counted for each subject and grouped by anatomical ROI. Percentages of voxels (those exceeding the 1.96 *t* threshold/total number of voxels within the ROI) were plotted for the average across all subjects for SFG (A), MFG (B), IFG (C), and IPS (D). The *t* statistic map was interpolated to the resolution of the high-resolution anatomical images used for the ROI definition. This interpolation influenced the false positive rate. We therefore used an empirically derived value from white matter control regions to determine false positive rates. This value was 0.89% and is indicated by the dashed black line.

sensitive delay activity within the FFG as activity returned to prestimulus baseline levels for the final 15 sec of the delay period.

As in Experiment 1, peak amplitude and latency analyses revealed later latency peaks with increasing load within MFG ($F(1, 8) > 8.62, p < .02$), IPS ($F(1, 8) = 12.9, p < .007$), and IFG ($F(1, 8) = 11.93, p < .008$).

To further investigate whether a subset of voxels within the MFG, IPS, and other regions may have been significantly influenced by memory load conditions in the late delay activity, a t statistic was computed comparing the three- and one-face load conditions for each time point within the delay epoch (see Methods) for all voxels, regardless of whether they were significantly correlated with the reference waveform. Figure 10 presents the percentage of voxels having 'significant' t values ($t > 1.96$) for each time point and each ROI. A longer epoch is presented here rather than in other figures so that the entire intertrial interval (ITI) is shown. Consistent with our other analyses, the MFG was associated with the largest percentage of voxels differentiating the load conditions of all frontal ROIs. This effect peaked at 9 sec post-S1 and declined thereafter. By 21 sec, the number of significant activations was similar to that obtained in the ITI and likely reflective of chance. A similar pattern was obtained for other frontal ROIs, although the percentage of load-sensitive voxels was considerably smaller. The largest number of voxels differentiating the memory loads was obtained for the IPS at 9 sec post-S1—again consistent with our prior analyses. Although the temporal pattern of IPS activation is more variable than that of the frontal ROIs, there is no indication of load-sensitive late delay activity; indeed, more activity is noted in the ITI than in the 9 sec preceding the S2 probe.

DISCUSSION

The present study varied the number of faces to be remembered in a memory array and used a long delay period to temporally segregate the influence of this memory load manipulation into phasic and sustained activation. We reasoned that maintenance processes should be identified by their sensitivity to memory load over the full duration of the delay interval, while encoding effects should be most apparent in the initial interval following the onset of the memory array. Our results showed that memory load significantly influenced activation in the early portion of the delay interval. Furthermore, this difference in activation related to memory load peaked at 6–9 sec and then declined thereafter. No significant effects due to memory load were evident at the end of the delay interval in Experiment 1 at any ROI. In Experiment 2, where a longer delay interval was used, no significant effects due to memory load were evident over the final 12 sec of the delay interval.

The four ROIs defined within the PFC showed very similar patterns of activation with the MFG and IFG showing a greater number of activated voxels and larger amplitude activation time course waveforms than the SFG or CG. Although there were anterior-posterior gradients in the strength of activation for all of these ROIs, all waveform features were correlated across this gradient; i.e., within each ROI, those slices with large HDRs evoked by S1 were the same slices from which large HDRs were evoked by S2. Finally, the activation time courses were similar in waveshape whether derived from the average of all (unselected) voxels within the ROI, or as the average of activity for a subset of voxels designated as activated by their similarity to an empirically derived reference waveform.

The MFG showed the greatest number of activated voxels and the largest amplitude activation waveforms among frontal ROIs. Increasing memory load significantly increased activity in the post-S1 interval, consistent with recent reports of load-sensitive S1 responses within PFC (Rypma & D'Esposito, 1999). Moreover, the latency to peak activation here was significantly affected by memory load with peak activation occurring at 8 sec for the one-face condition and 10 sec for the three-face condition (as measured on activation waveforms fit with cubic spline functions) in both experiments. We were able to account for the form of this latency shift and subsequent slower return to baseline of the HDR by convolving an empirically derived HDR with simulated neuronal response profiles that varied in duration with increasing load. Thus, this latency shift, which was observed at the MFG and IPS in both experiments, suggests that the duration of processes engendered by the appearance of the memory array may have been influenced by the number of items to be remembered, such as the time required to encode, or recode, the memoranda. It is also possible that this activity reflects other executive functions—such as those involved in switching strategies for encoding the memoranda when there is one face (perhaps a holistic strategy) and more than one face (perhaps a feature-based strategy). The use of faces as memory items, while useful for our interest in fusiform activity during delay intervals, limits the present study in that it is difficult to verify that subjects use a consistent encoding strategy across all memory loads.

