

Attention Protects the Fidelity of Visual Memory: Behavioral and Electrophysiological Evidence

Jie Huang and Robert Sekuler

Volen Center for Complex Systems and Department of Psychology, Brandeis University, Waltham, Massachusetts 02454

Recall from visual memory is vulnerable to the influence of task-irrelevant information, including the remembered, prototypical value of stimuli seen previously. Wilken and Ma (2004) proposed that this distortion of recall was actually adaptive, with the task-irrelevant information compensating for imperfections of memory. We tested their proposal by using trial-by-trial oscillations in the electroencephalogram's alpha band (8–14 Hz) collected from human subjects as a marker for the strength of visual attention. Subjects' recall of stimulus spatial frequency showed a systematic error, namely a shift toward the prototypical value of previously seen stimuli. The magnitude of this prototype effect was strongly related to the amplitude of alpha band oscillations recorded at posterior sensor locations during the first 100 ms after stimulus onset. Our results support the hypothesis that the prototype effect is compensatory for imprecision in memory. Moreover, the attentional modulation of alpha activity during the encoding of the target stimulus is consistent with the view that attention sharpens the neural responses that are elicited by the task-relevant stimulus.

Introduction

Visual attention's importance for cognition lies in attention's ability to selectively foreground information that is task relevant while filtering out what is not. However, sometimes this selectivity is imperfect, allowing task-irrelevant information to infiltrate, influence, and distort visual working memory. Among the well documented sources of such task-irrelevant information is the average or prototypical value of stimuli seen in the recent past. Wilken and Ma (2004) proposed that this error on each trial is actually a sign of the prototypical stimulus' beneficial adaptive function, namely as a supplement to any imperfection of memory. Recently, Huang and Sekuler (2010) reported behavioral results that were consistent with this proposal. In their experiments, the recall of the spatial frequency seen on a trial was systematically shifted in the direction of the average spatial frequency that had been seen on previous trials. Moreover, subjects' deployment of selective attention reduced the "prototype effect," as the phenomenon has been called.

We put Wilken and Ma's hypothesis to a test, focusing on the possible origin of working memory's vulnerability to the seemingly task-irrelevant information represented by the prototypical stimulus. To support our test, electroencephalographic (EEG) signals were recorded from subjects who viewed pairs of Gabor stimuli and then tried to recall the spatial frequency of one of those stimuli. The amplitude of oscillations in EEG's alpha frequency-band (8–14 Hz) is dramatically reduced during various forms of cognitive activity, including selective attention (Kli-

mesch, 1999). We hypothesized that selective attention would modulate the fidelity with which a stimulus would be encoded, and that natural variation in the strength of selective attention would cause this fidelity of encoding to vary from trial to trial. Then, using alpha amplitude as a marker of selective attention, we examined the magnitude of the prototype effect, comparing that distortion of memory trial by trial against the strength of selective attention during various temporal epochs before, during, and after stimulus presentation. To preview our main finding, the prototype effect varies strongly with the amplitude of alpha oscillations recorded at posterior sensor locations, but only during the very earliest visual processing of the stimulus for storage and subsequent recall.

Although selective attention does exert a powerful selective effect, there is disagreement about the mechanism of its effect. One account emphasizes the facilitatory effect that attention has on task-relevant information. It assumes that, for example, alpha oscillations are a manifestation of some active inhibitory process. The suppression of alpha potentiates cortical responses to an attended stimulus (Jokisch and Jensen, 2007; Dijk et al., 2008). An alternative account emphasizes selective attention's function as a protective barrier against task-irrelevant information that might otherwise interfere with memory (Gazzaley et al., 2008; Berry et al., 2009). According to this account, attention-driven improvement in the quality of memory for task-relevant stimuli arises from the suppression of task-irrelevant information (Jokisch and Jensen, 2007; Tuladhar et al., 2007; Freunberger et al., 2009). Our study was designed in part to select between these competing accounts.

Materials and Methods

Subjects

Thirteen subjects participated in the experiment after giving written informed consent. Of these, two subjects' data were excluded from the analysis because of excessive EEG artifacts (epoch rejection rate >50%);

Received May 19, 2010; revised Aug. 13, 2010; accepted Aug. 15, 2010.

This work was supported by National Institutes of Health Grant MH-068404. We thank Lisa Payne for comments on an earlier version of this manuscript.

Correspondence should be addressed to Jie Huang, Mail Stop 062, 415 South Street, Waltham, MA 02454. E-mail: jiehuang@brandeis.edu.

DOI:10.1523/JNEUROSCI.2560-10.2010

Copyright © 2010 the authors 0270-6474/10/3013461-11\$15.00/0

nine of the 11 remaining subjects were female. Subjects' ages ranged from 19–25 years ($M = 23.2$, $SD = 6.46$). All subjects were right-handed, had normal or corrected-to-normal vision as measured with Snellen targets, and normal contrast sensitivity as measured with Pelli–Robson charts (Pelli et al., 1988). None had history of psychological or neurological disorder. All subjects were naive to the purpose of the experiment, and all were paid for their participation.

Apparatus and stimuli

One-dimensional (1D) Gabor stimuli were generated and displayed using MATLAB 7 (MathWorks), supplemented by extensions from the Psychophysics Toolbox (Brainard, 1997). Each Gabor comprised a vertical sinusoidal luminance grating that was windowed by a circular Gaussian carrier with a space constant of 1.14° visual angle. Each Gabor subtended 4.65° visual angle. The contrast of the stimulus' sinusoidal component was fixed at 0.2, a value well above detection threshold. To undermine the possibility that subjects might base their judgments on some local correspondence(s) between stimuli, the absolute phase of each Gabor's sinusoidal component was randomly perturbed on each trial $[0, \pi/2]$. Stimuli were presented on a 21 inch cathode ray tube monitor with a refresh rate of 99.8 Hz and a screen resolution of 1280×960 pixels. The screen luminance was linearized by means of software adjustment, and the mean luminance of the screen was maintained at 32 cd/m^2 . During testing, a subject viewed the computer display binocularly from a distance of 57 cm.

