

Top-down Enhancement and Suppression of the Magnitude and Speed of Neural Activity

Adam Gazzaley, Jeffrey W. Cooney, Kevin McEvoy,
Robert T. Knight, and Mark D'Esposito

Abstract

■ Top-down modulation underlies our ability to selectively attend to relevant stimuli and to ignore irrelevant stimuli. Theories addressing neural mechanisms of top-down modulation are driven by studies that reveal increased magnitude of neural activity in response to directed attention, but are limited by a lack of data reporting modulation of neural processing speed, as well as comparisons with a perceptual (passive view) baseline necessary to evaluate the presence of enhancement and suppression. Utilizing functional MRI (fMRI) and event-related potential recordings (ERPs), we

provide converging evidence that both the *magnitude* of neural activity and the *speed* of neural processing are modulated by top-down influences. Furthermore, both *enhancement* and *suppression* occur relative to a perceptual baseline depending on task instruction. These findings reveal the fine degree of influence that goal-directed attention exerts upon activity within the visual association cortex. We further document capacity limitations in top-down enhancement corresponding with working memory performance deficits. ■

INTRODUCTION

Human interaction with the environment involves a fluid integration of internally driven, goal-directed decisions concerning encountered stimuli (top-down modulation) and externally driven, perceptual influences that demand attention based on stimulus salience (bottom-up processes) (Bar, 2003; Frith, 2001). Top-down modulation thus underlies our ability to selectively focus on relevant stimuli and ignore distracting stimuli, establishing a foundation for attention and memory. The theoretical framework for the neural basis of top-down modulation relies on extensive evidence from single-cell physiology, functional neuroimaging, and electroencephalography (EEG) data, revealing increased activity in specialized posterior cortical regions—the presumed sites of neural representation—when attention is directed toward a stimulus or stimulus attribute (Pessoa, Kastner, & Ungerleider, 2003; Kastner & Ungerleider, 2001; Barcelo, Suwazono, & Knight, 2000; Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1990; Moran & Desimone, 1985; Hillyard, Hink, Schwent, & Picton, 1973). Despite the wealth of information characterizing this important phenomenon, basic mechanistic details of top-down modulation have not yet been completely elucidated, such as the temporal dynamics of activity modulation (i.e., the speed of cortical processing), the relationship of modulation to stimulus-driven,

perceptual activity, and neural correlates of limitations in top-down resources.

We tested three main hypotheses: (1) instruction to remember a stimulus results in modulation of both the magnitude of neural activity and the speed of cortical processing, (2) modulation occurs relative to a perceptual baseline, with both enhancement (above baseline) and suppression (below baseline) manifest depending on task instruction, and (3) top-down modulation is resource limited and, when capacity is exceeded, the degree of attainable enhancement is compromised. To explore these issues, we employed both event-related functional MRI (fMRI) and event-related potential recordings (ERP) to capitalize on the high spatial and temporal resolution attained when coupling these techniques. Eighteen participants performed the same experiment in two counterbalanced sessions while blood oxygen level dependent (BOLD) signal and electrical activity were recorded with fMRI and ERP, respectively. Although both measures are thought to reflect cortical activity driven by local cortical processing and the summation of postsynaptic potentials on synchronously active, large ensembles of neurons (Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001; Chawla, Lumer, & Friston, 1999; Silva, 1991), changes in BOLD signal can be localized to cortical regions separated by millimeters and ERP can resolve activity changes in the millisecond range. Thus, these techniques offer complementary but unique information to study the modulation of activity at the neuronal population level.

The experiment consisted of four tasks in which the participants viewed a series of four sequentially presented stimuli—two faces and two scenes—in a randomized order (Figure 1). The tasks differed in the instructions informing the participants on how to process the stimuli: (1) Remember Faces and Ignore Scenes, (2) Remember Scenes and Ignore Faces, (3) Remember Both faces and scenes (dual task), and (4) Passively View faces and scenes, without attempting to remember or evaluate them. All trials were followed by a delay period and, in the memory trials, a response period requiring them to report whether the stimulus matched one of the previously presented, relevant stimuli. It is important to note that all four tasks have equivalent perceptual, bottom-up information in the stimulus presentation stage, thus allowing us to selectively probe top-down processes.

Accordingly, only data from the stimulus presentation stage are discussed: for fMRI data, activity from the 4-sec block of stimuli, and for ERP data, waveforms time-locked to the onset of individual face stimuli. In addition, all data were obtained from select regions of interest (ROIs) that serve as our activity markers. For fMRI data, an independent functional localizer task was administered to identify regions in the visual association cortex (VAC) selective for faces (fusiform face area [FFA]; Kanwisher, McDermott, & Chun, 1997; Puce, Allison, Gore, & McCarthy, 1995) and scenes (parahippocampal place area [PPA]; Epstein & Kanwisher, 1998). The magnitudes of the BOLD signal within these ROIs were calculated and compared across tasks to address the specific hypotheses. For ERP data, limitations in spatial resolution make it impossible to measure activity specifically within the FFA and PPA. However, viewing faces induces a negative component localized to posterior occipital electrodes, the N170, which is specific for faces (Bentin, Allison, Puce, Perez, & McCarthy, 1996).

Although not reflecting activity limited to the FFA, it is a reliable measure of VAC activity with face specificity. Peak latency and peak amplitude of the N170 component generated in response to face stimuli were calculated and compared across tasks.

