

Functional Parcellation of Attentional Control Regions of the Brain

Marty G. Woldorff¹, Chad J. Hazlett¹, Harlan M. Fichtenholtz¹,
Daniel H. Weissman¹, Anders M. Dale², and Allen W. Song¹

Abstract

Recently, a number of investigators have examined the neural loci of psychological processes enabling the control of visual spatial attention using cued-attention paradigms in combination with event-related functional magnetic resonance imaging. Findings from these studies have provided strong evidence for the involvement of a fronto-parietal network in attentional control. In the present study, we build upon this previous work to further investigate these attentional control systems. In particular, we employed additional controls for nonattentional sensory and interpretative aspects of cue processing to determine whether distinct regions in the fronto-parietal network are involved in different aspects of cue processing, such as cue-symbol interpretation and attentional orienting. In addition, we used shorter cue-target

intervals that were closer to those used in the behavioral and event-related potential cueing literatures. Twenty participants performed a cued spatial attention task while brain activity was recorded with functional magnetic resonance imaging. We found functional specialization for different aspects of cue processing in the lateral and medial subregions of the frontal and parietal cortex. In particular, the medial subregions were more specific to the orienting of visual spatial attention, while the lateral subregions were associated with more general aspects of cue processing, such as cue-symbol interpretation. Additional cue-related effects included differential activations in midline frontal regions and pretarget enhancements in the thalamus and early visual cortical areas. ■

INTRODUCTION

A number of neuroimaging studies of attentionally demanding tasks have shown activity in areas of the frontal and parietal cortex, including dorsal midline frontal regions, such as the anterior cingulate (e.g., MacDonald et al., 2000; Milham, Banich, Webb, & Barad, 2001; Hopfinger, Buonocore, & Mangun, 2000; Shulman et al., 1999; Wojciulik & Kanwisher, 1999; d'Esposito et al., 1998; Nobre et al., 1997; Vandenberghe et al., 1997; Corbetta, Miezin, Shulman, & Petersen, 1993; Corbetta, Kincade, Ollinger, McAvoy, & Shulman, 2000; Posner & Petersen, 1990; for a review, see Corbetta and Schulman, 2002). Until recently, however, these findings were based on hemodynamic studies using block design approaches, which precluded the ability to delineate which of the activity observed in these frontal and parietal cortical regions was related to the control of attention, such as processes related to attentional orienting, and which was due to stimulus or target processing and any differential effects of attention on that processing. The separation of the brain activity related to these two sets of processes requires that an approach other than block design be used, such as event-related functional magnetic resonance imaging (fMRI) (d'Espo-

sito, Zarahn, & Aguirre, 1999; Rosen, Buckner, & Dale, 1998).

Several of the more recent studies (e.g., Corbetta et al., 2000; Hopfinger et al., 2000; Shulman et al., 1999) have applied event-related fMRI to cued-attention paradigms to separate cue-induced activity (which would include activity related to attentional orienting) from target-induced activity. To address the severe overlap of the cue and target hemodynamic responses that would occur at the typical stimulus onset asynchronies (SOAs) used in the behavioral (or event-related potential [ERP]) literature, these studies either focused on trials with long cue-target SOAs (~8 sec) (Hopfinger et al., 2000) or have used a combination of special trial types, signal extraction techniques, and moderate SOAs (4–5 sec) (e.g., Ollinger, Corbetta, & Shulman, 2001; Ollinger, Shulman, & Corbetta, 2001; Corbetta et al., 2000; Shulman et al., 1999, 2002). Several of these studies (e.g., Hopfinger et al., 2000; see also Kastner et al., 1999) reported that spatial attention-directing cues triggered enhanced activity in the visual sensory cortices contralateral to the cued direction of attention, prior to the occurrence of the target stimulus. This cue-related activity was proposed to reflect a biasing of sensory cortical activity in favor of the expected target stimulus (see Desimone & Duncan, 1995).

¹Duke University, ²Harvard University

The present experiment builds on the methods of these earlier studies for distinguishing cue- and target-related activity to further determine the contributions of various regions in the fronto-parietal network identified as active during attention-orienting tasks. This was done by extending these approaches in mainly two ways: (1) using SOAs (900–1900 msec) that are closer to those used in behavioral studies and (2) adding control conditions that allow the decomposition of cue-triggered activity into activity more specifically related to orienting and activity related to general aspects of cue processing.

First, regarding the SOAs, long delays between cue and target may make it more likely that subjects engage other cognitive processes, such as working memory, during the delay period. Thus it is possible that, due to the relatively long cue–target SOAs in these several previous studies (e.g., 8 sec in Hopfinger et al., 2000; 4–5 sec in Corbetta et al., 2000; 5–6 sec in Shulman et al., 1999), the broad network of areas observed may have included some regions more related to working memory rather than attentional orienting per se. Indeed, the frontal and parietal areas previously observed to be activated by spatial attentional cues included areas similar to those activated in working memory tasks (e.g., Jha & McCarthy, 2000; Nystrom et al., 2000; LaBar, Gitelman, Parrish, & Mesulam, 1999; d’Esposito et al., 1998). To reduce this possibility, we used a substantially shorter cue–target SOA (randomly, 900 and 1900 msec). This change not only reduces the likelihood of invoking working memory processes, but also better controls the timing of attention orienting (subjects learn that targets may come as early as 900 msec so that they would be more likely to begin those processes right away after the cue). In addition, using shorter SOAs makes the results more directly comparable to the wide body of behavioral and ERP work using similar SOAs in endogenous cueing tasks (e.g., 800–1200 msec in Posner, Snyder, & Davidson, 1980; 700 msec in Downing, 1988; 1100 msec in Hopf & Mangun, 2000; 700 msec in Eimer, 2000).

Second, as noted above, previous studies have successfully separated cue-related activity from target-related activity (e.g., Corbetta et al., 2000; Hopfinger et al., 2000; Shulman et al., 1999). However, within the brain activity that is triggered by cues, there is a need to further distinguish activity due to the actual orienting of attention from activity due to sensory and semantic operations required for processing and interpreting the cues themselves (e.g., cue-interpretation processes). Both Hopfinger et al. (2000) and Shulman et al. (1999) included a control condition for the cues (“cue control” in Hopfinger et al., 2000; “passive cues” in Shulman et al., 1999), but these trials were in a separate passively viewed run and thus were not intermixed with the other trial types and were not behaviorally relevant. While these runs may have controlled for the most basic

sensory stimulation of the cues, they did not require any sort of processing for meaning or interpretation, and thus they were limited in their effectiveness as controls for the general cue-processing operations that would be likely to occur during the experimental attention runs. Thus, in the present study, we included a “cue-interpretation” control trial type that was randomly intermixed with the other attention-directing trial types. On these trials, participants needed to perceive, process, and interpret a behaviorally relevant cue, which instructed them that orienting of their spatial attention was unnecessary on that particular trial. These cues, which we have termed “interpret-cues,” could not be ignored, as they were intermixed with all other trial types in the same runs, and needed to be interpreted and processed to achieve correct behavior. Thus, interpret-cue trials required both sensory and cue-interpretation processes similar to those engaged by attend-cues, but did not require attentional orienting to a particular spatial location in the visual field.

Thus, our paradigm and analysis structure is “hierarchical” in nature. At the first level, activity can be separated into cue- and target-related activity. This separation and assessment of target-triggered activity allows us to determine which sensory regions would be expected to be the site of any prestimulus biasing during cued orienting, as well as to assess and separate out activity due to motor-related processes. At the second level, cue-related activity can be further split into activity related to attentional orienting and activity more related to general cue-interpretation processes, allowing us to better isolate attention-orienting networks while determining the functional roles of different subregions of the fronto-parietal executive control network.

RESULTS

Subjects were presented with a series of trials, each of which began with an instructional cue presented at fixation (Figure 1). Active-attend (i.e., attention-directing) cues consisted of either the letter L or R, which instructed the subject to attend to a location in the left or right lower visual field to detect a possible faint target that might occur in that location. Some of the instructional cues consisted of the letter P. These were the “interpret-cue” control trials, instructing the subject to not orient their attention and to not attend for targets. In some active-attend trials, a target would occur at a randomized time (900 or 1900 msec) after the onset of the cue (attend-cue-plus-target trials). In other active-attend trials (attend-cue-only trials), as well as in the interpret-cue trials, no target was presented, so that the brain response would be due to the cue only. These various trial types were randomized with “no-stim” trials (i.e., periods of fixation only) so that the full event-related responses to the various trial types could be extracted using selective averaging (Buckner et al., 1998;

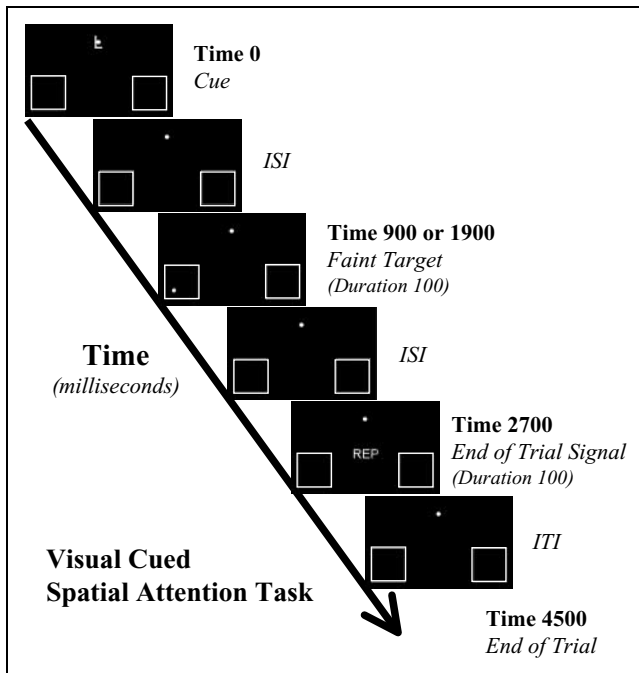


Figure 1. Schematic diagram of a compound event cued-attention trial. At the beginning of the trial, single-letter cues were presented (“L,” “R,” or “P”) at central fixation instructing the subject to either covertly attend to a location in the left or right lower visual field to detect a possible faint dot target there or to interpret the cue but then not attend for a target. In some trials, there was no target, so that the brain response would be due only to the cue. At 2700 msec, a brief visual stimulus noting the end of the trial (the letters “REP”) was presented in the midline lower visual field, at which point the subject was to press a button if they had detected a target. ISI = interstimulus interval; ITI = intertrial interval.