This load-dependent difference declined over successive time intervals and was no longer significant at the end of the delay interval. Indeed, in Experiment 2, the activation waveforms for the two load conditions were superimposed over the final 9 sec of the delay interval. It is notable that the MFG (and other PFC sites) showed clear sustained activity that spanned the entire delay interval. However, this activity was not influenced by memory load and, for this reason, we cannot conclude that this delay-spanning activity was related to maintenance in favor of any other process that may be operative in a delay interval, such as response preparation, mental timing, or expectancy. This conclusion is at odds

with that of Courtney et al. (1997) who reported maintenance related sustained activity within MFG. However, we note that the 8-sec delay interval used by Courtney et al. (1994) ends at about the same point in time where the early and, we assert, phasic activity evoked by S1 reached its peak in the present study. Our conclusions also differ from those studies of memory maintenance based upon the *n*-back task (Braver et al., 1997; Cohen et al., 1997). As discussed above, arguments for maintenance effects in PFC based upon *n*-back tasks are confounded by the simultaneous requirement of other executive and response selection processes.

A dorsal–ventral PFC distinction has been proposed with information manipulation primarily a function of dorsal PFC (D'Esposito et al., 1999; Postle et al., 1999) and maintenance more a function of more ventral PFC regions (Owen et al., 1996). With the exception of somewhat larger amplitudes measured at MFG, there was a striking similarity of effects at MFG and IFG and thus our data provide no support for this distinction. An alternative functional differentiation within PFC has recently been proposed by Rowe et al. (2000). Using a delayed-response task with an 18-sec delay, Rowe et al. (2000) proposed that BA 8—a cytoarchitectonic region that includes much of the posterior aspect of our anatomically defined SFG—supports memory maintenance during the delay interval. BA 46—a region that includes much of the MFG—was proposed to be mainly engaged by response selection during S2 presentation. Our data also fail to support this distinction. The pattern of load sensitive activation within SFG was similar, although somewhat smaller in amplitude to that obtained in MFG and sustained activity within SFG, like the MFG, was not load sensitive in the late delay interval. At all ROIs, S2 evoked large activation; indeed, S2 activation was greater in amplitude than S1 activation at all PFC ROIs, and there was no difference in S2 amplitude between SFG and MFG. In addition, S2 activation was not influenced by memory load at any ROI as one might expect if the choice of a response required a comparison of the probe to all elements within the memory buffer.

Sustained activation was also noted in the IPS, a region whose activation time course was similar to that obtained in MFG but larger in amplitude. In Experiment 1, there was no statistical difference related to memory load at the final time point of the delay interval. There was, however, a small increase in the level of sustained activity for the higher memory load conditions. We reasoned that this small difference might be maintained over a longer delay interval, and therefore conducted a second experiment with a longer interval. In this second study, the delay activity within the IPS actually returned to prestimulus baseline levels prior to the appearance of the probe stimulus for both the one-face and three-face load conditions. We therefore concluded that the sustained activity observed did not reflect memory maintenance.

In both experiments reported here, no sustained activity was observed within FFG in the late delay interval. This is contrary to previous fMRI studies that reported significant delay-period activity (Courtney et al., 1997) in this region, and to more recent studies that reported memory load sensitive delay activity within posterior perisylvian regions (Postle et al., 1999; Rypma & D'Esposito, 1999). In all of these studies, the delay interval was short and the phasic activity evoked by S1 may not have returned to baseline. We argue that the measurement of delay activity was confounded by load effects upon phasic activity evoked by S1. In the current studies, phasic S1 activity within FFG was load sensitive and more temporally restricted than the other ROIs investigated.