To take account of individual differences in visual discriminability, each subject's spatial frequency discrimination threshold was measured for the same type of 1D Gabor that would be later used in testing short-term memory. Each individual's discrimination threshold was then used to scale the spatial frequencies of stimuli used in the memory test. The discrimination thresholds were used also to normalize the recall errors that individual subjects made. Spatial frequency discrimination thresholds were estimated by two-alternative forced choice trials controlled by QUEST, an adaptive tracking algorithm (Watson and Pelli, 1983). On each trial, two vertical 1D Gabors were presented simultaneously, one to the left of fixation and one to the right, for 300 ms. These were the temporal and spatial conditions that would be subsequently used in the memory test. The center of each Gabor was 3.27° visual angle from the center of the computer display, to the left and right of the central fixation point. By pressing corresponding keys, a subject identified the Gabor, either the left or right, whose spatial frequency was the higher of the two. The higher spatial frequency Gabor was equally likely to occupy the left or right positions. A computer-generated tone provided feedback about response correctness.

The lower spatial frequency of each trial's pair was chosen randomly from a uniform distribution that spanned 0.5–5 cycles/degree. This covered the range of spatial frequencies that would be subsequently used in testing memory. The QUEST algorithm controlled the difference in spatial frequency between the Gabors in a pair. As implemented here, QUEST estimated the difference in spatial frequency that would produce correct judgments 80% of the time. Each subject's discrimination threshold was estimated in three separate, successive runs of QUEST. Of the three resulting threshold estimates, the lowest was used to represent the subject's discrimination threshold. Across subjects the resulting discrimination thresholds ranged from 11 to 22% ($M = 13.94$, $SD = 3.37$).

In the test of short-term memory, a pair of Gabors was presented simultaneously, one to left of fixation and one to right, as had been the case for threshold determinations. Then, after a short retention interval, the subject used a matching procedure to reproduce one Gabor's spatial frequency from memory. The spatial frequencies of the two Gabors dif-

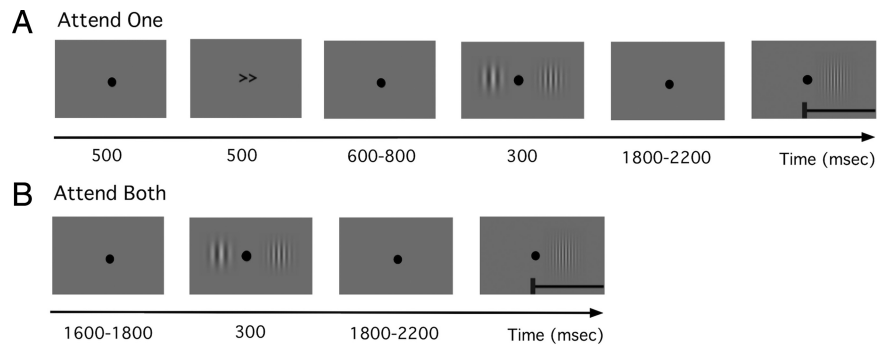


Figure 1. *A*, The sequence of events on a trial in the Att₁ condition. A subject first saw an arrow cue (\ll or \gg) at the center of the display. This was followed 600–800 ms later by the presentation of two horizontally aligned Gabors, which were equally distant from fixation. After the Gabors had been displayed for 300 ms, a random interval of 1800–2200 ms elapsed before the adjustable comparison Gabor was presented. In the example displayed, the cue was \gg , which signified that the Gabor on the right side was the target to be remembered and reproduced, and the one on the left side was nontarget and could be ignored. Subjects were instructed to direct their attention to the right side of the display without making eye movements. Here, the subject reproduced the target Gabor by adjusting the comparison Gabor's spatial frequency until it matched the remembered frequency of the target. *B*, The sequence of events on a trial in Att₂ condition. No visual cue was provided before the presentation of study stimuli; a subject viewed the pair of bilaterally presented Gabors and had to remember both of them. Then, after a short retention interval, the comparison Gabor appeared at one side of the display, signifying which study Gabor (left or right) was the target to be reproduced.

fered always by 4 just noticeable differences (JNDs). To explain how stimulus frequencies were scaled, letting x be the subject's Weber fraction for spatial frequency and f_L be the lower spatial frequency, then a frequency 4 JNDs higher than f_L is $f_L(1 + x)^4$. In the experiment, the lower of each trial's frequency was determined by summing (1) a random value drawn from a uniform distribution spanning 0.5–1 cycles/degree and (2) a second random value that was drawn from a set of discrete JND values [0 1 2 3 4 5 6]. The Gabor whose spatial frequency was the lower of the two study items appeared to the left or right of fixation with equal probability.

Subjects' accuracy of recall was tested under two main conditions, which we call “attend one” (Att₁) and “attend both” (Att₂). These conditions, which are described in detail below, were presented in a block design. The two conditions differed in the number of Gabors, one or two, that subjects had to attend to and remember on each trial.

Att₁. Figure 1*A* shows the sequence of events on a trial in the Att₁ condition. At the start of a trial, an arrow cue is presented at the center of display. This cue notifies the subject which side of the display, left (“ \ll ”) or right (“ \gg ”), should be attended to. Subjects were instructed to maintain fixation at the center of the display throughout a trial and to avoid eye movements. As explained below, compliance was monitored with electrodes that were part of the EEG recording configuration.

These two equally likely cues defined two corresponding types of trials, “attend left” (hereafter, Att_L) and “attend right” (hereafter, Att_R). The two were randomly intermixed within a block of Att₁ trials. The Gabor that was presented later in the attended hemifield was designated the target stimulus. The spatial frequency of this stimulus was to be remembered and then reproduced from memory. The Gabor in the uncued hemifield, which is the nontarget stimulus, was task irrelevant and could be ignored. Subjects were instructed to direct their attention to the target Gabor in the cued hemifield and ignore the nontarget Gabor in the uncued hemifield. Again, while shifting attention, subjects were encouraged to maintain their gaze on the centrally presented fixation point.