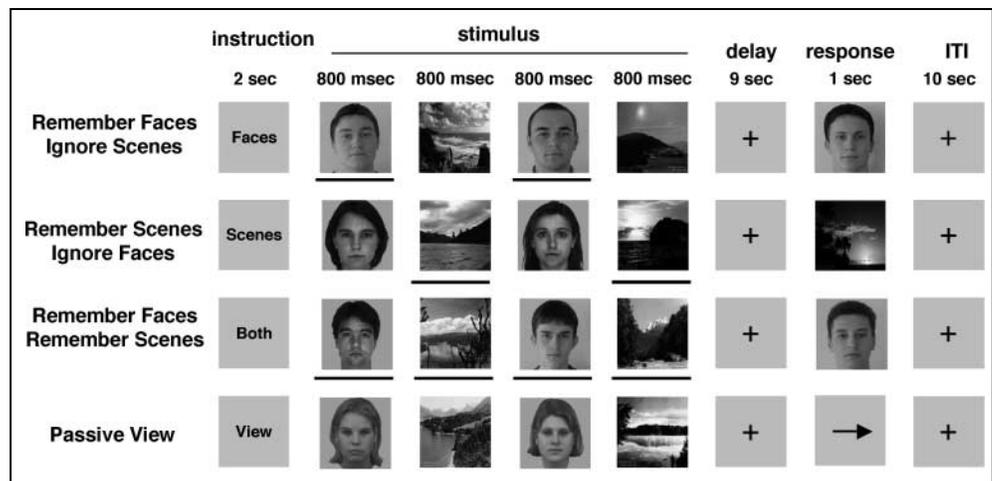
RESULTS

Magnitude and Speed of Neural Activity

In addressing our first hypothesis, analysis revealed that BOLD activity in the scene-selective PPA was greater for Remember Scenes than Ignore Scenes (left PPA and right PPA: $p < .00001$; Figure 2A and B) and activity in the right face-selective FFA was greater for Remember Faces than Ignore Faces ($p < .005$; Figure 2C and D). This activity difference within the stimulus-presentation stage must be driven by top-down modulation mediated by task instruction because the conditions were perceptually balanced. This effect was remarkably consistent, present in all 18 participants in the left PPA, 17 of 18 participants in the right PPA (Figure 3A), and 14 of 17 participants in the right FFA (Figure 3B). The left FFA was not consistently identifiable across subjects and so was not evaluated in the group analysis. Evidence that participants effectively ignored the stimuli when instructed to do so was obtained by way of a postexperiment recognition memory test, which revealed that faces and scenes to be ignored were less recognized than faces and scenes to be remembered (faces, $p < .001$; scenes, $p < .001$).

ERP analysis also revealed activity magnitude modulation, comparable to the fMRI FFA data, with an increase in the amplitude of the face-selective N170 component for Remember Faces versus Ignore Faces (bilateral N170, $p < .05$). There was no significant modulation

Figure 1. Experimental design of the fMRI and ERP tasks. Tasks differed in the instructions given at the beginning of each run and in the response requirements. Participants were instructed to (1) Remember Faces and Ignore Scenes, (2) Remember Scenes and Ignore Faces, (3) Remember Both Faces and Scenes (dual task), and (4) Passively View both Faces and Scenes, with no attempt to remember or evaluate them. In the memory trial response period, a face or scene stimulus was presented (depending on the condition), and participants were required to report with a button press whether the stimulus matched one of the previously presented stimuli. During the response period of the Passively View task, an arrow was presented and participants were required to make a button press indicating the direction of the arrow.