One issue that complicates our interpretation concerns the nature of the progressive decline in sustained activity over the course of the delay period. Others have shown that BOLD activation does not decline in amplitude in sensory experiments using longer intervals than those used here (Howseman, Porter, Hutton, Josephs, & Turner, 1998). However, it is possible that the decline in the load-sensitive component of the sustained activity represents a functional decline in memory maintenance. We feel that this is unlikely, in that the subjects' performance in all load conditions was high—89% for the three-face load condition in Experiment 1. However, we recognized that the overall performance in Experiment 2 that used a longer delay interval was somewhat poorer than that of Experiment 1 that used a shorter interval. Perhaps the further decline in load-sensitive delay activity in Experiment 2 reflects forgetting.

To test this possibility, we measured PFC activity during the delay period in the four highest performing subjects and five lowest performing subjects from Experiment 2. Delay period amplitude measures did not significantly differ between low and high performers, and neither group had significant late delay-interval activity that was load sensitive. Thus, we concluded that the decline in sustained activity was not related to small differences in memory performance observed. High performers in Experiment 2 performed at a comparable level with subjects in Experiment 1 (95% mean accuracy for one-face trials and 89% for three-face trials). MFG activation for this group is plotted in Figure 11.

To date, all studies that have investigated maintenance processes have used sensitivity to memory load as the experimental evidence. As we failed to find that relationship for delay-spanning-sustained activity, we have argued against an interpretation based upon maintenance. However, it is possible that the delay-spanning-sustained activity we did observe was indeed related to memory maintenance, but that the neural effort required by maintenance processes did not, and does not, scale with the number of items maintained as assumed by our design. In this case, the challenge is to distinguish sustained activity related to maintenance

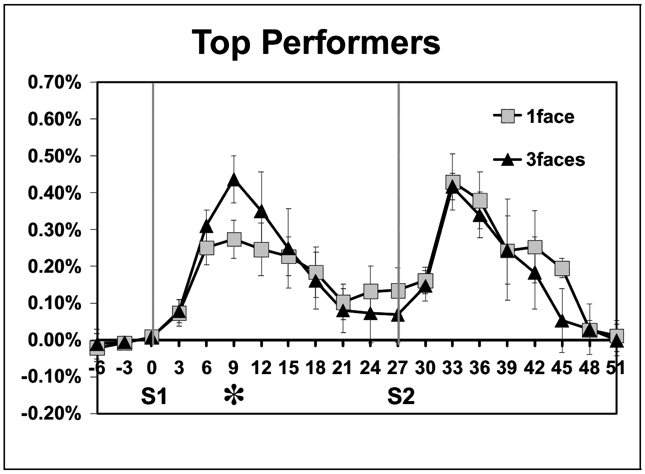

Figure 11. Activity within activated MFG voxels for the top four performers from Experiment 2.

from sustained activity related to any other processes that might occur over the span of a delay interval. For example, in most delayed-response tasks, the subject is waiting to press a button, is timing the likely appearance of S2, and may be actively inhibiting other overt or covert behaviors to reduce interference with the memory task.

Finally, we note with some concern the similarity of activation time profiles over large regions of brain. The improved sensitivity resulting from the use of signal averaging may help small activations emerge from the noise. Nevertheless, we are concerned that some of these activations may reflect nonlocal effects perhaps related to veins *en passage*. Pulse sequences that reduce the large vessel contribution to BOLD contrast may be necessary to improve functional resolution.

METHODS

Subjects

Ten male volunteers participated in Experiment 1 (age range 19–36 years, mean 26 years) and two female and eight male subjects participated in Experiment 2 (age range 20–29, mean 22.5 years). In each experiment, one subject was left-handed. One subject was excluded from Experiment 2 because of a data acquisition error. The Duke University Medical Center Institutional Review Board approved this study and each subject provided informed consent.

Experimental Tasks

Each trial consisted of a memory array (S1) presented for 3 sec, a delay interval free of all stimuli, and a single probe stimulus (S2) presented for 1.5 sec that terminated the delay interval. On half of the trials, the probe stimulus matched one of the items in the memory array while on the remaining half the probe was a new stimulus. The subjects made a choice-button response with the index finger of their left or right hand to indicate whether the probe did, or did not, match one of the memory items.