As can be seen in Figure 1*A*, a trial began with a central fixation point presented for 500 ms. The fixation point was then replaced by a cue (“ \ll ” or “ \gg ”), which remained visible for 500 ms. Then, 600–800 ms ($M = 700$) later, the two study Gabors were displayed for 300 ms. Finally, after a retention interval of 1800–2200 ms ($M = 2000$), a comparison Gabor appeared in the previously cued hemifield. This comparison stimulus, which occupied the position, left or right, at which the target stimulus had been displayed, was accompanied by a horizontal adjustment bar. Subjects reproduced the target stimulus' spatial frequency from memory by adjusting the comparison Gabor's frequency to match the

remembered frequency of the target. By moving and clicking the computer's cursor on the adjustment bar, the subject produced a real-time change in the comparison Gabor's spatial frequency. A change of one pixel along the horizontal adjustment bar changed spatial frequency by $2.28 \times e^{-5}$ cycles/degree, which was a small subthreshold change. A few iterations of this adjustment process were usually enough for subjects to be satisfied that the comparison stimulus' spatial frequency matched the remembered stimulus. Once satisfied, the subject pressed a key to have the match recorded, and the next trial began.

Att₂. Figure 1*B* illustrates the sequence of events comprising the other kind of trial, *Att₂*. The absence of an attention-directing visual cue forced subjects to attend to and remember the Gabors on both sides of fixation. Then, some 1800–2200 ms ($M = 2000$) later, a comparison Gabor appeared and its location (right or left of fixation) signaled the subject which of the two study Gabors (right or left) was the one that had to be reproduced from memory. In the example shown in Figure 1*B*, the subject would reproduce from memory the study Gabor that appeared to the right of fixation. Over trials, study Gabors to the right or left of fixation were equally likely to be the target, with left and right target conditions randomly intermixed within a block of *Att₂* trials. Note that until the comparison Gabor appeared, a subject had no information about whether the left or right study item was the one that would have to be reproduced. As a result, this condition provided a split-attention control, against which results from *Att₁* could be compared.

Procedure

Each subject served in two sessions. Session one comprised four blocks and session two comprised one single block. Session one's first three blocks were devoted to QUEST measurements of the subject's spatial frequency discrimination threshold. Each of these measurements took 75 trials. Session one's fourth block and all of session two were devoted to testing short-term memory. The two conditions, *Att₁* and *Att₂*, were presented in a block design. Five of the 11 subjects were tested first in *Att₁* and then in *Att₂*; the rest were tested in the reverse order. Each of the two short-term memory blocks comprised 336 trials, with the first 16 trials being practice trials that were excluded from analysis. This resulted in 320 trials for each of *Att₁* and *Att₂* conditions.

Electroencephalographic recording

Scalp EEG signals were recorded while subjects were performing the short-term memory tasks described above. Recordings were made with a high-density, 129 channel geodesic electrode net and high impedance NetAmps amplifiers (Electrical Geodesics). EEG signals were sampled at 250 Hz and were bandpass filtered [0.1–200 Hz]. All channels were adjusted for scalp impedance <30 k Ω . During data acquisition, EEG signals were online referenced to the vertex channel, Cz. Data were cleansed of bad electrode channels (identified by visual inspection) and were offline rereferenced to the grand average of the remaining electrodes. Finally, recordings were transformed into a montage of 19 electrodes defined by the International 10–20 system. These 19 electrode locations were the ones that corresponded most closely to the ones in the 10–20 system. These 19 electrodes were supplemented by bipolar eye channels above and below each eye and a pair of bipolar eye channels near the outer canthi.

We adjusted the intervals during which EEG was recorded on *Att₁* and *Att₂* trials to reflect differences between the trial structures for the two conditions. On *Att₁* trials, EEG signals were recorded from 1700 ms before the onset of the study stimuli until 1800 ms after the stimuli had disappeared from view. On *Att₂* trials, EEG signals were recorded from 600 ms before the presentation of the study stimuli until 1800 ms after the stimuli had disappeared. Then, offline, EEG data were passed through two filters, a 0.1–30 Hz bandpass filter and a 60 Hz notch filter. Blink and eye-movement artifacts were removed by rejecting trials in which the difference between maximum and minimum voltage at any of the three electrooculographic channel pairs exceeded 70 μ V. EEG data that survived this preprocessing were then visually inspected, and trials containing residual artifacts were excluded. These operations produced an average trial rejection rate of 14.04% ($SD = 7.53\%$) for *Att₁* condition, and 10.52% ($SD = 8.76\%$) for *Att₂* condition, which did not differ sig-

Table 1. Electrodes comprising ROIs for each hemisphere

ROI	Left hemisphere	Right hemisphere
Posterior	P3, P7, O1	P4, P8, O2
Central	C3, T7	C4, C8
Anterior	PF1, F3, F7	PF2, F4, F8

nificantly. Moreover, within *Att₁* trials, rejection rates did not differ significantly between the two subconditions, *Att_L* ($M = 13.80\%$, $SD = 7.56\%$) and *Att_R* ($M = 14.27\%$, $SD = 7.86\%$).

As each trial's study Gabors were presented to different hemifields, one to the left and one to the right, our analysis of EEG signals focused on the 16 electrodes that were lateralized and ignored the three midline electrodes. The 16 lateralized electrodes were grouped into three regions of interest (ROIs) for each hemisphere. We will refer to these ROIs by their approximate locations on the scalp: posterior, central, and anterior. These ROIs, together with the electrodes included in each, are shown in Table 1.

The time–frequency decomposition of EEG signals over a trial was performed with complex Morlet wavelets as implemented in EEGLAB (Delorme and Makeig, 2004). The wavelets' lengths increased linearly from 3 cycles at 4 Hz to 11.25 cycles at 30 Hz, in 27 steps. The wavelet transformation was meant to optimize the trade-off between good temporal resolution at lower frequencies and stability at higher frequencies. Alpha band amplitude estimate was calculated as the average oscillatory activity in the band 8–14 Hz.