Burock, Buckner, Woldorff, Rosen, & Dale, 1998; Dale & Buckner, 1997; see Methods). Each trial (including the “no-stim” trials) lasted 4500 msec.

Note the hierarchical structure of the trial types: no-stims, interpret-cue, left- and right-attend-cue-only, and left- and right-attend-cue-plus-target (Table 1a). This hierarchical structure was designed to allow various key contrasts to be performed that could isolate different brain responses associated with specific processing

Table 1a. Hierarchical Design of Trial Types

Attend-left-cue-plus-target	“L”.....Target.....EOT...
Attend-right-cue-plus-target	“R”.....Target.....EOT...
Attend-left-cue-only (not followed by target)	“L”.....EOT...
Attend-right-cue-only (not followed by target)	“R”.....EOT...
Interpret-cue	“P”.....EOT...
No-stim

components (Table 1b), while also subtracting out the overlapping hemodynamic responses. All stimulus trial types were also statistically contrasted with the lowest condition level, the no-stim trials. The results of the voxelwise analysis of these various contrasts will first be presented, followed by the regions-of-interest (ROI) analyses.

Figure 2 displays horizontal sections (at the Talairach z coordinate of +44; Talairach & Tournoux, 1988) showing areas whose activation varied significantly between the various trial types in the voxelwise analyses. Contrasts between each of the different trial types (interpret-cue, left and right-attend-cue-only, left- and right-attend-cue-plus-target) versus the no-stims are shown in Figure 2A. These contrasts reveal the full event-related activation maps for these trial types relative to a low-level baseline. Figure 2B shows the corresponding hierarchical subtraction images relative to the level closer below it in the hierarchy (e.g., left-attend-cue-only vs. interpret-cue). Also shown are maps derived for the activations collapsed over side of attentional focus (e.g., the average response to left-attend-cue-only and right-attend-cue-only vs. interpret-cues).

Cue-Related Activity: Effects in the Lateral versus Medial Regions of the Frontal and Parietal Cortex

Figure 2A reveals that, relative to the no-stim trials, all trials with instructional cues (including interpret-cue trials) activated the parietal and frontal areas. These

Table 1b. Functional Contrasts and Associated Processes

Contrast	Associated Processing Component(s)
Interpret-cue versus no-stim	Visual stimulation by foveal presentation of a letter Processing and interpreting a symbolic cue
Left-attend-cue-only versus interpret-cue	Attentional orienting to the left visual field Prestimulus biasing of right (contralateral) visual sensory cortex
Right-attend-cue-only versus interpret-cue	Attentional orienting to the right visual field Prestimulus biasing of left (contralateral) visual sensory cortex
Left-attend-cue-plus-target versus left-attend-cue-only	Processing of an attended target in the left visual field
Right-attend-cue-plus-target versus right-cue-only	Processing of an attended target in the right visual field

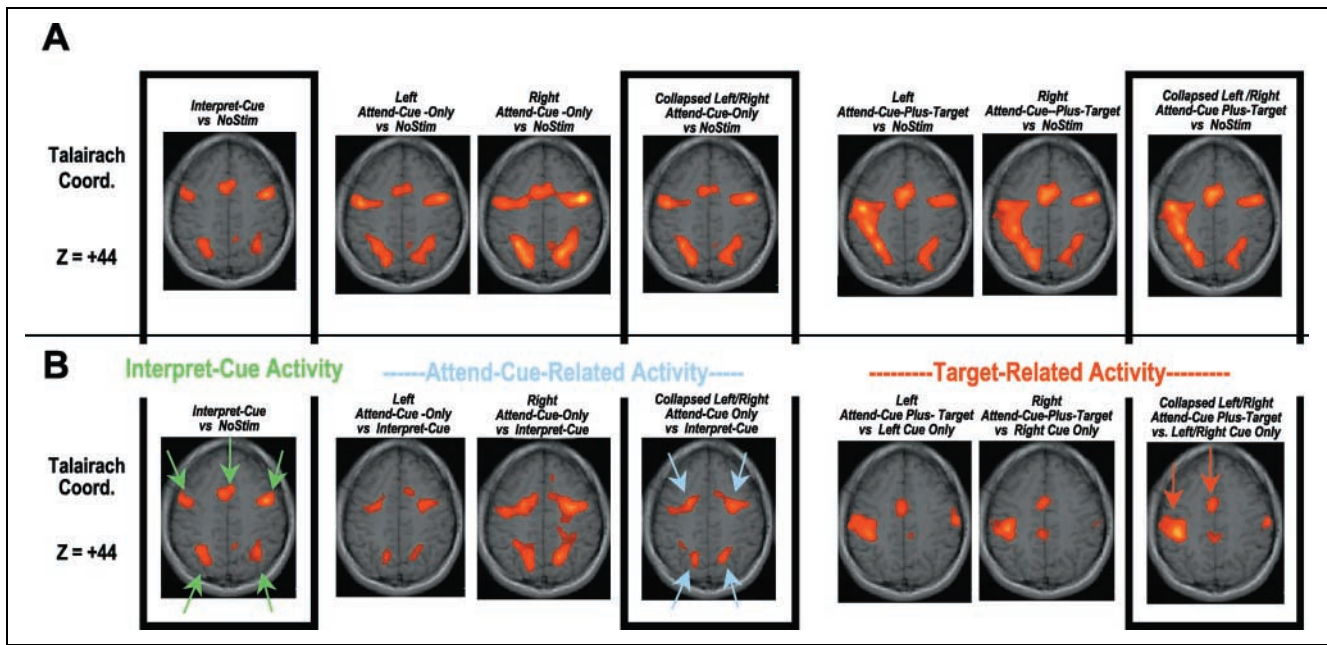


Figure 2. Activation maps for various contrasts between the trial types. **(A)** Contrasts of each of the trial types versus the no-stim at the level of superior cortex, showing the full event-related activation maps relative to a low-level baseline. The contrast images are overlaid on the T1-weighted structural images from a single subject that was normalized into the same Talairach space. Note that, relative to the no-stim trials, all of the instructional cues (including the interpret-cues) activated parietal and frontal areas, with the active-attention trial types appearing to elicit a greater and more extensive activation in these regions. **(B)** The corresponding hierarchical subtraction images relative to a level closer below it in the hierarchy of trial types revealed greater detail concerning the distributions of these effects. The frontal and parietal regions activated by the interpret-cues were quite lateral (green arrows), whereas the regions activated by the attend-cues relative to the interpret-cues, which were more specific to attentional orienting, were more medial (blue arrows). At this level of the brain through superior cortex, the target-related processing derived from the attend-cue-plus-target versus attend-cue-only contrast (right panels) revealed significant activity in left sensorimotor cortex and the supplementary motor area (SMA) (red arrows). For both (A) and (B), the color scale for these *t* value contrast images range from 3.25 to 8 (dark red to yellow) for the contrasts that are collapsed across the left/right factor (columns 4 and 7) and from 2.75 to 8 (dark red to yellow) for the uncollapsed contrasts (columns 1–3, 5–6).

areas are in the vicinity of those reported previously as being involved in the orienting of visual spatial attention (e.g., Corbetta et al., 2000; Hopfinger et al., 2000). Notably, however, even the interpret-cues, for which subjects did not need to orient their spatial attention, activated a portion of these putative attentional control areas in the frontal and parietal cortex. However, the active-attention trial types appeared to elicit a greater and more extensive activation in some of these regions.

The hierarchical contrasts shown in Figure 2B revealed important details about these activation patterns—namely, that the frontal and parietal areas activated by the interpret-cues and attend-cues did not have the same spatial distribution. Specifically, the activations elicited by the interpret-cues were quite lateral, whereas the active attend-cues triggered this lateral activity plus additional activation in more medial portions of the frontal and parietal regions. This pattern is reflected in the hierarchical contrast maps in that, after subtracting out the interpret-cue response from the attend-cue-only responses, the residual activation is considerably more medial than the activation for interpret-cues versus no-stims (see Figure 2).

Cue-Related Activity: Effects in the Midline Frontal Areas

All of the instructional cues (including interpret-cues) also activated several midline dorsal regions of the frontal cortex (Figure 2A). As before, the hierarchical contrasts (Figure 2B) revealed additional specificity. The interpret-cues activated an anterior medial dorsal area that appeared to be in the anterior cingulate cortex (ACC). In the activation map of the contrast of attend-cue-only versus interpret-cue, there was no additional activity apparent in this region (although ROI analyses presented below revealed that there was a small degree of additional activity). This response pattern differed sharply from that of the slightly more posterior supplementary motor area (SMA), which was not responsive to any of the cue-only trials (see below).

Target-Related Activity

The hierarchical contrast between attend-cue-plus-target and attend-cue-only trials was aimed at isolating activity related to target processing by subtracting out

the overlapping cue-related activity. To encourage subjects to begin attending immediately for these trials, targets were sometimes presented very soon after the cue (900 msec), although on other trials they were presented later (1900 msec). However, since no target would come in the attend-cue-only trials, subjects would attend continuously from the cue until the end of trial (EOT; 2700 msec) for these trials. Thus, the best contrast for identifying (attended) target-related activity was attend-cue-plus-late-target minus attend-cue-only, since the duration of attentional maintenance was most similar in these two conditions. The results of these contrasts are shown in Figure 2B (right panels) (see also Table 2).

In contrast to the cue-sensitive ACC mentioned above, the attend-cue-plus-target versus attend-cue-only contrast (a subtraction specific for isolating activity associated with processing the target and performing the right-hand button-press motor response) yielded a

slightly more posterior medial dorsal region that seems likely to be the SMA. Target-related activity was also observed in the left posterolateral inferior frontal cortex (likely SII), in the visual sensory cortex contralateral to the target, and in the thalamus (see below). Strong target-related activity was also observed in various additional motor structures, including the left sensorimotor cortex (Figure 2B), right cerebellum, and the caudate, consistent with the task requirement of pressing a button with the right hand upon detection of a target.