The memory array consisted of nine stimulus items arranged in three rows by three columns (see Figure 2A). In Experiment 1, each trial consisted of one, two, or three unfamiliar and moderately confusable faces. The faces were always presented in the middle row and were flanked by patterns that matched the faces in spatial frequency and overall luminance. Each pattern was created by computing a two-dimensional Fourier transform of a face, randomly permuting the phase spectrum, and then computing the inverse transform. The flanking patterns were never task relevant and were there simply to equate overall stimulus properties despite changes in number of faces. The delay interval, measured from the offset of S1 to the onset of S2, was 15 sec. The ITI, measured from the offset of S2 to the onset of S1 for the following trial, was 15 sec. All subjects performed nine runs with each run consisting of nine trials. In Experiment 2, stimulus presentation and timing were identical to Experiment 1 with the following exceptions: only one- and three-face trials were included, the (offset to onset) delay period was 24 sec, and the ITI was 18 sec. Seven subjects completed six runs with each run consisting of 10 trials. One subject completed four runs and another subject completed three runs.

All stimuli were presented using a LCD projector (XGA resolution, 900 lumens) equipped with a specially designed Buhl lens. Stimuli were projected upon a 10-in.-wide screen located within the magnet bore directly behind the subject's head and subjects viewed the stimuli through mirrored glasses. Behavioral responses were acquired using a button box incorporating a fiber-optic loop connected to a TTL driver circuit located outside of the magnet room. RTs were measured by the experimental control software.

Behavioral Data Analysis

Behavioral results were obtained from seven subjects in Experiment 1 (equipment failure precluded behavioral data collection in three subjects) and from nine subjects in Experiment 2. For each experiment, RT and percent correct measures were entered into separate repeated-measures ANOVA tests to determine the influence of memory load (one, two, and three faces in Experiment 1; one and three faces in Experiment 2). A second mixed-design analysis was performed to determine if RT and percent correct scores differed between Experiments 1 and 2 for the one- and three-face conditions.

MRI Acquisition

All scanning was performed on a General Electric 1.5 T NVi scanner equipped with an Advanced Development Workstation for real time echoplanar imaging. T1-

weighted sagittal localizer images were first acquired. The AC and posterior commissure (PC) were identified in the mid-sagittal slice and 34 (near axial) contiguous oblique slices were prescribed parallel to the AC–PC plane. High-resolution T1-weighted structural images were acquired (TR = 450 msec, TE = 20 msec, FOV = 24 cm, 256 × 256 matrix, slice thickness = 3.75 mm). A second series of (near-coronal) oblique T1-weighted images perpendicular to the AC–PC were then acquired using the same imaging parameters. Thirty-four contiguous echoplanar images (EPIs) sensitive to BOLD contrast were acquired parallel to the AC–PC plane using the same slice prescription described above for the near-axial structural images (TE = 40 msec, 24 cm FOV, 64 × 64 image matrix, 90° flip angle, TR = 3 sec, slice thickness = 3.75 mm resulting in 3.75 mm^3 isotropic voxels). The resulting EPI volumes could then be viewed with the same uninterpolated resolution in coronal, axial, and sagittal orientations. Images were superimposed on a slice-by-slice basis upon both the near-axial and near-coronal high-resolution anatomical images.

fMRI Data Analysis

The center of mass for each EPI image volume within each time series was computed and plotted to determine head motion during scanning. In both experiments, no subject was observed to have greater than a 1-mm deviation in the center of mass. The time series of EPI volumes was adjusted to compensate for the interleaved slice acquisition order within each TR interval. This was accomplished by fitting the time series of each voxel with a cubic spline and then resampling this function for all voxels at the onset of each TR interval.

Average time-epochs synchronized to S1 were then computed for each memory load. Following McCarthy, Luby, Gore, and Goldman-Rakic (1997), image time segments consisting of the five image volumes preceding and 10 (Experiment 1) or 14 volumes (Experiment 2) following each S1 stimulus were excised from the continuous time series of volumes for each run comprising each subject's data. These single trial volume epochs were then grouped by stimulus type (one-, two-, or three-face trials) and averaged with the temporal order of each volume relative to S1 carefully maintained. The average MR signal values were then converted to percent signal change relative to a pre-S1 baseline. The baseline was calculated as the mean intensity of the two volumes (6 sec) prior to S1. The averaged time-epochs were then subjected to the analytic procedures described below.