Results

Behavioral results

The raw error made on each trial is defined by the difference between (1) the spatial frequency of that trial's target Gabor and (2) the spatial frequency that the subject reproduced. Following the method of Huang and Sekuler (2010), we normalized the raw errors made by each subject relative to that subject's Weber fraction. Assume that x is the subject's Weber fraction, let f_T be the spatial frequency of target Gabor on some trial, and let f_R be the comparison Gabor's final adjusted frequency. The raw error in reproduction ($f_R - f_T$), is normalized by the subject's Weber fraction, yielding what we call the “normalized reproduction error” (nRE). Note that this variable was given a slightly different name in Huang and Sekuler (2010), where it was called “transformed reproduction error.” The revised designation is meant to produce greater terminological clarity:

$$nRE = \log_{(1+x)}\left(\frac{f_R}{f_T}\right) = \frac{\ln f_R - \ln f_T}{\ln(1+x)}. \quad (1)$$

Figure 2, *A* and *B*, present the distributions of nREs for the *Att₂* and *Att₁* conditions, respectively. Although the distribution of errors produced in *Att₁* is a bit narrower than the distribution from *Att₂*, the spreads of the distributions (as indexed by MAD, the median absolute deviation from the median) failed to reach statistical significance, $t_{10} = 1.983$, $p = 0.075$.

To assess further the overall accuracy with which a target Gabor was reproduced in each experimental condition, we took the medians of the absolute nRE values in the *Att₂* condition and in the *Att₁* condition. These values are shown by the leftmost pair of bars in Figure 2*C*. Subjects demonstrated consistency in the median absolute nRE made in two test conditions: a subject's median absolute value of nRE in *Att₂* was strongly correlated with that in *Att₁* condition, $r_9 = 0.80$, $p < 0.003$. The median absolute values of nRE made in reproducing the target in *Att₂* condition ($M = 1.322$, $SD = 0.094$) was significantly greater than that in *Att₁* condition ($M = 1.147$, $SD = 0.107$), $t_{10} = 2.741$, $p < 0.021$, paired t test. Given that the absolute nRE represents the size of

error made in recall regardless of its direction (higher or lower), this difference confirms that the differential deployment of selective attention in the two conditions does have a modest but statistically reliable effect on the accuracy of recall.

Huang and Sekuler (2010) showed that each nRE can be thought of as a sum of multiple influences, including two that arise from stimuli other than the target stimulus. To start the process of disentangling those influences, the sign of each nRE was transformed to reflect the error's direction relative to the spatial frequency of the nontarget stimulus. As can be seen in Equation 1, the calculation of nRE does not include a term for the nontarget. Because our experimental design made the nontarget spatial frequency equally likely to be four JNDs greater than or less than the target frequency, the sign of an nRE carries no information whatsoever about the relationship of nRE to the nontarget spatial frequency. To make that potentially important relationship manifest, the sign of each trial's nRE was adjusted relative to the trial's nontarget: a plus sign signified an nRE that was displaced from the target toward the nontarget, and a minus sign signified an nRE that was displaced from the target away from the nontarget. From these sign-adjusted values of nRE, we used an algebraic manipulation to extract (1) one component that reflects the influence of the spatial frequency of a trial's task-irrelevant nontarget Gabor, and (2) a second component that reflects the influence of the prototypical spatial frequency seen on the preceding trials.

The prototypical spatial frequency is defined by the mean frequency of all stimuli that a subject had seen (including both target and nontarget) on the preceding trials within the same block. In an earlier study, Bussemeyer and Myung (1988) used a matching task with random dot patterns and reported a clear recency weighting of the contribution that successive stimuli made to the prototypical stimulus. Additionally, Huang and Sekuler (2010) demonstrated that the prototype effect probably reflects a recency-weighted influence of preceding trials, with the influence of some single prior trial decreasing as the prior trial goes further back. As the stimuli in the present experiment varied randomly from trial to trial, we used the mean frequency of all previously seen stimuli in the block as a surrogate for the prototypical stimulus. Drawing upon the average influence of all preceding stimuli produced a more stable estimator of the effect in which we were interested.

Because of trial-to-trial variation in the nontarget stimulus' spatial frequency, there would be some trials on which the two putative influences would operate synergistically, both pulling the reproduction in a common direction, and other trials on which the two influences would work in opposition to one another. Mindful of this potential difference between types of trials, we divided the sign-adjusted nRE values from the trials within a block into two sets according to whether the two influences should operate in a common direction or in opposite directions. Specifically, for each subject we separated trials (1) on which the two influences would work in the same direction from trials (2)

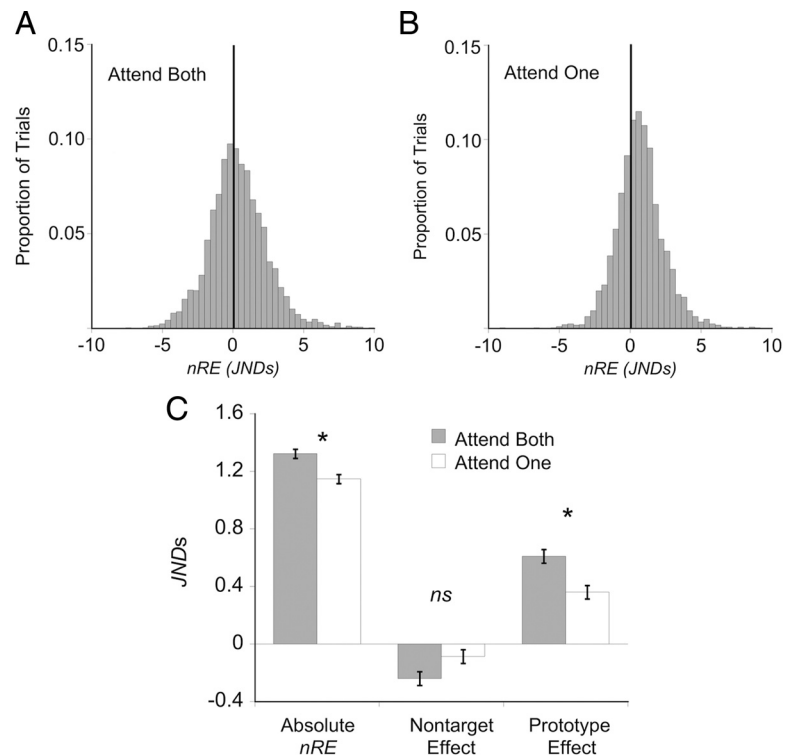


Figure 2. *A, B*, Frequency distributions of the normalized reproduction errors made in Att₂ (*A*) and Att₁ (*B*) conditions. Data are aggregated over all 11 subjects and, for Att₁, over both of its subconditions, Att_{1L} and Att_{1R}. *C*, Bar charts showing median absolute values of nRE (leftmost pair of bars), the effect of the nontarget stimulus (middle pair of bars), and the effect of the prototypical stimulus (rightmost pair of bars). In each pair of bars, results for Att₂ and Att₁ conditions are shown separately. Data are means over subjects. Error bars are ± 1 within-subject standard errors of the mean. A nonsignificant difference between two test conditions is indicated by ns; a difference significant at $p < 0.05$ is indicated by the asterisk (*).