Region-of-Interest Analyses: Frontal and Parietal Executive Control Areas

The main ROIs that were selectively activated in the hierarchical contrasts described above were analyzed further. More specifically, the event-related hemodynamic response functions (HRFs) generated by time-locked averaging for the various event types in those

Table 2. Talairach Coordinates and Cluster Sizes of Key ROIs Activated in the Voxelwise (SPM) Analyses, along with Results of the Statistical Contrasts between the Raw Event-Related Responses in These ROIs

ROI Number: Regions	Talairach Coordinates			Defining Contrast	Number of Voxels	P value for ROI (One-Tailed)
	x	y	z			
1: Left lateral frontal cortex	-48	3	43	interpret-cue versus no-stim	71	2.5E-05
2: Left medial frontal cortex	-23	-4	46	attend-cue-only versus interpret-cue	48	.0051
3: Left lateral parietal cortex	-28	-62	41	interpret-cue versus no-stim	79	6.2E-06
4: Left medial parietal cortex	-18	-58	48	attend-cue-only versus interpret-cue	72	.017
5: Right medial frontal cortex	27	1	46	attend-cue-only versus interpret-cue	73	.00089
6: Right lateral frontal cortex	46	6	43	interpret-cue versus no-stim	58	6.9E-05
7: Right medial parietal cortex	20	-57	50	attend-cue-only versus interpret-cue	66	.00071
8: Right lateral frontal cortex	32	-61	45	interpret-cue versus no-stim	106	1.7E-07
9: ACC	-1	11	47	interpret-cue versus no-stim	24	6.3E-07
10: SMA	-2	0	43	attend-cue-plus-target versus attend-cue-only	32	3.7E-08
11: Left sensorimotor cortex	-44	-23	50	attend-cue-plus-target versus attend-cue-only	242	2.1E-09
12: Left SII	-55	-25	20	attend-cue-plus-target versus attend-cue-only	75	2.7E-08
13: Left TPJ	-52	-52	14	attend-cue-plus-target versus no-stim	17	1.1E-06
14: Right TPJ	58	-45	16	attend-cue-plus-target versus no-stim	48	2.0E-08
Right cerebellum	21	-57	-17	attend-cue-plus-target versus attend-cue-only	124	8.8E-07
Left thalamus	-12	-21	0	attend-cue-plus-target versus no-stim	112	5.6E-09
Right thalamus	12	-21	0	attend-cue-plus-target versus no-stim	106	1.5E-08
Left dorsal occipital cortex	-21	-90	14	(all) left-attend-cue-plus-target versus (all) right-attend-cue-plus-target	57	.0012E-08
Right dorsal occipital cortex	30	-85	15	(all) left-attend-cue-plus-target versus (all) right-attend-cue-plus-target	119	2.5E-06

ROIs were extracted by subtracting the response to each of these types relative to the no-stims, yielding HRF time-courses (with the adjacent-trial overlap subtracted away) for each of the event types in these ROIs. The overlaying of these ROI response curves for the various event types enabled a closer examination of the relative specificity of the various cognitive processing functions in these various brain areas.

The ROIs in the superior cortical areas that were analyzed in this way are indicated in Figure 2B (arrows). They include the lateral areas in the frontal and parietal cortex (derived from the interpret-cue vs. no-stim contrast), the medial areas in the frontal and parietal cortex (derived from the attend-cue-only vs. interpret-cue contrast), the medial dorsal ACC (derived from the interpret-cue vs. no-stim contrast), and the SMA and left motor cortex (derived from the attend-cue-plus-target vs. attend-cue-only contrasts). In each of these ROIs, statistical analyses were performed on

the peak amplitudes of the raw time-locked averaged data, which confirmed that there were significant amplitude differences at the regional level between the relevant trial types (Table 2).

Figure 3 shows the HRF time-courses for the various event types, relative to the no-stims, in the lateral and medial areas of frontal and parietal cortex. The figure shows that, in the lateral areas, the responses to the interpret-cues were roughly as large as those for the attend-cue-only and attend-cue-plus-target trials. In contrast, in the medial areas the responses to the interpret-cues were considerably smaller than those for the attend-cue trials. These differential response amplitude patterns are reflected in the maps of Figure 2B. More specifically, the similar level of activity in the lateral areas for the different types of cues resulted in those regions subtracting out in the maps for the attend-cue-only versus interpret-cue contrast. However, in these same maps the differential activity

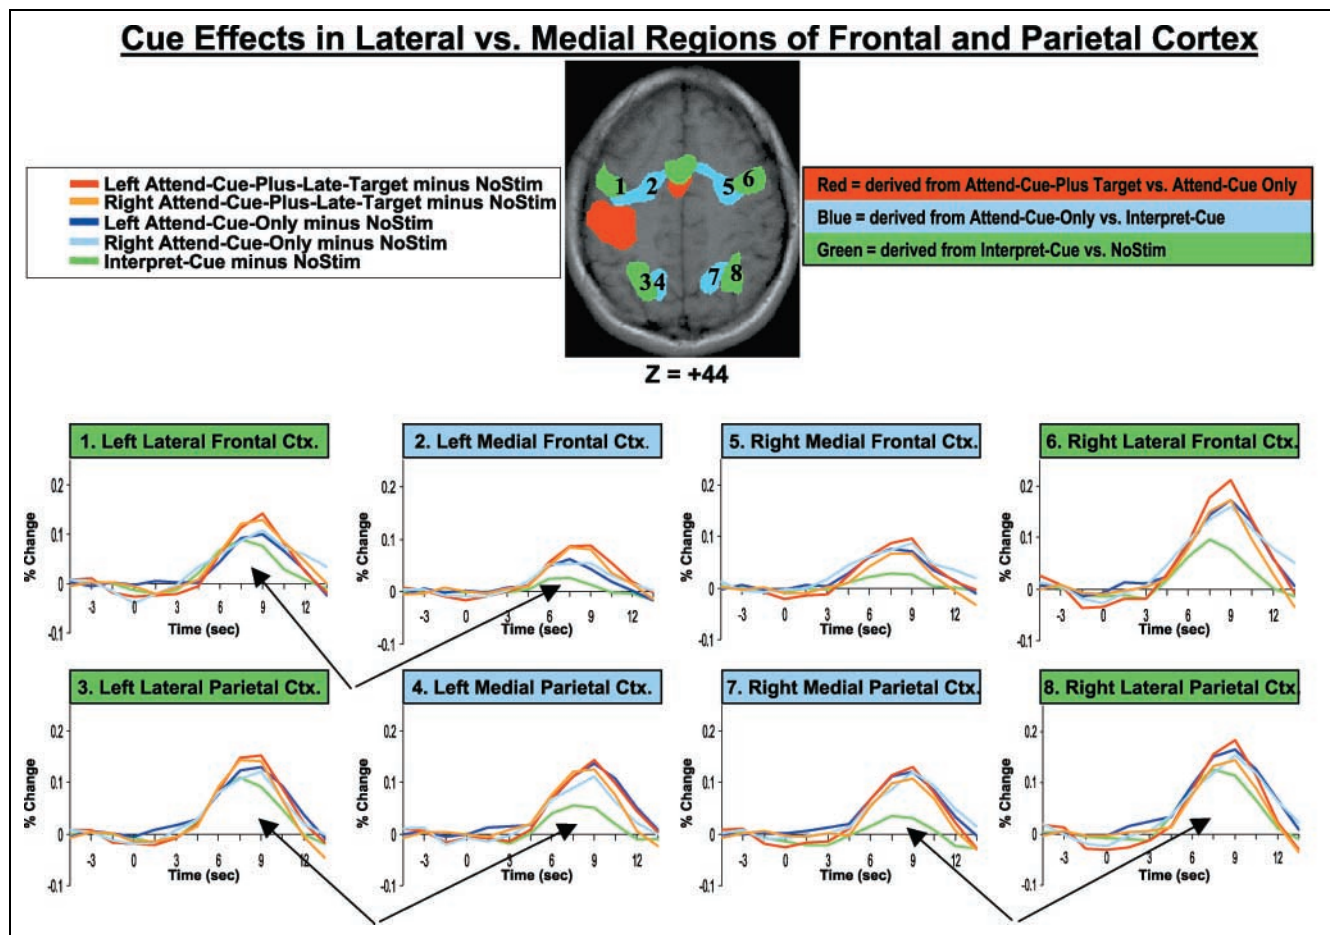


Figure 3. ROIs and hemodynamic time-courses in the superior frontal and parietal cortex. The main superior cortex ROIs that were selectively activated in the hierarchical contrasts are shown in different colors on a horizontal section ($z = +44$). The HRF time-courses for the various event types, relative to the no-stims, are shown for the lateral and medial ROIs in the frontal and parietal cortex (ROIs 1–8, green and blue areas). By subtracting off the responses to the no-stims, the overlap from responses to adjacent trials in the sequence was removed. The time-courses show that, in the lateral areas, the responses to the interpret-cues were almost as large as those for the attend-cue-only and attend-cue-plus-target trials, but in the medial areas, the interpret-cue responses were considerably smaller than for the attend-cue trials (black arrows).

more medially for these trial types resulted in robust differences in the medial regions.

To statistically assess whether lateral and medial regions of parietal and frontal cortex responded differently to attend-cues and interpret-cues, we tested the data for an interaction between cue type (attend cue-only, interpret-cue) and region (medial, lateral) separately for the left frontal, right frontal, left parietal, and right parietal ROIs. This was accomplished by entering the peak amplitudes of the raw event-related responses to interpret-cues trials and attend-cue-only trials from these ROIs into a random effects analysis of variance (ANOVA) (Figure 4A). The statistical interaction confirming this differential distribution of activity was significant in both left and right parietal cortices ($p < .01$ in both cases) and in left frontal cortex ($p < .05$); this interaction did not hold in the right frontal cortex. Note that since the magnitude of interpret-cue and attend-cue-only activations is nearly the same in lateral regions, there should be no difference in spatial extent based on magnitude alone, and yet the interpret-cue activity drops off much more dramatically in the more medial regions. Thus, the interaction effect is most likely due to a differential distribution of attend-cue and interpret-cue activity and is not an artifact based on spatial extent or thresholding effects in the t maps.

Because there were differences in overall activation levels for the lateral and medial areas (the medial regions tended to be more weakly activated), these interaction analyses were also performed after normalizing the activation amplitudes to the mean level (across subjects) of the stronger condition—i.e., the attend-cue condition—in both lateral and medial brain regions (Figure 4B). After such normalization, the significance of these interactions became even stronger and more highly significant (left parietal = .01, right parietal = .005, left frontal = .005, right frontal = still not significant). The statistical interactions in these frontal and parietal subregions therefore provide statistical support for the view that lateral and medial regions have a differential selectivity for the spatial orienting of attention.