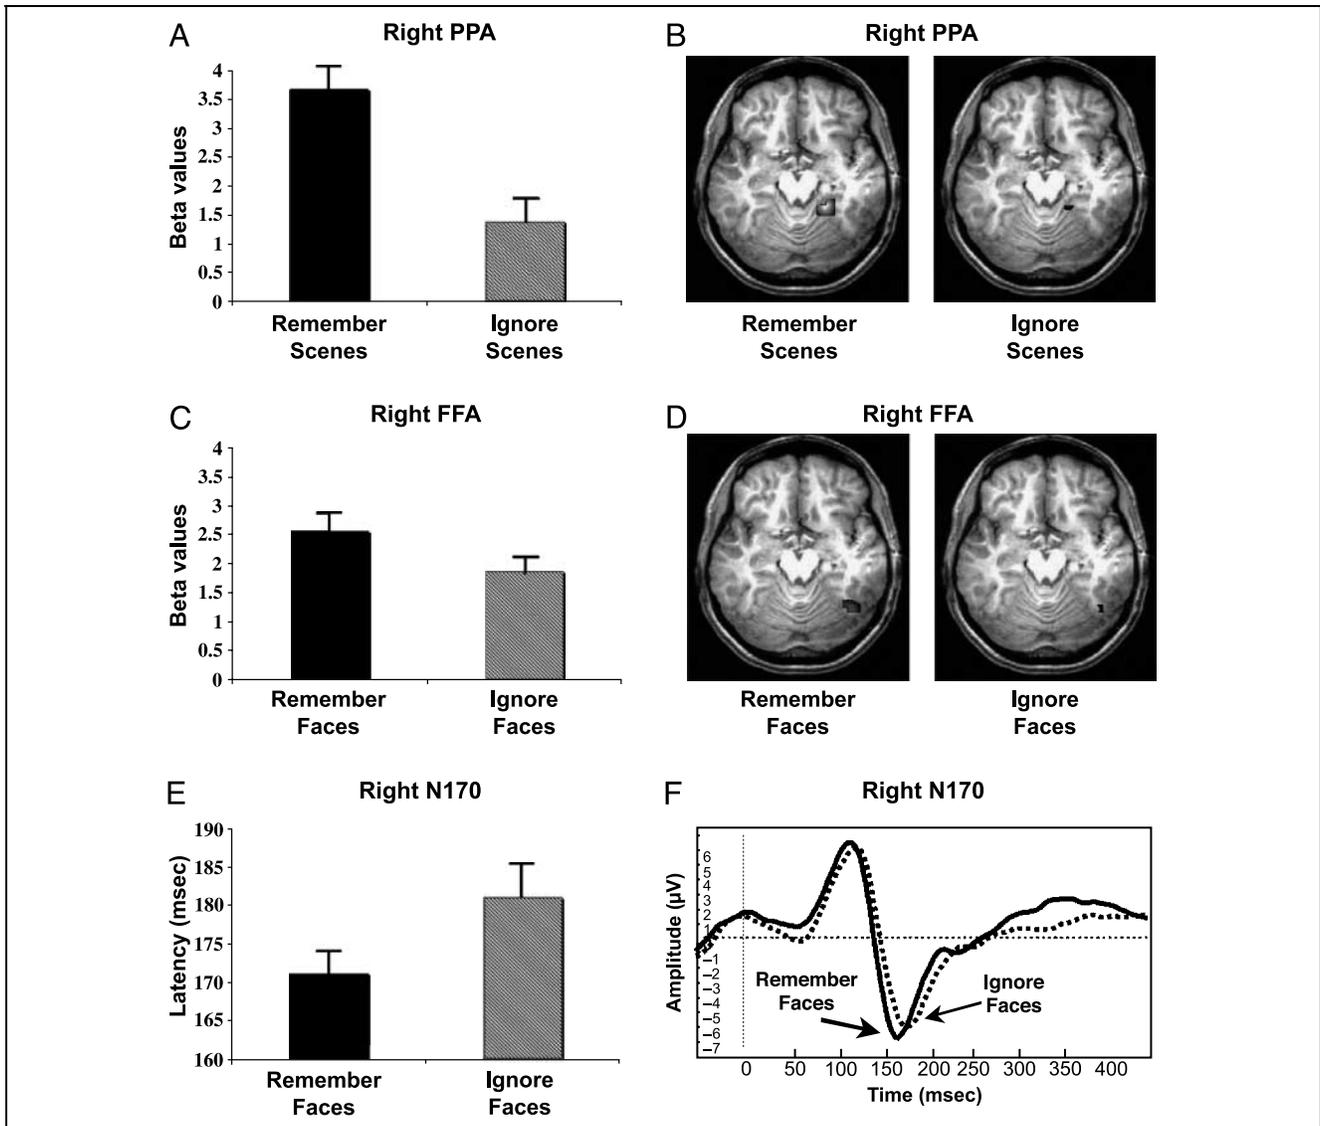


Figure 2. Activity data for Remember and Ignore conditions: fMRI and ERP. A and C: Group-average beta values in the right PPA and right FFA revealing greater activity in Remember conditions versus Ignore conditions. B and D: A representative subject demonstrating the BOLD signal level within the masked right PPA and right FFA in Remember conditions versus Ignore conditions. E: Group-average peak latency for the right N170 in PO8 electrode revealing earlier latency for Remember Faces versus Ignore Faces. F: Grand-averaged waveforms of the time-locked ERPs to face stimuli revealing earlier latency for Remember Faces versus Ignore Faces. Error bars indicate standard error of the mean.

of the P1 amplitude, possibly because this component is largely responsive to spatial attention (Hillyard & Anllo-Vento, 1998), which was not a factor in our paradigm. Because there was no identifiable scene-specific component and scene stimuli did not generate a consistent negative ERP in this latency range, waveforms generated in response to scenes were not further evaluated.

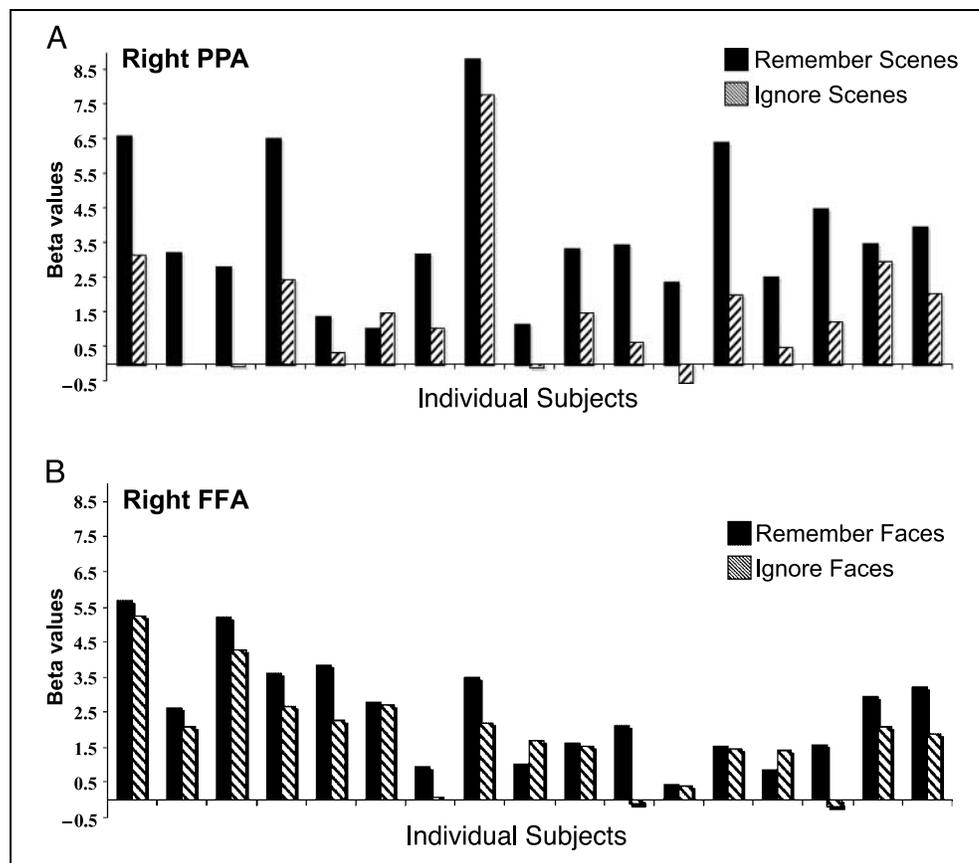
Notably, analysis of the ERP data also revealed a significant instruction-dependent shift in the right N170 peak latency (Figure 2E and F). The peak latency was, on average, 10 msec earlier for Remember Faces than Ignore Faces (171 vs. 181 msec, $p < .05$), a shift observed in 9 of 11 participants (Figure 4). No significant P1 latency shift was observed. N170 attention-dependent

latency modulation has not been previously reported and reveals that top-down influences can also modulate the time course of neural activity.

Enhancement and Suppression

We next investigated our second hypothesis that neural activity may be differentially enhanced or suppressed if a stimulus is, respectively, attended or ignored. The passive viewing task used in this experiment established a perceptual, bottom-up baseline from which activity in the remember tasks could be compared. Postexperiment recognition memory testing confirmed that passively viewed images were significantly less recognized

Figure 3. Individual subject data for Remember and Ignore conditions: fMRI. Activity in the right PPA for Remember Scenes and Ignore Scenes and the right FFA for Remember Faces and Ignore Faces for the *same* 17 participants. A: Average beta values in the right PPA reveals 16 of 17 participants have greater activity for Remember Scenes versus Ignore Scenes. B: Average beta values in the right FFA from the same participants reveal 14 of 17 have greater activity for Remember Faces versus Ignore Faces.



than those that were to be remembered during the Remember trials (faces, $p < .005$; scenes, $p < .001$).

fMRI data from the PPA demonstrated both significantly higher activity magnitude for Remember Scenes than Passively View (left, $p < .005$; right, $p < .05$; Figure 5A) and lower activity magnitude for Ignore Scenes than Passively View (left, $p < .001$; right, $p < .05$; Figure 5B), indicating the presence of both enhancement and suppression of activity relative to perceptual baseline. Enhancement was also observed in the FFA, as the activity was significantly higher for Remember Faces than Passively View ($p < .05$), but suppression effects did not reach significance (Figure 5C). In addition, ERP analysis of the N170 latency revealed the presence of both enhancement and suppression of processing speed associated with face stimuli, as the right N170 peak latency was earlier for Remember Faces than Passively View ($p < .05$) and later for Ignore Faces than Passively View ($p < .05$; Figure 5D).