on which the two influences would oppose one another. In the first set of trials, each trial's nontarget frequency and the prototypical frequency seen on the preceding trials bore the same ordinal relationship to the target stimulus' spatial frequency on that trial (e.g., both values were higher than target's frequency, or both were lower); in the other set of trials, the nontarget frequency and the mean, prototypical frequency had opposite ordinal relationships to the target stimulus' spatial frequency (e.g., one value was higher, and the other was lower than target's frequency).

Equations 2 and 3 present more formally the rationale behind our division of trials into two sets. In Equations 2 and 3 Error_{Sdir} and Error_{Odir} is the sign-adjusted nRE value from trials of each type; Bias_{Nontarget} is the magnitude of reproduction bias toward the nontarget, and Bias_{Proto} is the magnitude of the prototype effect, that is, the bias toward the mean spatial frequency seen on preceding trials. When the two component influences worked together, the observed reproduction error would be some sum of the two; when the two components influences operated in opposition to one another, the resulting reproduction error would be a difference between the two. As an nRE value with a positive sign signified a bias toward the nontarget stimulus, the difference between the two components' influences in Equation 3 can be found by subtracting Bias_{Proto} from Bias_{Nontarget}:

$$\text{Error}_{Sdir} = \text{Bias}_{Nontarget} + \text{Bias}_{Proto} \quad (2)$$

$$\text{Error}_{Odir} = \text{Bias}_{Nontarget} - \text{Bias}_{Proto} \quad (3)$$

Next, we used Equations 4 and 5 to separate the components and quantify each of them. To that end, the terms in Equations 2 and

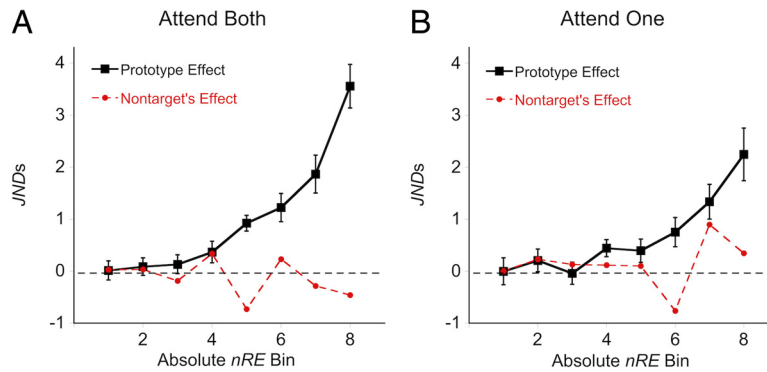


Figure 3. The prototype effect (black square) and the effect of the nontarget stimulus (red circle) are plotted against the absolute nRE. Absolute nRE is expressed as the ordinal value (1–8) of the bin into which each subject's absolute values of nRE were sorted (see Results for details). **A**, Results from Att₂ condition. **B**, Results from Att₁ condition. Data are averaged across the 11 subjects. For each subject, the trials from one condition were grouped into eight bins according to the absolute nRE made on that trial. Then, the prototype effect and the effect of the nontarget stimulus were computed for each bin. In both test conditions, only the prototype effect demonstrated a significant linear trend of increase from bin 1 to bin 8. Error bars are ± 1 within-subject standard errors of the average value.

3 were rearranged to produce Equations 4 and 5. Equation 4 isolates the contribution of the nontarget, nulling the contribution from the prototypical stimulus, whereas Equation 5 does the opposite, isolating the effect of the prototypical stimulus while nulling the contribution of nontarget stimulus. Additional details of this decomposition are given in Huang and Sekuler (2010).

$$\text{Bias}_{\text{Nontarget}} = 0.5(\text{Error}_{\text{Sdir}} + \text{Error}_{\text{Odir}}), \quad (4)$$

$$\text{Bias}_{\text{Proto}} = 0.5(\text{Error}_{\text{Sdir}} - \text{Error}_{\text{Odir}}). \quad (5)$$

Each of the two biases, one from the trial's nontarget stimulus and the other from the prototypical stimulus, was independently compared with zero, using a one-sample *t* test. Then, each one's influence in Att₁ and Att₂ conditions was evaluated by means of paired-sample *t* test.

The nontarget effect

Figure 2C shows the mean size of the two memory-distorting influences extracted from the sign-adjusted nREs. As the middle pair of bars in that figure shows, the nontarget failed to have a reliable effect on the fidelity with which the target Gabor was reproduced. In fact, the effect of the nontarget was not significantly different from zero on either the Att₁ trials, when attention could be directed toward a single target ($M = -0.087$, $SD = 0.22$, $t_{10} = -1.312$, $p = 0.219$), or the Att₂ trials, when attention was divided between two study items ($M = -0.241$, $SD = 0.376$, $t_{10} = -2.129$, $p = 0.06$). The very small effect of a trial's nontarget stimulus is consistent with what we expected when we chose to present target and nontarget stimuli to different hemifields.

The prototype effect

In Figure 2C the rightmost pair of bars shows the prototype effect that was seen in Att₁ and Att₂ conditions. Unlike what was seen for the effect of the nontarget stimulus, the prototype effect was significantly greater than zero in each condition: $M = 0.361$, $SD = 0.45$, $t_{10} = 2.664$, $p < 0.024$ in Att₁ and $M = 0.609$, $SD = 0.437$, $t_{10} = 4.623$, $p < 0.001$ in Att₂. Individual subjects again demonstrated consistency, with subjects' prototype effect in the two conditions being strongly correlated, $r_9 = 0.715$, $p < 0.013$. Of particular theoretical interest is the fact that the prototypical stimulus' influence was greater on Att₂ trials, when attention was divided between the two study stimuli, than on Att₁ trials, when

the timing of the cue made it possible to direct attention toward one of the study items ($t_{10} = 2.606$, $p < 0.026$).