Region-of-Interest Analyses: Midline Dorsal Frontal Areas

Figure 5A shows the time-courses of the time-locked-averaged response functions for the ACC, SMA, and left sensorimotor cortex ROIs. For the ACC (ROI 9), there appeared to be a gradation of responses for the interpret-cue, attend-cue-only, and attend-cue-plus-target trial types. ROI analyses of the peak amplitudes confirmed these relationships in the ACC, indicating some significant activity for interpret-cue trials alone (relative to no-stims) ($p < .0001$), significantly more activity for attend-cue-only trials than for interpret-cues ($p < .01$),

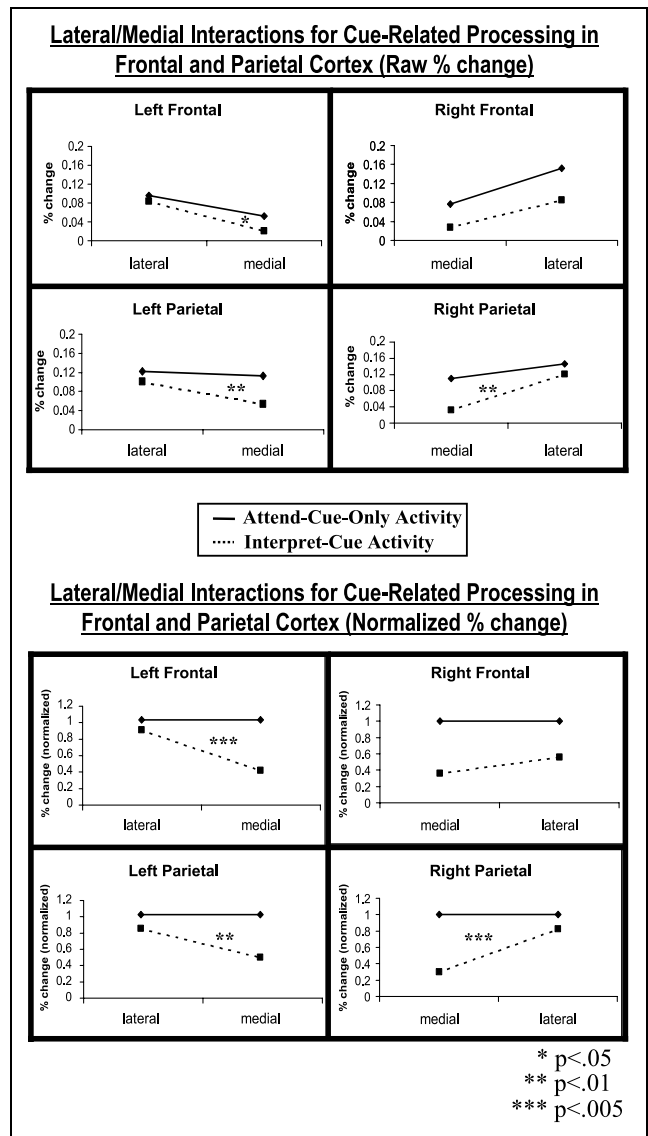
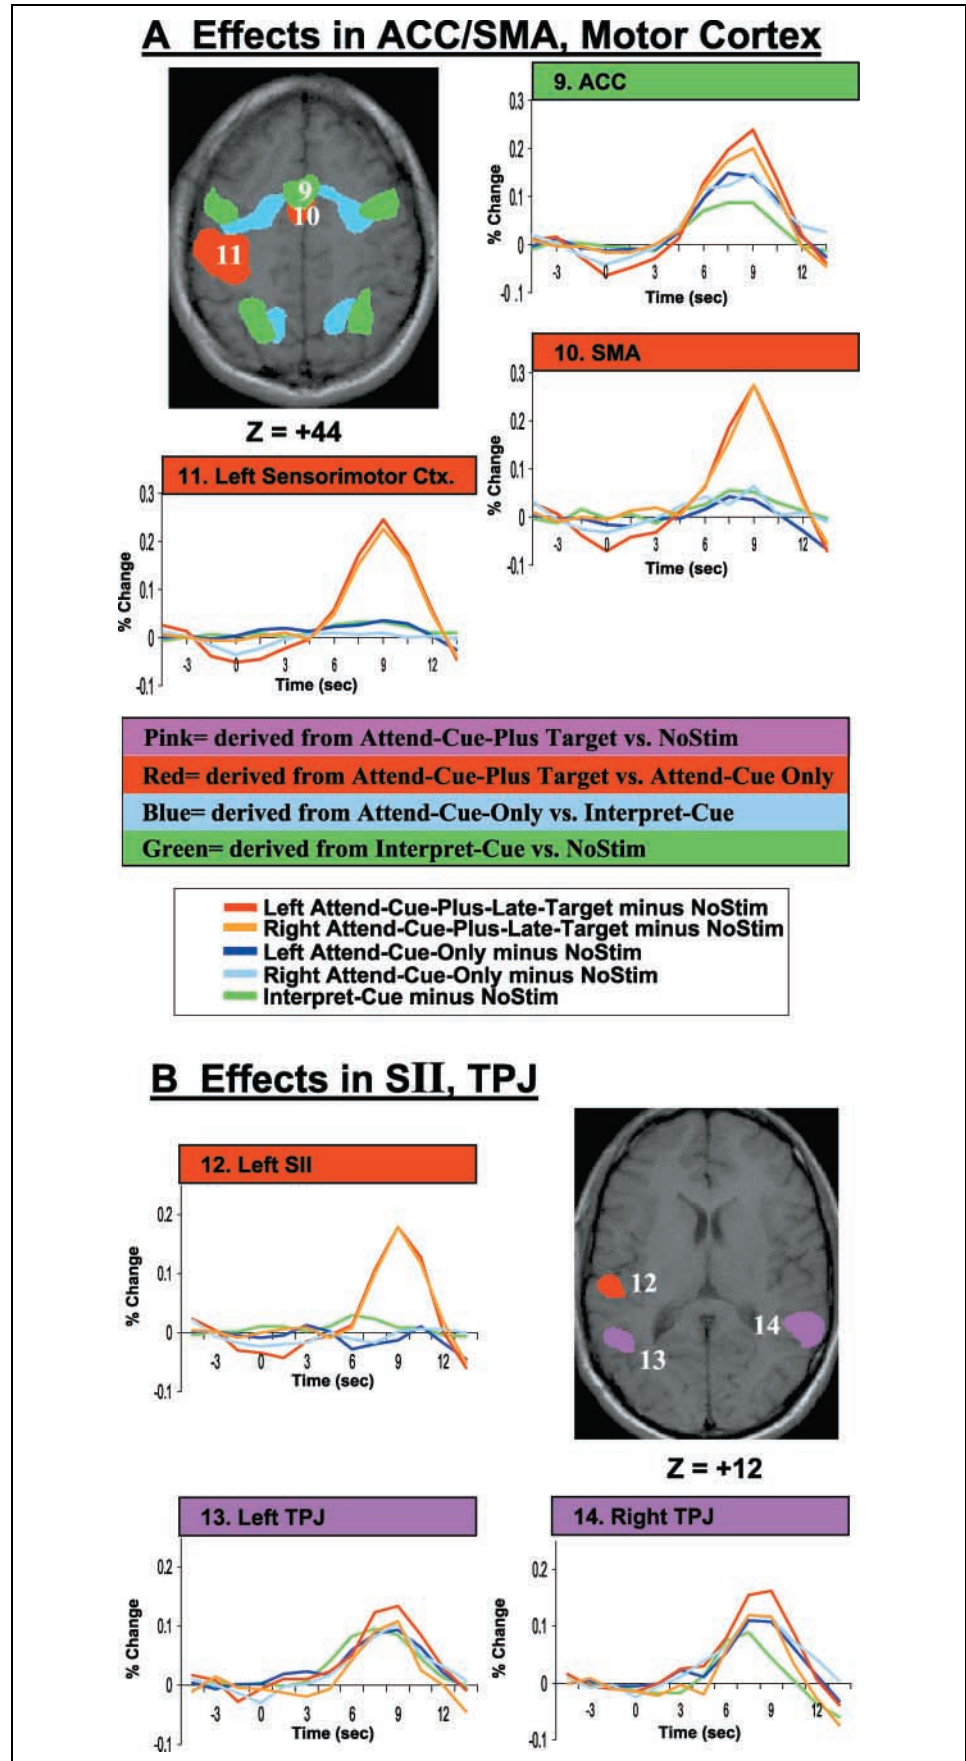


Figure 4. Lateral/medial interactions in the superior frontal and parietal cortex for cue-related activity. Peak amplitudes for the time-locked average responses to interpret-cue and attend-cue-only trials (after subtracting out the no-stim averages) were entered into repeated measures ANOVAs. These measures and the statistical analysis of them assessing for a significant interaction are shown both for the raw amplitude values and for the amplitude values after being normalized to the mean level (across subjects) of the stronger condition (attend-cue). This interaction was significant for both the parietal lobes and the left frontal cortex. The statistical interactions in these frontal and parietal subregions provide additional statistical support for the differential selectivity of these areas for the spatial orienting of attention.

and significantly more still for attend-cue-plus-target trials than for attend-cue-only trials ($p < .0002$). In sharp contrast, just posterior to this ACC region, in what would appear to be SMA (ROI 10), the response was almost completely selective for those trials with targets, with little response for the interpret-cues or the attend-cue-only trials.

Figure 5. ROIs and hemodynamic time-courses in the medial dorsal cortex, SII, and TPJ. As in Figure 3, various ROIs selectively activated in the hierarchical contrasts are shown on two horizontal sections ($z = +44$ and $z = +14$), with the HRF time-courses for the various event types, relative to the no-stims, in these ROIs. (A) Time-courses in the ACC, SMA, and left sensorimotor cortex. Notice the graded response across trial types in the ACC and the high selectivity in the SMA and in the left sensorimotor cortex for trials with targets. (B) Time-courses in left SII, and in left and right TPJ. Notice the high selectivity in left SII for target trials. SII = secondary somatosensory cortex; TPJ = temporal-parietal junction.



Region-of-Interest Analyses: Target-Related Activations

Like the medial dorsal SMA region, the activations in left sensorimotor cortex were present almost exclusively for those trials with targets (Figure 5A), consistent with the task requiring a button-press only on those trials. Similarly, a left inferior parietal region (likely the left secondary somatosensory area, SII) also was active only for trials with targets (Figure 5B), with flat responses for the other trial types.

Because of previous reports indicating target-selective activity in the left and right temporal–parietal junctions (TPJs) (e.g., Corbetta et al., 2000), additional analyses were performed on these brain areas. Figure 5B shows these regions as activated by the attend-cue-plus-target minus no-stim contrast. The time-courses from these ROIs reveal that all the trial types appeared to elicit some activation in these regions (relative to no-stims). In the left TPJ, neither the SPM maps nor the ROI analyses indicated any significant differences in the activation level between the various trial types, including those with targets. In the right TPJ, the ROI analysis indicated that there was slightly more activity for attend-cue-only than for interpret-cue trials ($p < .02$) and also slightly more activity for attend-cue-plus-target than for attend-cue-only trials ($p < .05$).