Competition for Top-down Resources

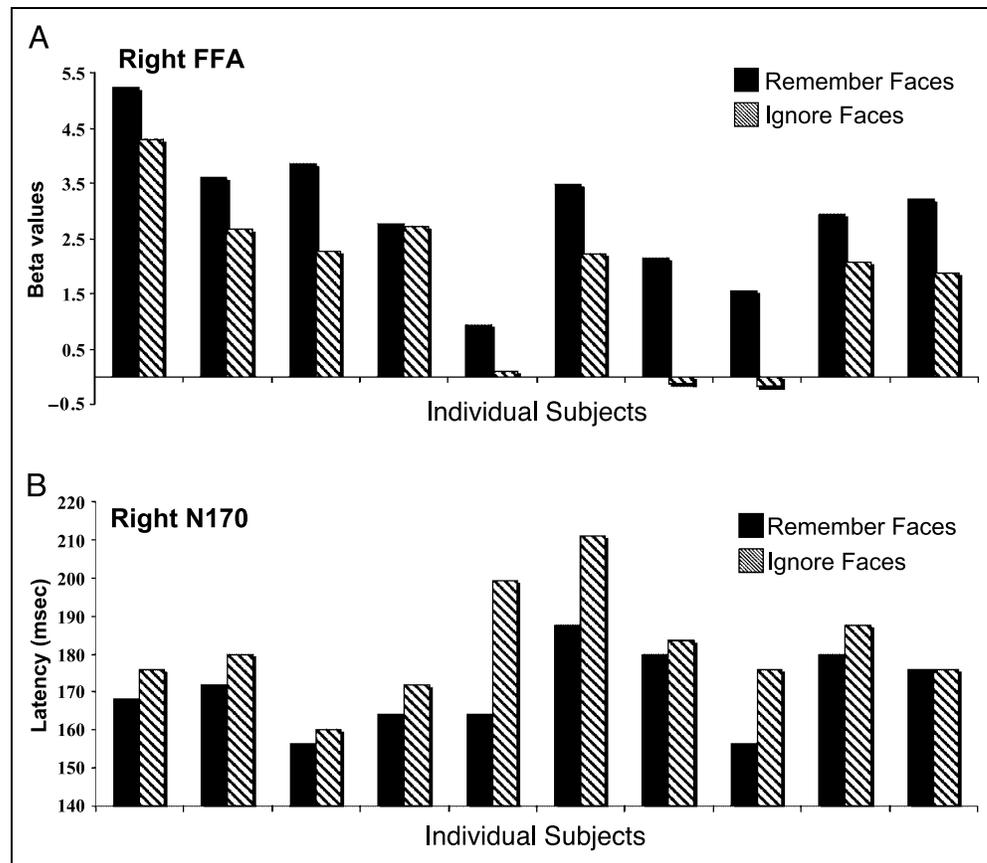
Lastly, we addressed our third hypothesis that top-down resources are limited and, when exceeded, modulation is compromised. Based on this hypothesis we predicted that instruction to remember both faces and scenes during the same trial (Remember Both—dual task)

would result in poorer working memory performance and diminished top-down enhancement compared to remembering faces or scenes alone. Supporting our hypothesis, we documented a performance deficit on the dual task trials and observed significant enhancement decrements for two of our activity markers, the left PPA magnitude (Figure 6A) and the right N170 latency (Figure 6B).

For the left PPA, the level of activity for Remember Both was not significantly elevated above Passively View ($p = .11$), indicating a lack of enhancement in the dual task above the perceptual baseline (Figure 6A). In addition, significantly lower BOLD activity was observed for Remember Both compared to Remember Scenes ($p < .05$), even though the same number of scenes were to be remembered in both tasks. This corresponded with a performance deficit in the fMRI working memory task for Remember Both compared to Remember Scenes (RT: 1452 vs. 1342 msec, $p < .001$; accuracy: 82% vs. 87%, $p < .05$). Post-experiment recognition testing also revealed that scenes to be remembered in Remember Both were less recognized than scenes to be remembered in Remember Scenes ($p < .005$).

Complementary evidence of both diminished enhancement and performance deficits in the dual task trials was present in the ERP experiment (Figure 6B). No significant N170 peak latency difference was evident

Figure 4. Individual subject data for Remember and Ignore conditions: fMRI & ERP. Activity in the right FFA for Remember Faces and Ignore Faces and the right N170 (PO8 component) for Remember Faces and Ignore Faces in the same 10 participants with sufficient data in these ROIs. A: Average beta values in the right FFA reveal 9 of 10 participants have greater activity for Remember Faces versus Ignore Faces. B: Average N170 peak latency values from the same participants reveals 9 of 10 have an earlier latency for Remember Faces versus Ignore Faces.



between Remember Both and Passively View (177 vs. 178 msec, $p = .83$), such as existed for Remember Faces and Passively View (171 vs. 178 msec, $p < .005$), indicating a lack of enhancement in the dual task above the perceptual baseline. Accordingly, the N170 peak latency for Remember Both was not as early as the latency for Remember Faces (177 vs. 171 msec, $p < .005$). This is a remarkable finding considering that the ERP data *only* reflect activity in response to faces, and in both tasks the participants were remembering faces. The only difference between the two tasks was that during the Remember Both condition, participants had also been instructed to remember scenes during the same trial. Again, a corresponding working memory performance deficit was observed in the ERP working memory task, with performance on Remember Both impaired relative to Remember Faces (RT: 1263 vs. 1105 msec, $p < .001$; accuracy: 91% vs. 97%, $p < .05$).