The prototype effect and the absolute nRE
Huang and Sekuler (2010) hypothesized that when the fidelity of memory was compromised, recall was supplemented by drawing on the prototypical stimulus' spatial frequency. This hypothesis predicts a strong relationship between the absolute error, an index of memory's imprecision, and the magnitude of prototype effect. To test this prediction, we first divided each subject's trials into eight equally populous bins based on the absolute nREs made on trials in a particular condition. This maneuver forces the median absolute nRE of each bin into a monotonically increasing relationship with the ordinal designators of the bins (1–8). Then, from the trials contained in each of the eight bins we computed the

magnitude of the prototype effect and the nontarget effect. These computations were done separately for Att₂ and Att₁ conditions. The black data points in Figure 3, A and B, show the results. As can be readily seen, the prototype effect (■) increases from bin 1 to bin 8. The strong association between prototype effect and bin ordinal number was confirmed by a statistically significant linear increase of prototype effect for Att₂ and Att₁ conditions: $F_{(1,10)} = 171.179$, $p < 0.001$ and $F_{(1,10)} = 9.957$, $p < 0.01$, respectively. Higher order components of the association were nonsignificant. The difference between the sizes of effects for Att₂ and Att₁ is consistent with the result noted earlier: in general, Att₂ produces a larger prototype effect than Att₁ does. To examine further the strength of association between the prototype effect and the absolute error, we looked at the correlation (Pearson's *r*) between the prototype effect in each of the eight bins and the absolute nRE in each bin. For this analysis, the values of the prototype effect and of the absolute nRE were averaged across all subjects within a bin. As expected, the prototype effect and the absolute nRE were significantly correlated with one another in each test condition: $r_6 = 0.975$, $p < 0.001$ in Att₂ and $r_6 = 0.964$, $p < 0.001$ in Att₁. These results confirm what was noted above, namely that the prototypical stimulus exerts a greater influence on memory-based reproductions when the overall error in memory for the target stimulus is large.

We also computed the effect of the nontarget stimulus from the trials comprising each of the eight ordered bins and looked for evidence of an association between the size of that nontarget effect and the magnitude of the absolute error. The red data points in Figure 3, A and B, show the result. For neither Att₂ nor Att₁ were the nontarget effect and the absolute error significantly related: $F_{(1,10)} = 0.171$, $p = 0.688$ and $F_{(1,10)} = 0.107$, $p = 0.75$, respectively.

Huang and Sekuler (2010) showed that the reproduction of a target stimulus tended to be attracted toward the task-irrelevant, nontarget stimulus that was presented on that same trial. We saw no such effect here, regardless of whether the nontarget was attended to (Att₂) or selectively filtered out (Att₁). We should note a difference in test procedures used by Huang and Sekuler and here. On each trial, we presented the two study items simultaneously to two different hemifields, with a separation of 6.54°

visual angle; in contrast, Huang and Sekuler presented each trial's study items sequentially to the same central location in the visual field. The absence of a nontarget effect in the present experiment is consistent with Huang and Sekuler's hypothesis that the influence of the nontarget stimulus requires some spatial overlap between the nontarget and target stimuli.

Overall, then, subjects' reproductions of a target stimulus were more accurate when a precue made it possible to attend just to the target Gabor (Att_L) instead of having to attend to and remember both study items (Att_R). Additionally, the prototype effect was greater in Att_R condition when attentional resource was divided between target and nontarget stimuli. These findings show that timely deployment of selective attention can have a strong beneficial influence on the fidelity of subjects' recall. This attentional effect is distinct from but not incompatible with previously demonstrated attentional effects on early vision (Reynolds and Chelazzi, 2004; Pestilli et al., 2009). Although attention had been shown to boost the detectability of a stimulus at an attended location and enhance the perceived contrast of an attended Gabor stimulus (Carrasco et al., 2004), those effects are distinct from the increased precision that selective attention has shown in our study. For example, our stimuli were well above detection threshold, so any further attention-driven boost of contrast would have been unlikely to show a discernible effect on the extraction of spatial frequency information (Smith and Thomas, 1989).

EEG results

Alpha activity and the prototype effect

The behavioral results just described included two findings of particular theoretical importance. First, the accuracy of recall improved and the prototype effect diminished if a timely cue directed attention to the hemifield in which a target stimulus would appear; second, as the absolute nRE in recall grew, so too did the influence of the prototypical stimulus. Together, these findings suggest a connection between the magnitude of prototype effect and the imprecision of memory. If the prototype effect in our behavioral results were compensatory for imprecision in memory, as has been suggested (Huang and Sekuler, 2010), that imprecision could have arisen during any one of a number of stages of stimulus processing. To select among these multiple stages, we turned to the EEG signals collected in parallel with our behavioral measures. Suspecting that trialwise fluctuations in attention might be at fault, we focused on signals in the alpha band collected while subjects were performing the short-term memory task.