Some previous studies have also found that the frontal or parietal regions activated by cues were also activated by the targets themselves (i.e., intraparietal sulcus in Corbetta et al., 2000; precentral sulcus/superior frontal sulcus in Shulman et al., 1999). As noted above, t contrasts at the voxelwise level do not indicate this in the present experiment. At the ROI level, however, the time-courses shown in Figure 3 suggest that cue-plus-target trials elicited slightly more activity than cue-only trials in some of these regions. Although this additional activity for the targets in these regions was fairly small (on the order of 10–15% of the total activity in these regions), it did reach significance in the ROI analyses for three of the eight regions (left lateral frontal, left medial frontal, and left lateral parietal cortex) based on a one-tailed, random effects t test of two time points around the response peak. However, it is clear that the bulk of activity observed in these areas, at least in the present experiment, was triggered by the cues, rather than the targets. It is possible that target activity was much less pronounced in our study relative to previous ones because we used faint, small targets.

Effects in the Visual Cortex and Thalamus

Both attend-cue-plus-target and attend-cue-only trials enhanced activity in the occipital cortex contralateral to the direction of attention, an effect that is best seen in a comparison of left-attend versus right-attend conditions (Figure 6A). The contrast for left-attend-cue-plus-

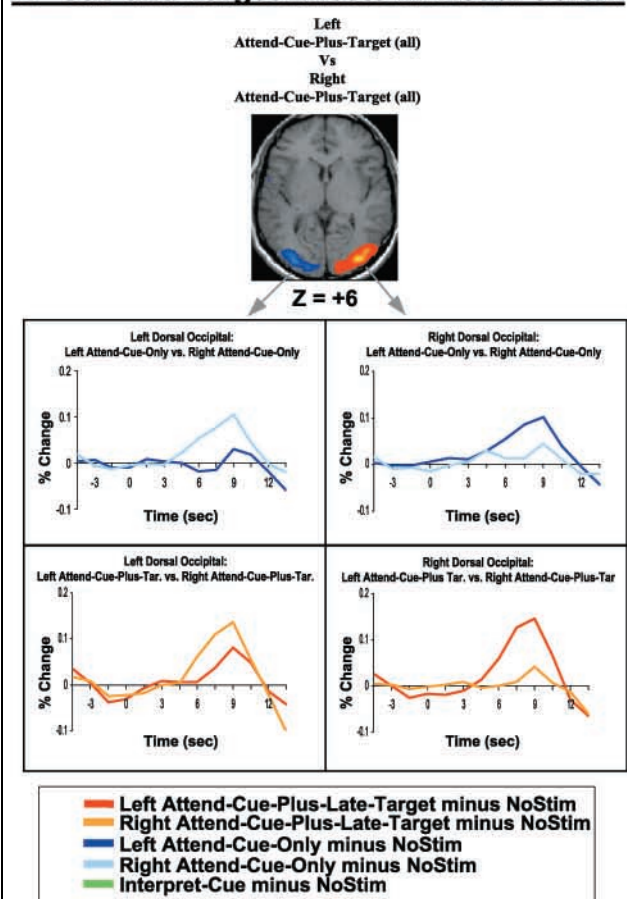
target versus right-attend-cue-plus-target (Figure 6A, top) showed particularly strong contralateral effects in the dorsal occipital cortex, presumably in part because the targets in these trials were unilateral stimuli in the lower visual field. To assess for possible pretarget biasing effects in visual sensory cortex in response to the attention-directing cues, this contrast of left-attend-cue-plus-target versus right-attend-cue-plus-target was used to define ROIs for these low-level, retinotopically organized, visual sensory areas. The HRF time-courses in these ROIs (after subtracting out the no-stim responses) were then extracted for these conditions, as well as for the left- and right-attend cue-only conditions (Figure 6A). The ROI time-courses show the clear relative enhancement of the response level in the dorsal occipital areas contralateral to the direction of attention for both the attend-cue-only trials and attend-cue-plus-target trials. The contralaterality of attend-cue-plus-target activity was confirmed statistically by ANOVA of the peak amplitudes of the raw time-locked averages in these dorsal occipital ROIs ($p < .000003$). In addition, the same ANOVA test applied to the peak amplitudes for the attend-cue-only trials in these dorsal occipital ROIs confirmed that the centrally presented attention-directing cues also produced contralateral effects in these same target processing areas even when no targets were presented ($p < .0003$).

Robust activations were also seen bilaterally in the thalamus, particularly for the attend-cue-plus-target trials (Figure 6B). ROIs defined from the attend-cue-plus-target versus no-stim contrast were used to extract time-courses for all the event types in these areas and to perform ROI analyses. The time-course response functions indicated that all the event types, including the interpret-cues, activated the thalamus. The attend-cue-plus-target trials showed a particularly large response, but even the attend-cue-only trials appeared to have enhanced activity relative to the interpret trials (arrows in Figure 6B). A statistical analysis of the peak amplitudes in these thalamic ROIs indicated that this difference (i.e., attend-cue-only vs. interpret-cue) was significant ($p < .025$ and $p < .01$ for the left and right thalamus, respectively).

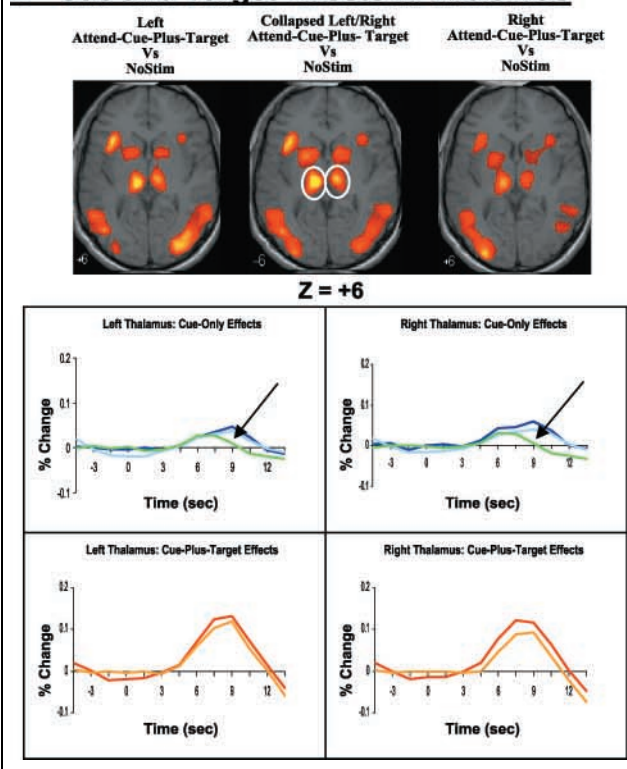
DISCUSSION

In this study, event-related fMRI was used to study the brain networks involved in the voluntary control of spatial attention. A spatial cueing paradigm was employed, in which an initial cue instructed subjects as to whether and where to direct their attention for a possible upcoming target stimulus. In contrast to previous event-related neuroimaging studies, the parameters used for this study, including the cue–target SOAs and intertrial intervals, were very similar to previous behavioral and ERP studies using cued-attention paradigms with endogenous cueing (e.g., Hopf & Mangun, 2000;

A Cue and Target Effects in Visual Cortex



B Cue and Target Effects in Thalamus



Posner et al., 1980). Furthermore, we extended the methodologies for separating cue- and target-related activity introduced in prior studies (cf. Corbetta et al., 2000; Hopfinger et al., 2000) in ways that allowed us to distinguish not only between target- and cue-related activity, but also between cue-related processes that enable attentional orienting versus those related to more general aspects of cue processing, such as cue-symbol interpretation. Using this method, activity related to attentional orienting, presumably “pretarget biasing,” was found in the early visual cortical areas and thalamus. The ACC, an area with a controversial role in attentional orienting, showed a graded response across cue types, with some degree of activity during cue-interpretation processes and additional activity above and beyond this in response to attention-directing cues. Most importantly, the data revealed functional parcellation within the frontal and parietal control regions, with the more medial regions being more specific for attentional orienting and the more lateral regions being involved in more general aspects of cue processing, such as cue-symbol interpretation.

Executive Control and Functional Specificity in the Frontal and Parietal Regions

A key result in this study was that all of the instructional cues (including the interpret-cues) activated the parietal and frontal areas that were near the same areas previously implicated as being involved in the orienting of spatial attention (e.g., Corbetta et al., 2000; Hopfinger et al., 2000). However, in both frontal and parietal cortex, the activity elicited by the interpret-cues was found in the more lateral of the frontal and parietal regions, whereas the attend-cues triggered this activity PLUS additional activation in adjacent subregions that were more medial. This activity pattern was explored in greater detail by extracting the time-locked averaged time-courses in these ROIs for the attend-cue-only and interpret-cue trial types and analyzing their peak ampli-

Figure 6. ROIs and response time-courses in the occipital cortex and thalamus. (A) ROIs activated in the dorsal occipital cortex in the left-attend-cue-plus-target versus right-attend-cue-plus-target conditions, along with the corresponding HRF time-courses in these ROIs for the various event types relative to the no-stims. Note the larger responses contralateral to the direction of attention. For target trials, this would include an enhanced response to the unilateral targets at attended positions. However, for attend-cue-only trials, no target is presented and, therefore, the activation of these contralateral regions suggests a retinotopically based “biasing” signal in expectation of a possible target in the attended region. (B) ROIs activated in the thalamus in the hierarchical contrasts, along with the corresponding HRF time-courses for the various event types relative to the no-stims. The response for attend-cue-plus-target trials was substantially larger than for the other trial types, but even the attend-cue-only responses were enhanced relative to the interpret-cue responses (arrows). The *t* value color mapping is identical to that used in Figure 2.

tudes. These additional analyses revealed statistically significant lateral/medial interactions for these ROIs and confirmed that the differential patterns observed in the SPM activation maps were not an artifact of the particular threshold at which these maps were created. Note also that the attend-cue activity found in the medial regions was likely not due to motor preparation for two reasons. First, motor preparation for a right-handed button-press would likely be lateralized and these activations were not. Second, the lack of cue-period activation in areas much more sensitive to motor preparation (e.g., SII, left sensorimotor cortex, and even SMA) strongly argues that activation seen in other areas is unlikely to be motor preparation.