DISCUSSION

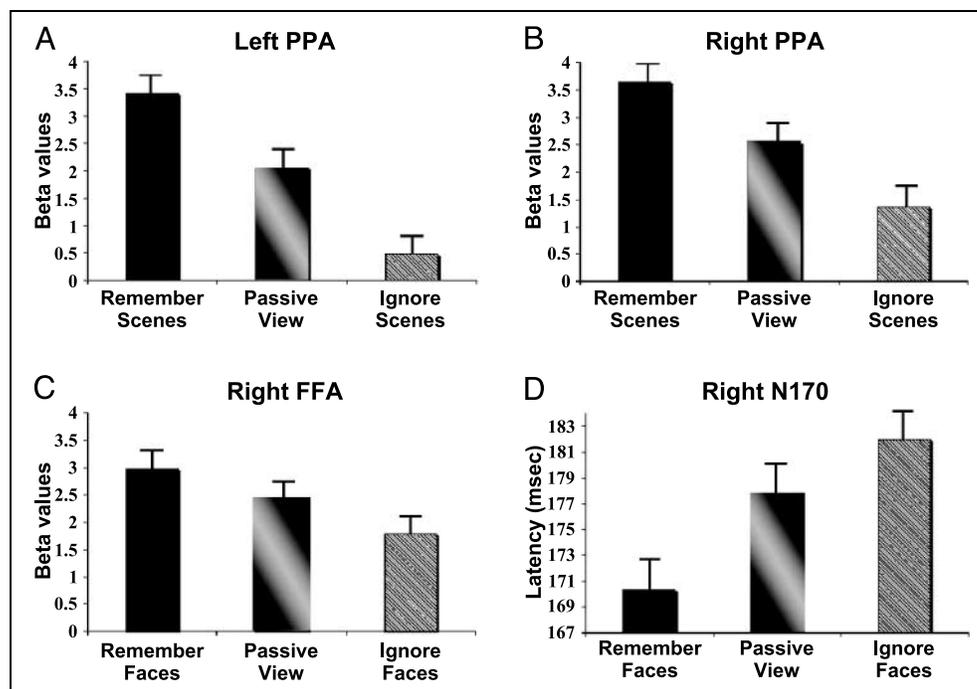
Top-down modulation of the magnitude of neural activity, as revealed through diverse experimental paradigms and physiological techniques, has represented the foundation for our understanding of the mechanism of goal-directed control of cortical activity (Pessoa et al., 2003; Barcelo et al., 2000; Corbetta et al., 1990; Moran &

Desimone, 1985; Hillyard et al., 1973). In this study, we extend the mechanistic basis of top-down modulation in three regards by revealing that (1) both activity magnitude and the speed of processing are modulated by top-down influences, (2) modulation occurs both above (enhancement) and below (suppression) a perceptual baseline, and (3) modulation of both activity magnitude and speed of processing exhibits capacity limitations that correspond with impaired working memory performance.

Magnitude and Speed of Neural Activity

Our fMRI and ERP evidence of greater VAC activity magnitude for the remember conditions versus the ignore conditions is consistent with reports from previous studies. fMRI experiments have revealed activity magnitude differences between attend and ignore conditions within category-specific VAC (Vuilleumier, Armony, Driver, & Dolan, 2001; O'Craven, Downing, & Kanwisher, 1999; Wojciulik, Kanwisher, & Driver, 1998) and ERP studies have reported attention-mediated amplitude changes in the N170 (Holmes, Vuilleumier, & Eimer, 2003; Eimer, 2000), as well as other early event-related components such as the P1 and N1 (Hillyard & Anllo-Vento, 1998). However, studies of selective attention frequently assess activity modulation by using task

Figure 5. Activity data for Remember, Passively View, and Ignore conditions: fMRI and ERP. A and B: Group-average beta values reveal greater activity in the right and left PPA for Remember Scenes versus Passively View and lower activity for Ignore Scenes versus Passively View. C and D: Group data for the right FFA and the right N170 (P08 electrode) from the same 10 participants with sufficient data in both fMRI and ERP experiments. C: Group-average beta values reveal greater activity in the right FFA for Remember Faces versus Passively View and a trend (not significant) to lower activity for Ignore Faces versus Passively View. D: Group-average peak latency of the N170 in P08 electrode reveals an earlier latency for Remember Faces versus Ignore Faces and a later latency for Ignore Faces versus Passively View. Error bars indicate standard error of the mean.



relevancy to resolve competition between multiple *simultaneously* presented stimuli (Vuilleumier et al., 2001; O'Craven et al., 1999; Wojciulik et al., 1998), and so modulation is inferred to be driven by competition for limited perceptual processing resources. In contrast, our findings reveal that modulation of activity magnitude can occur based on differential attention to *sequentially* presented stimuli, without competition for perceptual processing resources. In our design, task relevancy was used to resolve competition for limited working memory resources. This finding supports a more general model of top-down modulation in which goal-directed decisions modulate activity levels to resolve competition for limited resources, whether perceptual or mnemonic (Ranganath, DeGutis, & D'Esposito, 2004; Luck, Chelazzi, Hillyard, & Desimone, 1997).

Our ERP analyses also revealed a latency modulation of the N170, suggesting that in addition to modifying activity magnitude, top-down influences modulate the speed of neural processing as reflected by a shorter time to reach maximal synchronized neural activity. An ERP latency shift reflects the summed contribution of large ensembles of neurons in the VAC, a population measure of processing speed that may not be readily observable at the single-neuron level (Silva, 1991). This may account for the lack of latency differences in attention studies using single-cell-recording techniques, which often record spike activity and not local field potentials. The

absence of a documented N170 latency shift in previous ERP selective attention studies (Holmes et al., 2003; Eimer, 2000) suggest that our findings may either be specific for attention directed at memory encoding or driven by a task demanding greater attention.

It has been proposed that amplification of activity magnitude improves the signal-to-noise ratio, allowing more information to be extracted from relevant stimuli (Hillyard, Vogel, & Luck, 1998). Likewise, faster processing speed reflects an augmentation in the efficiency of neural processing, further facilitating information extraction. By modulating activity magnitude and processing speed, top-down signals bias the likelihood of successful representation of relevant information in a competitive system.

Enhancement and Suppression

Inherent to theories of top-down modulation is the concept that neural activity is modulated relative to a level of activity generated by the bottom-up, perceptual influence of a stimulus. Neural activity in response to viewing a stimulus may be differentially enhanced or suppressed relative to this level of activity if it is attended or ignored. Despite this logic, modulation relative to a stimulus-present baseline has rarely been evaluated, and comparisons are usually made between

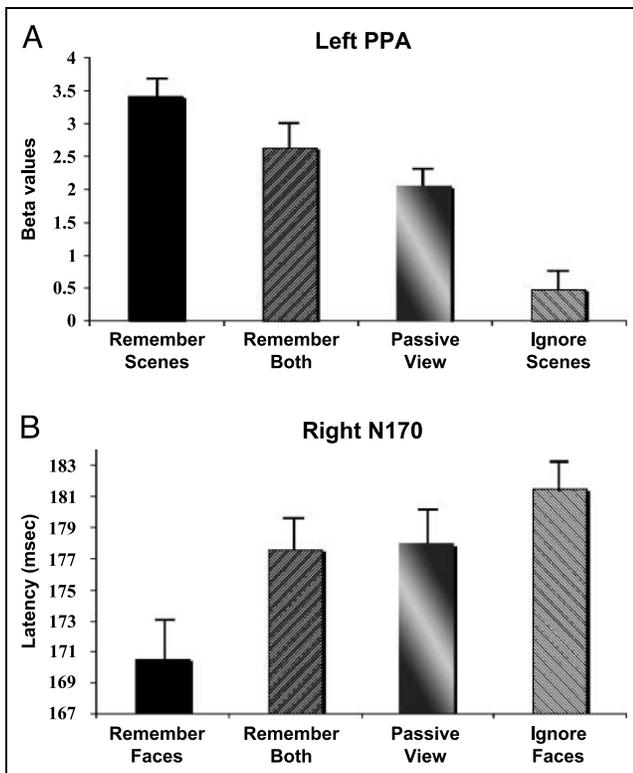


Figure 6. Activity data for Remember Both, Remember, Passively View and Ignore conditions: fMRI and ERP. A: Group-average beta values in the left PPA reveal no significant difference between Remember Both and Passively View and greater activity for Remember Scenes versus Remember Both, and so there is no significant enhancement when both faces and scenes were to be remembered. B: Group-average peak latency of N170 in P10 electrode reveals no significant difference between Remember Both and Passively View and an earlier latency for Remember Faces versus Remember Both, thus no significant enhancement exists for faces when scenes were also remembered during the same trial. Error bars indicate standard error of the mean.