Our primary goal was to identify ROIs and temporal epochs during which alpha oscillations were strongly correlated with the

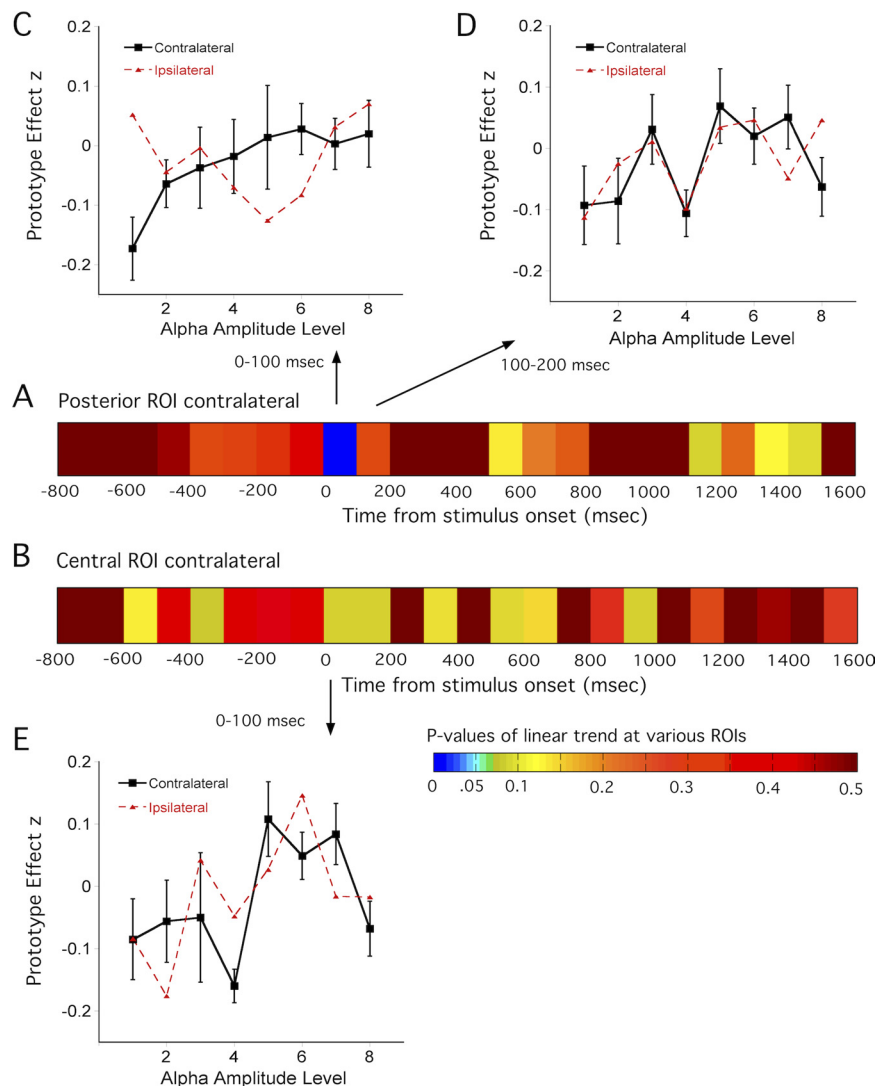


Figure 4. *A, B*, The strength of the linear relationship between alpha amplitude contralateral to the target and the prototype effect at various times before and after stimulus onset. The colors represent p values for the linear trend in the prototype effect across eight alpha-defined bins within successive 100 ms intervals. *A* shows results from posterior ROI and *B* shows results from central ROI. As explained in the text, for each analysis trials were sorted into eight equally populous bins according to the ongoing alpha amplitude within the 100 ms time window; bin 1 comprises a subject's trials with the lowest alpha amplitude, and bin 8 includes trials on which alpha amplitude was highest. A significant linear increase in the prototype effect with alpha amplitude is evident immediately after stimulus onset at posterior ROI contralateral to target. *C–E*, The prototype effect in selected 100 ms time windows plotted against ordinal values of alpha amplitude (based on the eight bins described earlier). In each panel, black squares are for alpha values measured contralateral to the target stimulus, and red triangles are for alpha values ipsilateral to the target. Error bars are ± 1 within-subject standard errors of the average value. Values in *C* and *D*, based on posterior ROI, are for 0–100 and 100–200 ms after stimulus onset, respectively; values in *E*, based on central ROI, are for 0–100 ms after stimulus onset.

magnitude of the prototype effect. We began by examining alpha amplitude at the posterior ROI, whose electrodes sit over the brain regions that are most likely to have participated in the initial stages of visual encoding. We hypothesized that if errors in recall did arise from imprecision in the initial visual encoding of the stimulus that had to be remembered, alpha activity shortly after that stimulus' onset, during the period of its visual encoding, should be predictive of subsequent memory recall.

To test this prediction, all values of nRE and alpha amplitude were separately converted to standardized values (z -scores). These standardizations were done separately for each subject and for the Att_L and Att_R subconditions of Att_L . This transformation muted the impacts of (1) between-subject differences in the overall precision of recall or (2) baseline, hemispheric asymmetries in

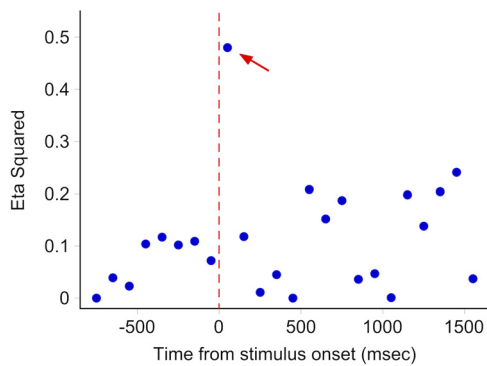


Figure 5. Effect size, defined as η^2 , of the linear relationship between the prototype effect and alpha amplitude. The data points represent η^2 values based on the alpha amplitude at posterior ROI contralateral to the target stimulus within successive 100 ms intervals, both before and after stimulus onset. The x-axis represents time from stimulus onset and maps to the time scale in Figure 4, A and B. Time windows extend from 800 ms before stimulus onset to 1600 ms after, in steps of 100 ms. The red vertical dashed line marks the onset of target. The data point next to the arrow shows the value of η^2 during 0–100 ms after target onset.

alpha amplitude. Then, the z-transformed alpha amplitudes were sorted into eight equally populous bins for each subject and condition. By numbering the bins according to the ongoing alpha amplitude contained in each, bin 1 comprises a subject's trials with the lowest alpha amplitude, and bin 8 includes trials on which alpha amplitude was highest. A set of eight such bins was generated for every combination of ROI location (posterior vs central), time window relative to target onset (800 ms before to 1600 ms after stimulus onset), and hemisphere (contralateral or ipsilateral to the attended target stimulus). After a set of eight alpha-based bins had been generated from alpha amplitudes alone, the prototype effect associated with the trials in each of the bins was calculated. Finally, we evaluated the strength of the relationship over the eight ordered bins between the prototype effect and the alpha amplitude. Note that a number of previous studies used this same kind of binning as a basis for examining any linear trends in their data (Thut et al., 2006; Dijk et al., 2008; Busch et al., 2009). Various panels of Figure 4 illustrate the results.