The similar activation for interpret-cue and attend-cue-only trials in the lateral frontal and parietal regions suggests that these areas are activated by the engagement of some common cognitive process (or processes) for these different cue types. One clear candidate process is cue-symbol interpretation. All the cues in this study, including the “interpret” ones, were relevant and needed to be attended and processed by subjects to interpret the meaning of the cue stimulus and decide what to do on that trial. In contrast, attend-cue-only trials produced substantially greater activity than interpret-cue trials in the medial frontal and parietal regions, suggesting that these regions play a relatively specific role in the orienting of visual spatial attention.

Note that the present design allowed distinctions between lateral and medial regions of the frontal and parietal cortex that would be more difficult without an intermediate cue-condition level that was randomized into the sequence (e.g., the “interpret-cue” trials in the present study). Hemodynamic responses triggered by the instructional cues themselves—for example, relative to the prestimulus baseline rather than to an intermediate condition level—could reflect processes that are not specific to visual attentional orienting. Moreover, the inclusion of interpret-cue trials in separate sensory control blocks (e.g., Hopfinger et al., 2000) without a task is not sufficient to control for such processes. When presented in a separate block, these cues are much less likely to be fully processed and to engage cue-interpretation processes. By using a randomized trial sequence that included both attend-cue and interpret-cue trials, we were able to distinguish functionally specific activation patterns that likely correspond to distinct aspects of cue processing.

As noted above, HRF time-courses for the various event types, relative to the no-stims, provided further evidence for a functional dissociation between lateral and medial regions of the frontal and parietal cortex. These time-courses, however, revealed that although the interpret-cue response amplitudes in the medial areas were considerably lower than those for attend-cue-only trials, they were not flat. This may be due to

some spatial smearing of the hemodynamic responses between the regions (due, e.g., to the averaging of the data across subjects), although it could simply reflect that the activations in these adjacent executive control areas are not “all or none.” In either case, it is clear that the medial areas likely perform a function that is relatively specific for the orienting of visual spatial attention while the lateral areas likely perform a more general function, such as cue-symbol interpretation.

Given the present evidence that lateral and medial regions of the frontal and parietal cortex perform somewhat different functions, one might wonder why they would be so physically near to each other. One possibility is that the cue-symbol interpretation required for all the cue types involves higher level executive processes, which includes evaluating the cue meaning and making the decision of what to do in response to that cue, even if it does not involve the spatial orienting of attention. It may be that there is a functional advantage (e.g., speed of communication) in having neuroanatomical proximity between higher level brain areas performing interpretive and decisional processes related to meaningful environmental stimuli and higher level areas that perform orienting of attention in response to such stimuli when appropriate.

The frontal areas activated by cues appeared to overlap with the frontal eye fields. Such results are consistent with previous studies reporting neuroanatomical overlap of the frontal areas involved in the covert focusing of attention (as used here) and those areas involved in overt control of eye position (Nobre, Gitelman, Dias, & Mesulam, 2000; Corbetta, 1998). Note, however, this result does not suggest that subjects were moving their eyes, as it is known that the frontal eye fields can be activated without concomitant eye movements (Gitelman, Parrish, LaBar, & Mesulam, 2000).

Finally, the current findings of cue-triggered activity in the frontal and parietal cortex occurred at relatively short SOAs that are comparable to those used in the behavioral and ERP literatures. This helps to rule out the possibility that cue-triggered activity observed in these areas in some previous studies (e.g., Hopfinger et al., 2000) occurred mainly because long SOAs engaged working memory processes to a greater degree than is normally the case. However, the Hopfinger et al. (2000) study reported cue-related frontal activation that included more anterior regions of left lateral prefrontal cortex that were not activated in the present experiment. Thus, it is possible that activation of these anterior regions was associated with high working memory demands (due to the long SOAs) rather than visual spatial attention per se. We acknowledge, however, that attention-directing cues may engage working memory processes even at short SOAs. Indeed, some evidence indicates that working memory may be used to store representations of task-relevant (i.e., to-be-attended) stimuli (e.g., de Fockert, Rees,

Frith, & Lavie, 2001) and spatial locations (e.g., Awh, Anillo-Vento, & Hillyard, 2000).

Executive Control and Functional Specificity for Response Selection and Target Processing

By contrasting cue-plus-target trials with cue-only trials, we were able to extract target-related activity. As expected, targets selectively activated classical motor-related areas associated with a right-handed button-press—namely, the left sensory-motor cortex, the right cerebellum, the caudate, and the SMA (see Figure 5 and Table 2). HRF time-courses for each of the trial types (after subtracting out the overlap by contrasting to the no-stims) revealed that each of these ROIs was selectively activated by target processing. In particular, activation occurred when there was a target (and hence a button-press) on attend-cue-plus-target trials, while the HRFs for attend-cue-only and interpret-cue trials, in which no motor response was required, were relatively flat.

Target effects in the midline dorsal regions of the frontal cortex were somewhat more complex. Two adjacent regions were robustly activated, but distinguished themselves by behaving differently as a function of trial type. The more anterior of the two, which appeared to be dorsal ACC, showed a graded response to the different trials types, with some response for the interpret-cues, more for attend-cue-only, and still more for attend-cue-plus-target. In sharp contrast, a slightly more posterior region, which seems likely to be SMA, was activated strongly when there were targets (and button-presses), but exhibited little response in the cue-only conditions.

The more anterior area is likely to include the same ACC area that has been activated in numerous experiments in which various cognitive processes are engaged (reviewed in Paus, 2001). Some recent findings are consistent with the view that this area is important mainly for detecting and/or resolving response conflict (e.g., Botvinick, Braver, Barch, Carter, & Cohen, 2001; Van Veen, Cohen, Botvinick, Stenger, & Carter, 2001; MacDonald, Cohen, Stenger, & Carter, 2000; Carter, Botvinick, & Cohen, 1999). However, it has been shown that this area is activated just by anticipation of targets (Murtha, Chertkow, Beauregard, Dixon, & Evans, 1996). Moreover, in block-design attention experiments in which the task was to detect target stimuli in a stream of nontarget stimuli and the target frequency was manipulated, the ACC showed substantial activation relative to passive viewing regardless of whether there were many (16%) or very few (2%) targets, but with little difference in activity between the many- and few-target conditions (Woldorff, Matzke, Zamarripa, & Fox, 1999). These studies suggest that the functional contribution of the ACC to selective attention may not be limited to detecting response conflict. Consistent with this view, the interpret-cue trials in the present event-related study, which should not have evoked either motor

preparation or conflict detection processes, activated the ACC (although not the slightly more posterior SMA region). Importantly, these cues were relevant stimuli that needed to be attended to and processed (e.g., like the nontarget attended “standards” in attentional stream experiments, such as Woldorff et al., 1999). Active attentional orienting without targets in the present study activated this same ACC area somewhat more, and trials with targets produced still stronger activation in this area. These findings, especially those for the interpret-cue trials suggest that the ACC contributes to selective attention at nonresponse levels of processing.

It should also be noted that in the more posterior SMA area trials with targets elicited strong responses, whereas both the attend-cue-only and interpret-cue trials elicited little activity. This pattern of results suggests that there was not substantial motor preparation occurring for these other conditions or, at the very least, it was not sufficient to produce much activation in this putative motor-planning area.

Lastly, significant effects of trial type in the TPJ region identified by Corbetta et al. (2000) were rather limited in the present study. One of the goals of Corbetta et al.’s study was to examine the effects of manipulating cue validity. The authors reported enhanced activity in both the left and right TPJs for targets, with the right TPJ being more activated for invalidly cued than for validly cued targets. In the present study, there was some activation in the region of the TPJ on both sides of the brain. However, this activation was elicited by all the cue trial types (relative to the no-stims), did not differ at all between the trial types in the left TPJ ROI, and differed only slightly between the trial types, even for those trials having targets, in the right TPJ ROI. We speculate that this very limited target-related activation in the TPJ in the present study might have been due to the targets consisting of very faint dots that were always valid when they occurred and the task being just to detect their occurrence.

Effects in the Visual Cortex and Thalamus

Consistent with the results previously reported in Hopfinger et al. (2000), the attend-cue-only trials enhanced activity in the occipital cortex contralateral to the direction of attention and diminished activity ipsilaterally. This activity may reflect biasing of activity in favor of expected target stimuli in sensory cortices (also see Kastner et al., 1999). In the present study, this enhancement was contralateral to the direction of attention and mainly present in the dorsal occipital cortex, as expected for attention directed laterally and to the lower visual field, which is represented in that part of cortex (Serenio et al., 1995). In addition, these cue-triggered attention effects were in the same visual cortex locations activated by the contralateral unilateral lower-visual-field targets, adding support for the view that the effects of attention are

retinotopically organized (Noesselt et al., 2002; Hopfinger et al., 2000; Brefczynski & DeYoe, 1999; Martinez et al., 1999; Tootell et al., 1998; Woldorff et al., 1997). Note also that the contralaterality of these effects verifies that subjects did not merely “expand the spotlight of attention” to encompass the target areas, but rather shifted the spotlight to the target positions, as has also been verified in similar experiments using ERPs (reviewed in Hillyard, Mangun, Woldorff, & Luck, 1995). In addition, since the cue was centrally presented, regions that are contralateral or ipsilateral to the instructed shift of attention should have experienced the same activation in response to the sensory properties of the cue. Thus, the observed difference between contralateral and ipsilateral time-courses must be due to an effect of attentional orienting in response to the directional cue, rather than the sensory properties of the cue itself.

Although we and others (Hopfinger et al., 2000; Kastner et al., 1999) have interpreted cue-triggered activity in visual sensory cortices contralateral to the direction of attention as reflecting attentional biasing of those cortices in favor of expected target stimuli, other studies have either observed different results or have interpreted cue-triggered or cue-period activity in sensory cortices differently. For example, Corbetta et al. (2000) found transient activation of visual sensory cortex in response to cues, but they concluded that this effect might have been due simply to the sensory encoding of cue stimuli. Shulman et al. (1999) observed a pretarget transient increase in motion area MT+ in response to informative versus noninformative cues in a motion cueing task; however, they interpreted this effect as possibly reflecting “encoding and/or maintenance of instruction signals.” In a later study, Shulman et al. (2002) observed transient cue-period activity in both sensory and nonsensory cortices that appeared to be triggered by the EOT signal occurring 4–7 sec after the cue rather than by the cue itself. This result led to the proposal that certain brain regions were “reactivated” by the EOT signal due to the “turning-off” of the preparatory state that had been maintained while attending. In the present study, it is possible that the later portions of activity in visual sensory cortices during cue-only trials (Figure 6) could include some activity due to this kind of “reactivation” response triggered by the EOT signal. However, the time-courses contralateral and ipsilateral to the direction of attention (Figure 6) diverge as soon as hemodynamic activity begins, rather than being delayed 2.7 sec as would have been expected if it was associated with the EOT signal. In addition, it is difficult to understand how “reactivation” activity due to the ending of a preparatory state would be larger contralateral to the direction of attention, while the initial triggering of the preparatory state would not be contralateral. Thus, in the present study, we conclude that the centrally presented attend-cues triggered enhanced activity in a visual sensory area that was contralateral to the instructed direction of attention (the same

visual sensory area that would process a possible upcoming visual target stimulus if it were to occur), and that this cue-period activity is evident even in the absence of the target stimulus actually occurring. In accord with Hopfinger et al. (2000) and Kastner et al. (1999), we interpret this pattern as being consistent with an attention-induced, pretarget biasing of the visual sensory cortices.