Regions of Interest (ROIs)

Anatomical ROIs were then drawn on each subject's high-resolution coronal anatomical images (Figure 2B) by a trained observer who was blind to other statistical analysis of the data. Each subject contributed 140 ROIs. ROIs were drawn on left and right superior frontal gyri (SFG), middle-frontal gyri (MFG), inferior frontal gyri (IFG), cingulate gyri (CG), IPS, FFG, and white matter control regions (WHM) on each slice on which the structure appeared. ROI slice labels indicated the distance in millimeters from each subject's AC so that summary activation waveforms could be created across subjects. ROIs for SFG, MFG, IFG, CG, and WHM were drawn on 11 slices ranging form 41.25 to 3.75 mm anterior to the AC. IPS and FFG were drawn on seven slices ranging from 41.25 to 63.75 mm posterior to the AC.

The ROIs were used for two different analyses. In the first, the mean signal change for all voxels within each ROI was computed for each time point and plotted to visualize the time course of the HDR for each ROI during each memory load condition. This analysis was not dependent upon choice of reference waveform. The HDR time course was separately examined for each slice and hemisphere within each ROI, so that regional and task-related differences in the form of the HDR would not be missed. For example, it is possible that some slices within a ROI would only evince a phasic response to S1 without any activity evoked to S2 or during the delay interval. In the second analysis, the ROIs were used to group and count voxels determined to be active by correlation analysis with a reference waveform (see below).

Correlation with Reference Waveform

While the ROI approach described above is useful to measure the mean activity within an anatomically circumscribed region without recourse to predetermined reference waveforms, it is possible that the majority of voxels within the ROI will not be activated by the task. Thus, the mean activity measured within the ROI may represent a dilution of the activity attributed to a subset of voxels. A correlation analysis was therefore conducted in an attempt to identify this subset of activated voxels. The reference waveform was derived empirically from a pilot study using a similar design (see inset of Figure 4A) (Jha & McCarthy, 1999). Active voxels were defined as those that significantly correlated with the reference waveform ($t > 1.96$, $p < .05$, uncorrected). Using the ROIs defined above, we tabulated the number of active voxels within each anatomical ROI for each load condition. These counts were converted to percentages relative to the number of voxels in each ROI and were analyzed by repeated-measures ANOVA to determine if the number of activated voxels differed across gyrus, slice, hemisphere, and memory load condition.

Time-Activation Waveforms for Activated Voxels

Time-activation waveforms were computed for the activated voxels located within each anatomical ROI. As

there were differences in the number of activated voxels as a function of memory load, the time waveforms were computed for the union of active voxels for all load conditions for each ROI (collapsing across slice). Thus, each subject contributed 12 waveforms (six anatomical regions by two hemispheres) for each memory load condition. ANOVAs were performed to determine if activity differed as a function of memory load at each time point in the activation time-waveforms for each ROI as computed above.

We were concerned that the analyses described above would not appropriately capture apparent differences in the latency to peak amplitude of the HDRs evoked by the S1 as a function of memory load. A further measure of peak amplitude and latency was therefore made of the HDR evoked by S1 (0–12-sec window) and S2 (18–30-sec window in Experiment 1 and 27–42-sec window in Experiment 2). To compute this analysis, time-activation waveforms from each subject, hemisphere, ROI, and memory load were first fit with a cubic spline that was then sampled at 1-sec resolution. The peak amplitudes and latencies were then measured within the windows described above. Repeated-measures ANOVAs were performed to determine if the amplitude or latency varied for ROI, hemisphere, and load condition for each experiment.

Temporal Profile of t Statistics

It is possible that our choice of reference waveform would bias against the discovery of voxels that significantly differed in their response to the load manipulation, but had time-activation profiles that differed significantly from the reference waveform, and were too few in number to be identified by the analysis of mean activity from the anatomical ROI analysis. We therefore computed a t statistic for each voxel, comparing the one- and three-face load conditions from Experiment 2. This analysis was performed for each time point and computed across the single trial epochs for each subject. The number of voxels at each time point whose t value exceeded 1.96 ($p < .05$, uncorrected) was counted for each subject and grouped by anatomical ROI. It is important to note that a voxel identified as significant at any individual time point need not be the same voxel identified as significant at another time point.