attend and ignore tasks or relative to a resting baseline without visual stimulation (Pinsk, Doniger, & Kastner, 2004; Holmes et al., 2003; Vuilleumier et al., 2001; Eimer, 2000; O’Craven et al., 1999; Wojciulik et al., 1998; Rees, Frith, & Lavie, 1997). Without establishing a perceptual baseline level of activity, it is not possible to interpret top-down influences as representing enhancement or suppression. Although modulation both above and below a perceptual baseline has not yet been reported with neuroimaging data, the presence of enhancement and suppression has been suggested in ERP studies of spatial attention, which document a decreased amplitude of the P1 component for ignored locations and an increased amplitude of the N1 component for attended locations, both relative to a baseline obtained with “neutral” trials (Luck & Hillyard, 1995; Luck et al., 1994). Our fMRI and ERP data reveal that top-down modulation of activity magnitude and processing speed occur both above and below a perceptual baseline depending on task instruc-

tion. This combination of enhancement and suppression provides a contrast for sculpting neural processes.

Competition for Top-down Resources

Behavioral studies suggest that many cognitive processes have limited capacity, such that increased demands (e.g., dual task and load paradigms) result in diminished performance (Kahneman, 1973). Both activity and working memory performance data from our fMRI and ERP experiments converge to suggest that the increased memory load in the dual task condition exceeds the resources needed to complete the task on par with the simpler memory tasks, and is associated with a reduction in the level of attainable top-down enhancement. This finding was obtained by holding the perceptual load constant and varying only the mnemonic load, differentiating this study from others that have explored capacity limitations of working memory by varying perceptual load (Todd & Marois, 2004; Vogel & Machizawa, 2004). In addition to revealing a neural correlate of limitations in top-down resources, these data imply a link between top-down enhancement generated during working memory encoding and subsequent memory performance. Given that stimuli were presented sequentially rather than simultaneously, the enhancement deficit is likely the result of competition for limited top-down or mnemonic resources, rather than competition for limited perceptual processing resources as proposed by the biased competition model for simultaneously presented stimuli (Desimone & Duncan, 1995).

Although this study focused exclusively on a site of top-down modulation—the VAC—future analysis employing multivariate techniques may help identify cortical regions differentially associated with enhancement and suppression in the frontal and parietal cortices, the presumed sources of such modulation (Pessoa et al., 2003; Barcelo et al., 2000). These data serve to extend our understanding of the mechanisms of top-down modulation as well as to establish important biological markers to investigate limitations in top-down enhancement and/or suppression, which may underlie memory and attention impairments associated with normal aging and neurological disease.

METHODS

Participants

Eighteen participants (10 men) with ages ranging from 19 to 30 years (average = 23.3, *SD* = 2.4) participated in both the fMRI and ERP experiments after providing informed consent. Participants were prescreened, and none used any medication with psychoactive, cardiovascular, or homeostatic effects. All participants had normal-to-corrected vision and were right handed, although two of the participants used their left hands for some activities.

Stimulus Material

Grayscale images of faces and natural scenes were used as stimuli. All images were novel across all runs of both the fMRI and ERP experiments. Stimuli were 225 pixels wide by 300 pixels tall, and subtended approximately 5 by 6 degrees of visual angle. Although scene stimuli had identical attributes in each experiment, the face stimuli used in the fMRI experiment were blurred along the contours of the faces, so that only the faces themselves were visible. Faces used in the ERP experiment did not have blurred contours. Both male and female faces were used, although the sex of the face stimuli used within each trial was held constant. Only faces with neutral expressions were used as stimuli.

Experiment (General)

All participants were read the same set of instructions and shown the same sample trials to help explain the task. The experiment consisted of four tasks, presented in an event-related design. In all tasks, participants viewed four sequentially presented images (2 faces and 2 scenes). Each image was presented for 800 msec, with a 200-msec blank-screen interstimulus interval. Faces and scenes were presented in a randomized order. The tasks differed in the instructions given at the beginning of each run. Participants were instructed to (1) Remember Faces and Ignore Scenes, (2) Remember Scenes and Ignore Faces, (3) Remember Both Faces and Scenes (dual task), and (4) Passively View both Faces and Scenes—with no attempt to remember or evaluate them. Presentation of stimuli was followed by a 9-sec delay period, after which a fifth stimulus was presented for 1 sec. In the memory trials, a face or scene stimulus was presented (depending upon the condition), and participants were required to report with a button press whether the stimulus matched one of the previously presented stimuli. In the Passively View task, an arrow was presented, and participants were required to make a button press indicating the direction of the arrow. Presentation of the probe stimulus was followed by a 10-sec ITI in the fMRI experiment, and by a 3-sec ITI in the ERP experiment.

Experiment (fMRI Specific)

The general protocol described above was followed for all participants performing the fMRI sessions. Participants performed the experiment while lying on the scanner bed in the Varian 4T scanner. Head movement was restricted by using foam cushioning. Stimuli were back-projected onto a screen located at the subjects' waist (~90 cm from the subjects' eyes), and were viewed via a mirror located inside the head coil. Data were acquired during 12 runs lasting 4.5 minutes each, with a total of 120 trials.

Experiment (ERP Specific)

The general protocol described above was followed. Participants were seated in an armchair in a dark, sound-attenuated room and were monitored by camera during all tasks. The screen was approximately 125 cm from the subjects' eyes. Data were recorded during 12 runs lasting 5.5 minutes each, for a total of 240 trials.