Robust effects of attention were also seen in the thalamus. The attend-cue-plus-target trials showed a particularly large response, but even the attend-cue-only trials had enhanced activity relative to the interpret-cue trials. Notably, these effects in the thalamus were not contralateral to the direction of attention for the attend-cue-only trials, nor were they contralateral to the direction of the unilateral targets in the attend-cue-plus-target trials. In addition, they were not contralateral to the right-hand button-press. These results suggest an overall enhancement of both sides of the thalamus when cued to attend, with still additional activity bilaterally when processing a detected attended target, even when it is unilateral. Since these effects were not left-sided, and thus not contralateral to the right-hand button-presses, these activations seem unlikely to be due to motor preparation, although this cannot be completely ruled out.

Conclusions and Summary

In the present study, brain networks underlying the voluntary control of visual spatial attention were studied using a recently developed event-related fMRI approach in combination with a cued-attention paradigm characterized by having additional control conditions and the relatively short cue–target SOAs that are more similar to those often used in behavioral (and ERP) work. Even with these shorter cue–target SOAs, enhanced activity was seen in the frontal and parietal brain regions in response to the cues. This argues against the possibility that cue-related effects previously reported in these areas in long-SOA experiments were actually only the result of increased working memory demands or other cognitive operations induced by those long cue–target SOAs. Moreover, by including interpret-cue trials in the present study that were randomized into the sequence, we observed a pattern of activation that supports the view that there is functional specialization within the medial and lateral subregions of the parietal and frontal cortex. Activation in the more medial regions was relatively specific to the orienting of visual spatial attention, while that in the lateral regions was relatively general in nature, perhaps reflecting cognitive functions that occur for all cue stimuli, such as cue–symbol interpretation. Accordingly, we also hypothesize that it is the more medial frontal and parietal areas that control the observed retinotopically based biasing of activity in visual sensory cortex prior to the occurrence of an expected target.

METHODS

Participants

Twenty young adults (13 men and 7 women, mean age 27 years) were paid US\$20 per hour for their participation. Each one gave his/her informed consent prior to participating, in accordance with the rules of the Institutional Review Board at Duke University.

Design Rationale

At the short SOAs and fast stimulus rates needed for this study, overlap of the hemodynamic responses from adjacent events in the stimulus sequence is a major issue. To address this problem, we used a combination of trial-type randomization and hierarchical structuring of trial types.

Randomization of trial types is particularly useful for fast-rate event-related fMRI, in that it results in the overlapping responses due to adjacent trials in the sequence being about the same for the different trial types. Thus, contrasts between the trial types subtract out the overlap, leaving the event-related differential response activity (Buckner et al., 1998; Dale & Buckner, 1997; Woldorff, 1993). The fast-rate approach can be facilitated with the use of so-called no-stim events, points in time in the sequence that are randomized just as if they were a real stimulus event type, but during which no stimulus actually occurs (Buckner et al., 1998; Burock et al., 1998). Because no-stims then also contain the same overlap from adjacent trials, contrasts between them and the other trial types also subtract out this overlap, while also providing a low-level control condition for the various trial types. However, neither this randomization approach, in and of itself, nor the use of no-stims, solves the separation of the cue- and target-triggered activity, because the sequence of cues and targets cannot be randomized.

Thus, to separate the cue- and target-related activity, we combined the randomization approach (including no-stims) with a hierarchical structure of the trial types in which cues sometimes occur without targets. This approach also provided a natural mechanism for including an interpret-cue trial type within the hierarchy as a means of controlling for, and assessing, the contribution of cue-symbol processing and interpretation. (For previously reported alternative approaches using multiple regression, see Ollinger, Corbetta, et al., 2001; Ollinger, Shulman, et al., 2001; Hinrichs et al., 2000.)

Stimuli and Task

Participants were presented with a series of trials, each of which began with an instructional cue letter presented at fixation (Figure 1). The attention-directing instructional cues consisted of the letter L or R, which instructed the subject to attend for a possible visual target (a faint dot)

at a location in either the left or right lower visual field (3° lateral and 3° below horizontal meridian). The “interpret-cue” instructional trials began with the letter “P,” instructing the subject to not attend for a target on that trial. In some of the attention directing trials (attend-cue-plus-target trials), a target would occur at a randomized time (900 or 1900 msec) after the onset of the cue. In attend-cue-only (i.e., attend-left-cue-only or attend-right-cue-only) trials, as well as in interpret-cue trials, no target was presented, so that the brain response would be due to the cue only. Finally, “no-stim” trials (i.e., periods of fixation only) were randomized with these other trial types, so that the full event-related responses to the various trial types could be extracted using selective averaging (Buckner et al., 1998; Burock et al., 1998). In all trial types (other than no-stims), an end-of-trial (EOT) stimulus (the letters “REP”) was presented 2700 msec after trial onset. Participants were instructed to press a button after the EOT stimulus appeared if a target had been presented during that trial. The EOT stimulus was presented in all attend-cue-plus-target trials, attend-cue-only trials, and interpret-cue trials to equate sensory processing demands across conditions. Each trial lasted 4500 msec, and the sequence of trial types was randomized.

Note the hierarchical structure of the trial types, as illustrated in Table 1a. This hierarchical structure was designed to allow various key contrasts to be performed that could isolate different brain responses associated with specific processing components (Table 1b), while subtracting out the hemodynamic response overlap from adjacent trials in the sequence (cf. Ollinger, Corbetta, et al., 2001; Ollinger, Shulman, et al., 2001; Corbetta et al., 2000; Shulman, et al., 1999).

Image Acquisition

While the subjects were engaged in the visual attention task, functional magnetic resonance images of their brains were recorded with the GE 1.5-T Signa LX MRI system at the Brain Imaging and Analysis Center at Duke University. Eighteen T2*-weighted, echo-planar image slices, each 5 mm thick with 1-mm skip, were acquired. Slice direction was oblique axial (close to horizontal), set to be parallel to the anterior commissure–posterior commissure line, with the sixth slice through that line. Imaging parameters were TR = 1.5 sec, TE = 40 msec, and flip angle = 90° , with in-plane resolution of 64×64 (3.75×3.75 mm) and a field of vision of 24 cm. With a TR of 1.5 sec, three brain volumes were acquired in each 4.5-sec trial.

For most subjects, high-resolution T1-weighted 2-D structural scans were also acquired, using the same slice thickness and orientation as the functional images. These were then used for overlaying the functional activation maps. These also had a 24-cm field of vision, with an in-plane resolution of 256×256 (0.94×0.94 mm).

Image Analysis

Image preprocessing was performed using routines in SPM99 (Friston et al., 1995). The images were corrected for subject motion and for slice acquisition timing variation within the TR, normalized to the MNI template, and spatially smoothed using an 8-mm FWHM isotropic kernel. Voxelwise regression analyses were performed using SPM99 (Friston et al., 1995) with individual regressors containing the onset times for different trial types, as well as regressors for subject motion.

Voxelwise analyses

The procedures described above were followed by across-subject random-effects voxel-level analyses, which included contrasts between the responses to the different trial types. These included hierarchical contrasts (see Results) as well as contrasts between each trial type relative to the lowest level trial type, the no-stims. Additional contrasts included left versus right comparisons (e.g., left-attend-cue-only vs. right-attend-cue-only). In some cases, conditions were collapsed across for contrasting to another condition or group of conditions (e.g., the average response to left-attend and right-attend-cues vs. the response to interpret-cues) to increase statistical power.

Threshold values for the voxelwise statistical parametric maps were set at a t value of either 2.75 ($p < .012$, one-tailed) or 3.25 ($p < .006$, one-tailed), the latter being for higher power contrasts in which two conditions (e.g., left and right) were collapsed together. To reduce false positives from noise, these thresholds were used in conjunction with a cluster extent threshold of 11 contiguous voxels (Forman et al., 1995; Xiong, Gao, Lancaster, & Fox, 1995). In addition, results from these voxelwise analyses were confirmed by ROI analyses of the time-locked averaged responses, derived from selective averaging, as described below.

Region-of-Interest Analyses

ROIs were functionally defined based on the various contrasts, and additional functional analysis of these ROIs was performed using custom in-house software. In each ROI, we used selective averaging to extract the average hemodynamic response across all voxels in the ROI that was elicited by the various trial types (in terms of percent change from the prestimulus baseline). The time-locked average response to no-stim trials was then subtracted from the time-locked averaged responses to the other trial types, thereby subtracting off the hemodynamic overlap from adjacent responses and leaving the “pure” event-related transient response to each trial type (Burock et al., 1998; Dale & Buckner, 1997) within these ROIs. The average peak amplitudes (i.e., the average of two time points around the peak) of these

responses were contrasted with random effects t tests to determine whether and how distinct trial types differentially activated the various ROIs. In addition, the peak amplitude values across ROIs in the frontal and parietal cortex were analyzed using repeated measures ANOVA to directly test whether the medial and lateral regions of the frontal and parietal cortex were differentially involved in the processing of interpret-cues and attention-directing cue stimuli.

Conversion of Coordinates from MNI to Talairach Space

The neuroanatomical locations of activated areas were later converted from MNI space to Talairach space using a nonlinear transformation that has been previously described (<http://www.mrc-cbu.cam.ac.uk/Imaging/mnispac.html>). More specifically, coordinates above the anterior commissure were transformed by $x = 0.99x$, $y = 0.9688y + 0.0460z$, and $z = -0.0485y + 0.9189z$, and coordinates below the anterior commissure were transformed with: $x = 0.99x$, $y = 0.9688y + 0.0420z$, and $z = -0.0485y + 0.8390z$. Talairach coordinates in the Table 2 indicate the location(s) of peak activity within each region.