Acknowledgments

We wish to thank Martin McKeown, Teresa Mitchell, and Daniel Weissman for their comments on an earlier version of this manuscript. In addition, we wish to thank Brian Miller, Timothy Hingston, Atheen Venketaramani, and Jeremy Goldstein for their assistance in data acquisition and analysis. This research was supported by NIMH grant MH-05286, the Department of Veterans Affairs, and the McDonnell-Pew Foundation.

Reprint requests should be sent to Amishi P. Jha, Box 3918, Duke University Medical Center, Durham, NC 27710; email: amishi.jha@duke.edu.

This data reported in this experiment have been deposited in National fMRI Data Center (http://www.fmridc.org). The accession number is 2-2000-1115T.

REFERENCES

Awh, E., Jonides, J., Smith, E. E., Buxton, R. B., Frank, L. R., Love, T., Wong, E. C., & Gmeindl, L. (1999). Rehearsal in spatial working memory: Evidence from neuroimaging. *Psychological Science, 10,* 433–437.

Baddeley, A. (1996). The fractionation of working memory. *Proceedings of the National Academy of Sciences U.S.A., 93,* 13468–13472.

Belger, A., Puce, A., Krystal, J. H., Gore, J. C., Goldman-Rakic, P., & McCarthy, G. (1998). Dissociation of mnemonic and perceptual processes during spatial and nonspatial working memory using fMRI. *Human Brain Mapping, 6,* 14–32.

Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. *Neuroimage, 5,* 49–62.

Callicott, J. H., Mattay, V. S., Bertolino, A., Finn, K., Coppola, R., Frank, J., Goldberg, T. E., & Weinberger, D. H. (1999). Physiological characteristics of capacity constraints in working memory as revealed by functional MRI. *Cerebral Cortex, 9,* 20–26.

Chafee, M. V., & Goldman-Rakic, P. S. (2000). Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. *Journal of Neurophysiology, 83,* 1550–1566.

Cohen, J. D., Forman, S. D., Braver, T. S., Casey, B. J., Servan-Schreiber, D., & Noll, D. C. (1994). Activation of prefrontal cortex in a non-spatial working memory task with functional MRI. *Human Brain Mapping, 1,* 293–304.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature, 386,* 604–608.

Constantinidis, C., & Steinmetz, M. (1996). Neural activity in posterior parietal area 7a during the delay periods of a spatial memory task. *Journal of Neurophysiology, 76,* 1352–1355.

Courtney, S. M., Ungerleider, L. G., Keil, K., & Haxby, J. V. (1997). Transient and sustained activity in a distributed neural system for human working memory. *Nature, 386,* 608–611.

Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences U.S.A., 93,* 13494–13499.

Desmond, J. E., Gabrieli, J. D. E., Wagner, A. D., Ginier, B. L., & Glover, G. H. (1997). Lobular patterns of cerebellar activation of verbal working-memory and finger-tapping tasks as revealed by functional MRI. *Journal of Neuroscience, 17,* 9675–9685.

D'Esposito, M., Detre, J. A., Alsop, D. C., Shin, R. K., Atlas, S., & Grossman, M. (1995). The neural basis of the central executive system of working memory. *Nature, 378,* 279–281.

D'Esposito, M., Postle, B. R., Ballard, D., & Lease, J. (1999). Maintenance versus manipulation of information held in working memory: An event-related fMRI study. *Brain and Cognition, 41,* 66–86.

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolat-

eral prefrontal cortex. *Journal of Neurophysiology, 61,* 331–349.

Fuster, J. M. (1997). Network memory. *Trends in Neurosciences, 20,* 451–459.

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In F. Plum (Ed.), *Handbook of physiology, the nervous system, higher functions of the brain* (pp. 373–417). Bethesda, MD: American Physiological Society.

Goldman-Rakic, P. S. (1996). The prefrontal landscape: Implications of functional architecture for understanding human mentation and the central executive. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences, 351,* 1445–1453.