Post-experiment Recognition Testing

Following completion of the fMRI experiment, participants were given an unexpected memory test to assess recognition of stimuli presented during the various experimental conditions. Participants were presented with a sequence of 240 images (120 novel, 120 previously viewed), and were asked to indicate on a 4-point scale whether each stimulus was novel or previously viewed, as well as how confident they were with each judgment. All previously viewed stimuli used in the postexperiment recognition test had been presented only once during the course of the experiment (i.e., stimuli that were presented as both cue and probe were omitted), and represented a balanced combination of stimuli taken from each task condition. Only stimuli presented during the fMRI experiment were employed in the recognition testing.

fMRI Acquisition and Processing

MR data were acquired with a Varian INOVA 4T scanner (Palo Alto, CA) equipped with a transverse electromagnetic (TEM) send-and-receive radio frequency (RF) head coil. Functional data were obtained using a two-shot T2*-weighted echo-planar imaging (EPI) sequence sensitive to BOLD contrast (TR = 2000 msec, TE = 28 msec, FOV = 22.4 cm², matrix size = 64 × 64, in-plane resolution = 3.5 × 3.5 mm). Each functional volume consisted of eighteen 5-mm axial slices separated by a 0.5-mm interslice gap and provided nearly whole-brain coverage. Two T1-weighted anatomical scans were also acquired. In the first, anatomical images coplanar with the EPI data were collected using a gradient-echo multislice (GEMS) sequence (TR = 200 msec, TE = 5 msec, FOV = 22.4 cm², matrix size = 256 × 256, in-plane resolution = 0.875 × 0.875 mm). In the second, high-resolution anatomical data were acquired with an MP-FLASH 3-D sequence (TR = 9 msec, TE = 5 msec, FOV = 22.4 × 22.4 × 19.8 cm, matrix size = 256 × 256 × 128, resolution = 0.875 × 0.875 × 1.54 mm).

Following acquisition, MRI data were converted to ANALYZE format. Data were corrected for between-slice timing differences using a sinc interpolation method and were interpolated to 1-sec temporal resolution (half of the total repetition time) by combining each shot of half *k* space with the bilinear interpolation of the two flank-

ing shots. Subsequent processing was performed using SPM2 software (<http://www.fil.ion.ucl.ac.uk>) run under Matlab 6.5 (www.mathworks.com). Functional data were realigned to the first volume acquired and were spatially smoothed with a 3-D Gaussian kernel (8 mm full width half maximum).

fMRI Data Analysis

A detailed description of the methods used for the analysis of brain activation (i.e., BOLD signal) within a trial have been presented elsewhere (Postle, Zarahn, & D'Esposito, 2000) and are summarized below. Changes in BOLD signal were modeled with independent regressors for each stage of each condition of the cognitive tasks (e.g., Remember Faces—stimulus-presentation period). Because the stimulus-presentation period of each trial lasted 4 sec, the regressor for this period consisted of a 4-sec boxcar function. Only trials with correct responses were incorporated in the analysis; trials with incorrect responses were modeled separately and excluded. Each regressor was convolved with the canonical hemodynamic response function (HRF) provided by SPM2, with the result entered into the modified generalized linear model (Worsley & Friston, 1995) instantiated in SPM2. Data were proportionally scaled on a voxelwise basis to account for differences in mean signal values and were corrected for linear signal drift within scanning runs. A high-pass filter (cutoff period = 128 sec) was applied to remove low-frequency effects from the data. Maps of parameter estimates (β values) were computed from the generalized linear model to assess the magnitude of activation during each stage of the task. The mean parameter estimate of each regressor was then calculated within functionally defined ROIs for each subject. Statistical analyses with planned comparisons (two-sided, paired-sample *t* tests) were used to test significance of differences ($p < .05$) between the parameter estimates from the tasks for a priori hypotheses.

fMRI ROI Selection

ROIs representing face- and scene-selective voxels were functionally defined for each subject using an ROI-localizer task. Before beginning the memory experiments, subjects performed a 1-back task in which they were presented with seven 16-sec blocks of face stimuli, scene stimuli, and rest. Subjects were instructed to attend to the stimuli and to indicate when each 1-back match occurred by depressing a button with each forefinger. MR data were processed as described above, and the face- and scene-stimuli regressors were contrasted to generate SPM[T] images. A face-selective ROI (FFA) was then identified within the right fusiform gyrus of each subject. A face-selective region within the left fusiform gyrus was weak or absent in many subjects and was not included in the group analysis. Scene-selective

ROIs (PPAs) were identified bilaterally within the parahippocampal gyrus and/or anterior lingual gyrus. Each ROI was defined as the cluster of seven contiguous voxels (within the predetermined anatomical structures) with the highest *t* values from the face–scene or scene–face contrast, respectively. On the rare occasions when the seven voxels with the highest *t* values did not form one contiguous cluster, the ROI was defined as the seven contiguous voxels with the highest mean *t* value. PPAs were identified bilaterally for all 18 subjects, and right FFAs were identified in 17 of the 18 subjects. One subject was excluded from the FFA analysis due to lack of an identifiable face-selective region within the fusiform gyrus.

EEG Recording

EEG was recorded simultaneously with a sampling rate of 256 Hz from 64 scalp locations using Ag–AgCl-tipped electrodes attached to an electrode cap using the 10/20 labeling system. Recordings were made with the Active-Two system (BioSemi, Amsterdam, Netherlands), which does not require impedance measurements, an online reference, or gain adjustments. Additional channels recorded the electrooculogram (EOG) from the infra-orbital region of the right eye, the outer canthus of the right and left eye, and the right and left mastoid process. Both the EEG and the EOG were referenced off-line to an electrode placed on the tip of the nose.

ERP Analysis

Segments were time-locked to the onset of face and scene stimuli automatically recorded by triggers sent to the recording system. ERPs were segmented separately for face and scene stimuli. All data presented in this article are for responses time-locked to faces. Artifacts due to eye blinks or movements were excluded from analysis by using an automatic rejection algorithm that removed segments with an 80 μ V maximum difference over a 150-msec period for vertical eye movements (EOG), 70 μ V maximum difference over a 100-msec period for horizontal eye movements (EOG), a 110- μ V maximum difference over 200 msec for all channels (EEG), or for an amplitude greater than ± 100 μ V (EEG). Only recordings from correct trials were included in the analysis. Epochs without artifacts beginning 100 msec before a stimulus and ending 800 msec after were averaged, filtered using a band pass of 1 Hz to 20 Hz (12 dB/octave), and baseline corrected. Analysis of the data with several different levels of filtering (0.1–30, 0.1–70, 0.5–30 and 0.5–70 Hz) determined that the latency of the N170 was unaltered by the level of filtering. Peak latency and amplitude magnitudes were collected for all cap electrodes by automatically selecting the most positive peak between 75 and 140 msec (P1), and the most negative peak between 140 and 205 msec (N170). Sub-

jects with an average of less than 45 segments of a single type (e.g., face stimuli during correct trials) were excluded from averaging, leaving 11 subjects (6 men) with sufficient usable data (average age = 23.7, $SD = 3.4$). The average number of segments left for all subjects was 80.2 ($SD = 22.2$). Three of the removed subjects had large alpha waves resulting in too few usable segments or unidentifiable P1 and N170 peaks; the other 4 removed subjects had too few segments from data contaminated by a large number of blinks.