Acknowledgments

This work was supported by grants from NIMH (MH60415), NINDS (P01 NS41328, Proj. 2), and the ARO (DAAD 19-00-1-0530) to M.G.W.

Reprint requests should be sent to Marty G. Woldorff, Center for Cognitive Neuroscience, Duke University, Room B203, LSRC Building, Box 90999, Durham, NC 27708-0999, USA, or via e-mail: woldorff@duke.edu.

The data reported in this experiment have been deposited in the fMRI Data Center (<http://www.fmridc.org>). The accession number is 2-2003-11421.

REFERENCES

- Awh, E., Anllo-Vento, L., & Hillyard, S. A. (2000). The role of spatial selective attention in working memory for locations: Evidence from event-related potentials. *Journal of Cognitive Neuroscience*, *12*, 840–847.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*, 624–652.
- Brefczynski, J. A., & DeYoe, E. A. (1999). A physiological correlate of the ‘spotlight’ of visual attention. *Nature Neuroscience*, *2*, 370–374.
- Buckner, R. L., Goodman, J., Burock, M., Rotte, M., Koutstaal, W., Schacter, D., Rosen, B. R., & Dale, A. M. (1998). Functional-anatomic correlates of object priming in humans revealed by rapid presentation event-related fMRI. *Neuron*, *20*, 285–296.
- Burock, M. A., Buckner, R. L., Woldorff, M. G., Rosen, B. R., & Dale, A. M. (1998). Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. *NeuroReport*, *9*, 3735–3739.

- Carter, C. S., Botvinick, M. M., & Cohen, J. D. (1999). The contribution of the anterior cingulate cortex to executive processes in cognition. *Reviews in the Neurosciences*, *10*, 49–57.
- Corbetta, M. (1998). Frontoparietal cortical networks for directing attention and the eye to visual locations: Identical, independent, or overlapping neural systems? *Proceedings of the National Academy of Sciences, U.S.A.*, *95*, 831–838.
- Corbetta, M., Kincade, J. M., Ollinger, J. M., McAvoy, M. P., & Shulman, G. L. (2000). Voluntary orienting is dissociated from target detection in human posterior parietal cortex. *Nature Neuroscience*, *3*, 292–297.
- Corbetta, M., Miezin, S., Shulman, G. L., & Petersen, S. E. (1993). A PET study of visuospatial attention. *Journal of Neuroscience*, *13*, 1020–1026.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 215–229.
- Dale, A. M., & Buckner, R. L. (1997). Selective averaging of rapidly presented individual trials using fMRI. *Human Brain Mapping*, *5*, 329–340.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.
- D'Esposito, M. D., Aguirre, G. K., Zarahn, E., Ballard, D., Shin, R. K., & Lease, J. (1998). Functional MRI studies of spatial and nonspatial working memory. *Cognitive Brain Research*, *7*, 1–13.
- D'Esposito, M. D., Zarahn, E., & Aguirre, G. K. (1999). Event-related functional MRI: Implications for cognitive psychology. *Psychological Bulletin*, *125*, 155–164.
- De Fockert, J. W., Rees, G., Frith, C. D., & Lavie, N. (2001). The role of working memory in visual selective attention. *Science*, *291*, 1803–1806.
- Downing, C. J. (1988). Expectancy and visual-spatial attention: Effects on perceptual quality. *Journal of Experimental Psychology: Human Perception and Performance*, *14*, 188–202.
- Eimer, M. (2000). The time course of spatial orienting elicited by central and peripheral cues: Evidence from event-related brain potentials. *Biological Psychology*, *53*, 253–258.
- Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., & Noll, D. C. (1995). Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): Use of a cluster-size threshold. *Magnetic Resonance in Medicine*, *33*, 636–647.
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. P., Frith, C. D., & Frackowiack, R. S. J. (1995). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, *2*, 189–210.
- Gitelman, D. R., Parrish, T. B., LaBar, K. S., & Mesulam, M. M. (2000). Real-time monitoring of eye movements using infrared video-oculography during functional magnetic resonance imaging of the frontal eye fields. *Neuroimage*, *11*, 58–65.
- Hillyard, S. A., Mangun, G. R., Woldorff, M. G., & Luck, S. J. (1995). Neural systems mediating selective attention. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 665–681). Cambridge: MIT Press.
- Hinrichs, H., Scholtz, M., Tempelmann, C., Woldorff, M. G., Dale, A. M., & Heinze, H.-J. (2000). Deconvolution of event-related fMRI responses in fast-rate experimental design: Tracking amplitude variations. *Journal of Cognitive Neuroscience*, *12*, S76–S89.
- Hopf, J., & Mangun, G. R. (2000). Shifting visual attention in space: An electrophysiological analysis using high spatial resolution mapping. *Clinical Neurophysiology*, *111*, 1241–1257.
- Hopfinger, J. B., Buonocore, M. H., & Mangun, G. R. (2000). The neural mechanisms of top-down attentional control. *Nature Neuroscience*, *3*, 284–291.
- Jha, A. P., & McCarthy, G. (2000). The influence of memory load upon delay-interval activity in a working-memory task: An event-related functional MRI study. *Journal of Cognitive Neuroscience*, *12*, 90–105.
- Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*, *22*, 751–761.
- LaBar, K. S., Gitelman, D. R., Parrish, T. B., & Mesulam, M. M. (1999). Neuroanatomic overlap of working memory and spatial attention networks: A functional MRI comparison within subjects. *Neuroimage*, *10*, 695–704.
- MacDonald, A. W., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science*, *288*, 1835–1838.
- Martinez, A., Anillo-Vento, L., Sereno, M. I., Frank, L. R., Buxton, R. B., Dubowitz, D. J., Wong, E. C., Hinrichs, H., Heinze, H. J., & Hillyard, S. A. (1999). Involvement of striate and extrastriate visual cortical areas in spatial attention. *Nature Neuroscience*, *2*, 364–369.
- Milham, M. P., Banich, M. T., Webb, A., & Barad, V. (2001). The relative involvement of anterior cingulate and prefrontal cortex in attentional control depends on nature of conflict. *Cognitive Brain Research*, *12*, 467–473.
- Murtha, S., Chertkow, H., Beauregard, M., Dixon, R., & Evans, A. (1996). Anticipation causes increased blood flow to the anterior cingulate cortex. *Human Brain Mapping*, *4*, 103–112.
- Nobre, A. C., Gitelman, D. R., Dias, E. C., & Mesulam, M. M. (2000). Covert visual spatial orienting and saccades: Overlapping neural systems. *Neuroimage*, *11*, 210–216.
- Nobre, A. C., Sebestyen, G. N., Gitelman, D. R., Mesulam, M. M., Frackowiak, R. S., & Frith, C. D. (1997). Functional localization of the system for visuospatial attention using positron emission tomography. *Brain*, *120*, 515–533.
- Noesselt, T., Hillyard, S. A., Woldorff, M., Hagner, T., Jaencke, L., Tempelmann, C., Hinrichs, H., & Heinze, H.-J. (2002). Delayed striate cortical activation during spatial attention. *Neuron*, *35*, 575–587.
- Nystrom, L. E., Braver, T. S., Sabb, F. W., Delgado, M. R., Noll, D. C., & Cohen, J. D. (2000). Working memory for letters, shapes, and locations: fMRI evidence against stimulus-based regional organization in human prefrontal cortex. *Neuroimage*, *11*, 424–446.
- Ollinger, J. M., Corbetta, M., & Shulman, G. L. (2001). Separating processes within a trial in event-related functional MRI. *Neuroimage*, *13*, 218–229.
- Ollinger, J. M., Shulman, G. L., & Corbetta, M. (2001). Separating processes within a trial in event-related functional MRI. *Neuroimage*, *13*, 210–217.
- Paus, T. (2001). Primate anterior cingulate cortex: Where motor control, drive and cognition interface. *Nature Reviews Neuroscience*, *2*, 417–424.
- Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, *13*, 25–42.
- Posner, M. I., Snyder, C. R. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, *109*, 160–174.
- Rosen, B. R., Buckner, R. L., & Dale, A. M. (1998). Event-related functional MRI: Past, present, and future. *Proceedings of the National Academy of Sciences, U.S.A.*, *95*, 773–780.
- Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R., & Tootell, R. B. H. (1995).

- Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, *258*, 889–893.
- Shulman, G. L., Ollinger, J. M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Petersen, S. E., & Corbetta, M. (1999). Areas involved in encoding and applying directional expectations to moving objects. *Journal of Neuroscience*, *19*, 9480–9496.
- Talairach, J., Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain*. Thieme: New York.
- Tootell, R. B., Hadjikhani, N., Hall, E. K., Marrett, S., Vanduffel, W., Vaughan, J. T., & Dale, A. M. (1998). The retinotopy of visual spatial attention. *Neuron*, *21*, 1409–1422.
- Vandenberghe, R., Duncan, J., Dupont, P., Ward, R., Poline, J.-B., Bormans, G., Michiels, J., Mortelmans, & Orban, G. A. (1997). Attention to one or two features in left and right visual field: A positron emission tomography study. *Journal of Neuroscience*, *17*, 3739–3750.
- Van Veen, V., Cohen, J. D., Botvinick, M. M., Stenger, V. A., & Carter, C. S. (2001). Anterior cingulate cortex, conflict monitoring, and levels of processing. *Neuroimage*, *14*, 1302–1308.
- Wojciulik, E., & Kanwisher, N. (1999). The generality of parietal involvement in visual attention. *Neuron*, *23*, 747–764.
- Woldorff, M. G. (1993). Distortion of ERP averages due to overlap from temporally adjacent ERPs: Analysis and correction. *Psychophysiology*, *30*, 98–119.
- Woldorff, M. G., Fox, P. T., Matzke, M., Lancaster, J. L., Veeraswamy, S., Zamarripa, F., Seabolt, M., Glass, T., Goa, J. H., Martin, C. C., & Jerabek, P. (1997). Retinotopic organization of early spatial attention effects as revealed by PET and ERPs. *Human Brain Mapping*, *5*, 280–286.
- Woldorff, M. G., Matzke, M., Zamarripa, F., & Fox, P. T. (1999). Hemodynamic and electrophysiological study of the role of the anterior cingulate in target-related processing and selection for action. *Human Brain Mapping*, *8*, 121–127.
- Xiong, J., Gao, J., Lancaster, J. L., & Fox, P. T. (1995). Clustered pixels analysis for functional MRI activation studies in the human brain. *Human Brain Mapping*, *3*, 287–301.