Haxby, J. V., Ungerleider, L. G., Horwitz, B., Rapoport, S. I., & Grady, C. L. (1995). Hemispheric differences in neural systems for face working memory: A PET-rCBF study. *Human Brain Mapping, 3,* 68–82.

Howseman, A. M., Porter, D. A., Hutton, C., Josephs, O., & Turner, R. (1998). Blood oxygenation level dependent signal time courses during prolonged visual stimulation. *Magnetic Resonance Imaging, 16,* 1–11.

Huettel, S. A., & McCarthy, G. (2000). Evidence for a refractory period in the hemodynamic response to visual stimuli as measured by MRI. *Neuroimage, 11,* 547–553.

Jackson, S. R., Jackson, G. M., & Rosicky, J. (1995). Are nonrelevant objects represented in working memory? The effect of non-target objects on reach and grasp kinematics. *Experimental Brain Research, 102,* 519–530.

Jha, A. P., & McCarthy, G. (1999). Responses in middle frontal gyrus during working memory for faces—an event-related fMRI study. *Society for Neurosciences Abstracts, 25*, 1551.

Jonides, J., Schumacher, E. H., Smith, E. E., Lauber, E. J., Awh, E., Minoshima, S., & Koeppe, R. A. (1997). Verbal working memory load affects regional brain activation as measured by PET. *Journal of Cognitive Neuroscience, 9,* 462–475.

Manoach, D. S., Schlaug, G., Siewert, B., Darby, D. G., Bly, B. M., Benfield, A., Edelman, R. R., & Warach, S. (1997). Prefrontal cortex fMRI signal changes are correlated with working memory load. *NeuroReport, 8,* 545–549.

McCarthy, G., Blamire, A. M., Puce, A., Nobre, A. C., Bloch, G., Hyder, F., Goldman-Rakic, P., & Shulman, R. G. (1994). Functional magnetic resonance imaging of human prefrontal cortex activation during a spatial working memory task. *Proceedings of the National Academy of Sciences U.S.A., 91,* 8690–8694.

McCarthy, G., Luby, M., Gore, J., & Goldman-Rakic, P. (1997). Infrequent events transiently activate human prefrontal and parietal cortex as measured by functional MRI. *Journal of Neurophysiology, 77,* 1630–1634.

Mellers, J. D., Bullmore, E., Brammer, M., Williams, S. C., Andrew, C., Sachs, N., Andrews, C., Cox, T. S., Simmons, A., & Woodruff, P. (1995). Neural correlates of working memory in a visual letter monitoring task: An fMRI study. *NeuroReport, 7,* 109–112.

Miller, E. K., Li, L., & Desimone, R. (1991). A neural mechanism for working and recognition memory in inferior temporal cortex. *Science, 254,* 1377–1379.

Miller, E. K., Li, L., & Desimone, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *Journal of Neuroscience, 13,* 1460–1478.

Owen, A. M., Evans, A. C., & Petrides, M. (1996). Evidence for a two-stage model of spatial working memory processing within the lateral frontal cortex: A positron emission tomography study. *Cerebral Cortex, 6,* 31–38.

Petrides, M., Alivisatos, B., Meyer, E., & Evans, A. C. (1993). Functional activation of the human frontal cortex during the performance of verbal working memory tasks. *Proceedings of the National Academy of Sciences U.S.A., 90,* 878–882.

Postle, B. R., Berger, J. S., & D'Esposito, M. (1999). Functional neuroanatomical double dissociation of mnemonic and executive control processes contributing to working memory performance. *Proceedings of the National Academy of Sciences U.S.A., 96,* 12959–12964.

Rowe, J. B., Toni, I., Josephs, O., Frackowiak, R. S. J., & Passingham, R. E. (2000). The prefrontal cortex: Response selection or maintenance within working memory? *Science, 288,* 1656–1660.

Rypma, B., & D'Esposito, M. (1999). The roles of prefrontal brain regions in components of working memory: Effects of memory load and individual differences. *Proceedings of the National Academy of Sciences U.S.A., 96,* 6558–6563.

Sammer, G. (1999). Working memory load and EEG-dynamics as revealed by point correlation dimension analysis. *International Journal of Psychophysiology, 34,* 89–101.