N170 peak latency and peak amplitude analyses focused on P10 and P08 electrodes in the right hemisphere and P9 and P07 electrodes in the left hemisphere because they exhibited the largest N170 amplitude regardless of task. Accordingly, the P1 analyses also focused on these electrodes because P1 also exhibited the highest amplitude in these electrodes. Statistical analyses with planned comparisons (two-sided, paired-sample t tests) were used to test significance of differences ($p < .05$) between the peak latency and the peak amplitude between the tasks in a hypothesis-driven manner.

Acknowledgments

This study was supported by National Research Service Awards (NRSA) and American Federation for Aging Research (AFAR) fellowships (A. G.), and NIH NS21135 (R. K.), NS 40813 (R. K and M. D.), and MH63901 (M. D.). We thank D. Pino, J. Rissman, and L. Deouell for helpful discussions and suggestions and J. Kelley for assistance with ERP data analysis.

Reprint requests should be sent to Adam Gazzaley, University of California, 32 Barker Hall, Berkeley, CA 94720-1650, or via e-mail: adamgazz@comewander.com.

The data reported in this experiment have been deposited in the fMRI Data Center (www.fmridc.org). The accession number is 2-2004-117NB.

REFERENCES

- Bar, M. (2003). A cortical mechanism for triggering top-down facilitation in visual object recognition. *Journal of Cognitive Neuroscience*, *15*, 600–609.
- Barcelo, F., Suwazono, S., & Knight, R. T. (2000). Prefrontal modulation of visual processing in humans. *Nature Neuroscience*, *3*, 399–403.
- Bentin, S., Allison, T., Puce, A., Perez, E., & McCarthy, G. (1996). Electrophysiological studies of face perception in humans. *Journal of Cognitive Neuroscience*, *8*, 551–565.
- Chawla, D., Lumer, E. D., & Friston, K. J. (1999). The relationship between synchronization among neuronal populations and their mean activity levels. *Neural Computation*, *11*, 1389–1411.
- Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1990). Attentional modulation of neural processing of shape, color, and velocity in humans. *Science*, *248*, 1556–1559.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.
- Eimer, M. (2000). Attentional modulations of event-related brain potentials sensitive to faces. *Cognitive Neuropsychology*, *17*, 103–116.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*, 598–601.
- Frith, C. (2001). A framework for studying the neural basis of attention. *Neuropsychologia*, *39*, 1367–1371.
- Hillyard, S. A., & Anillo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. *Proceedings of the National Academy of Sciences, U.S.A.*, *95*, 781–787.
- Hillyard, S. A., Hink, R. F., Schwent, V. L., & Picton, T. W. (1973). Electrical signs of selective attention in the human brain. *Science*, *182*, 177–180.
- Hillyard, S. A., Vogel, E. K., & Luck, S. J. (1998). Sensory gain control (amplification) as a mechanism of selective attention: Electrophysiological and neuroimaging evidence. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *353*, 1257–1270.
- Holmes, A., Vuilleumier, P., & Eimer, M. (2003). The processing of emotional facial expression is gated by spatial attention: Evidence from event-related brain potentials. *Brain Research, Cognitive Brain Research*, *16*, 174–184.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice Hall.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17*, 4302–4311.
- Kastner, S., & Ungerleider, L. G. (2001). The neural basis of biased competition in human visual cortex. *Neuropsychologia*, *39*, 1263–1276.
- Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*, 150–157.
- Luck, S. J., & Hillyard, S. A. (1995). The role of attention in feature detection and conjunction discrimination: An electrophysiological analysis. *International Journal of Neuroscience*, *80*, 281–297.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, *77*, 24–42.
- Luck, S. J., Hillyard, S. A., Mouloua, M., Woldorff, M. G., Clark, V. P., & Hawkins, H. L. (1994). Effects of spatial cuing on luminance detectability: Psychophysical and electrophysiological evidence for early selection. *Journal of Experimental Psychology. Human Perception and Performance*, *20*, 887–904.
- Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, *229*, 782–784.
- O'Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, *401*, 584–587.
- Pessoa, L., Kastner, S., & Ungerleider, L. G. (2003). Neuroimaging studies of attention: From modulation of sensory processing to top-down control. *Journal of Neuroscience*, *23*, 3990–3998.
- Pinsk, M. A., Doniger, G. M., & Kastner, S. (2004). Push-pull mechanism of selective attention in human extrastriate cortex. *Journal of Neurophysiology*, *92*, 622–629.
- Postle, B. R., Zarahn, E., & D'Esposito, M. (2000). Using event-related fMRI to assess delay-period activity during performance of spatial and nonspatial working memory tasks. *Brain Research. Brain Research Protocols*, *5*, 57–66.

- Puce, A., Allison, T., Gore, J. C., & McCarthy, G. (1995). Face-sensitive regions in human extrastriate cortex studied by functional MRI. *Journal of Neurophysiology*, *74*, 1192–1199.
- Ranganath, C., DeGutis, J., & D'Esposito, M. (2004). Category-specific modulation of inferior temporal activity during working memory encoding and maintenance. *Brain Research, Cognitive Brain Research*, *20*, 37–45.
- Rees, G., Frith, C. D., & Lavie, N. (1997). Modulating irrelevant motion perception by varying attentional load in an unrelated task. *Science*, *278*, 1616–1619.
- Silva, L. d. (1991). Neural mechanisms underlying brain waves: From neural membranes to networks. *Electroencephalography and Clinical Neurophysiology*, *79*, 81–93.
- Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature*, *428*, 751–754.
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*, 748–751.
- Vuilleumier, P., Armony, J. L., Driver, J., & Dolan, R. J. (2001). Effects of attention and emotion on face processing in the human brain: An event-related fMRI study. *Neuron*, *30*, 829–841.
- Wojciulik, E., Kanwisher, N., & Driver, J. (1998). Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study. *Journal of Neurophysiology*, *79*, 1574–1578.
- Worsley, K. J., & Friston, K. J. (1995). Analysis of fMRI time-series revisited—again. *Neuroimage*, *2*, 173–181.