The colors within the horizontal bars of Figure 4, A and B, colors represent the p values for the linear trend of the prototype effect over the eight ordered bins that were based on the alpha amplitude in the hemisphere contralateral to the target stimulus. Figure 4A shows the results from posterior ROI, and results from central ROI are in Figure 4B. Proceeding left to right across either array, regions represent successive 100 ms time windows, starting at 800 ms before stimulus onset and extending to 1600 ms after stimulus onset. The solitary blue rectangle in Figure 4A indicates that during the first 100 ms after target onset, alpha amplitude at the posterior ROI contralateral to target is a strong predictor of the prototype effect ($F_{(1,10)} = 9.228, p < 0.013$). To highlight the temporal specificity of the effect, we calculated the effect size (η^2) associated with each of 100 ms interval in Figure 4A. The variation in effect size with time relative to stimulus onset is shown in Figure 5 for the posterior ROI contralateral to the target stimulus. In Figure 5, the data point next to the arrow shows that the effect size for this 0–100 ms time window after stimulus onset, $\eta^2 = 0.48$, is more than $2\times$ the effect size in any other time window. This confirms the singular character of that time window.

The black data points (■) in Figure 4C present a more detailed view of the linear trend represented by the blue rectangle in Figure 4A. In Figure 4C, the prototype effect during the interval corresponding to the blue rectangle's time window is plotted

against the ordinal values of the bins containing alpha amplitudes from individual trials. As the black data points show, the prototype effect grows nearly linearly with the amplitude of alpha measured at the posterior ROI contralateral to the target ($p < 0.013$), and the linear trend accounts for 48% of the variance in the prototype effect across the eight alpha-defined bins. As mentioned earlier, the magnitude of alpha necessarily increased monotonically with bin number, but that increase was not precisely linear. For another, more direct perspective on the relationship between prototype effect and alpha amplitude, we examined the correlation (Pearson's r) between alpha amplitude and the magnitude of the prototype effect across the bins. The result was a strong positive correlation between prototype effect and alpha amplitude: $r_6 = 0.855, p < 0.007$. It is noteworthy that an analysis of this very same time window produced no comparable association between the prototype effect and alpha amplitude at the posterior ROI ipsilateral to the target, $F_{(1,10)} = 0.046, p = 0.835$ (Fig. 4C, red triangle data points). In sum, the amplitude of alpha oscillations at the posterior ROI immediately upon stimulus presentation is a strong predictor of the fidelity with which the encoded target will subsequently be recalled several seconds later. Moreover, this relationship between alpha amplitude and the prototype effect is seen only over the hemisphere to which the retinal image of the target projects.

Although we do not claim that neural encoding of spatial frequency has run its entire course within just 100 ms from stimulus onset, our results suggest that a significant portion of that encoding seems to occur within that short period, the equivalent of about one cycle of oscillation with the alpha band. A clue to the importance of this one time period is the considerable temporal specificity shown by the black data points in Figure 4C. In further support of this idea, the black data points (■) in Figure 4D result from an analysis like the one that produced their counterparts in Figure 4C, but now for alpha amplitude recorded at the posterior ROI just 100 ms later, that is, during 100–200 ms after stimulus onset. Note that the relationship between alpha amplitude and the prototype effect during this period is much weaker than it was during the immediately preceding window ($F_{(1,10)} = 1.344, p = 0.272$). Moreover, the weak relationship seen contralateral to target (■) does not differ from the relationship seen in the other, ipsilateral hemisphere (▲): $F_{(1,10)} = 1.472, p = 0.253$.

Figure 4E shows the outcome of the corresponding analysis performed on alpha oscillations recorded at the central ROI. With alpha oscillations taken during the first 100 ms after stimulus onset, the linear trend of the prototype effect over the eight bins approaches statistical significance, but this nearly significant linear trend does not distinguish between alpha amplitude recorded at the central ROI contralateral to the target ($F_{(1,10)} = 3.60, p = 0.087$) and the amplitude at the central ROI ipsilateral to the target ($F_{(1,10)} = 4.35, p = 0.064$). Importantly, as can be appreciated from Figure 4E, the prototype effect's relationship to alpha amplitude at the central ROI is distinctly nonmonotonic. To emphasize the point that alpha amplitude at central ROI does not predict the influence of the prototypical stimulus, a correlation analysis showed no significant relationship between a bin's alpha amplitude and its prototype effect, $r_6 = 0.373, p = 0.363$ for the central ROI contralateral to the target, and $r_6 = 0.476, p = 0.233$ for the central ROI ipsilateral to the target.

Finally, we should comment on the relationship between alpha amplitude and the absolute nRE, which represents the magnitude of overall error in recall and, like the prototype effect, can be an index of imprecision in memory. When absolute nRE is included in an analysis like the ones reported above for the pro-

prototype effect (Fig. 4A,B), the magnitude of absolute nRE is not significantly related to alpha amplitude in any combination of times (100 ms windows, -800 to 1600 ms from stimulus onset) and electrode sites (posterior or central ROI), all $p > 0.10$. So, although absolute nRE and the prototype effect are well correlated (Fig. 3), the insignificant relationship between absolute nRE and alpha amplitude ratifies our decision to focus on the prototype effect rather than absolute nRE for possible correlations between recall and alpha oscillations. It stands to reason that some portion of the absolute nRE reflects influences that are unrelated to the precision with which the stimulus is encoded. Specifically, such influences could include any irreducible imprecision in controlling the computer mouse and any limits on the precision of visual perception that guides the matching process.

Attentional modulation of alpha activity

The preceding section demonstrated that alpha amplitude during a target's visual encoding was a good predictor of the prototype effect. We followed up by investigating the topographical distribution of alpha oscillations during that important initial period of visual processing. We were particularly interested in how the topography of alpha signals might vary with changes in selective attention.

The attention-related comparisons in which we were interested were potentially vulnerable to distortion by baseline differences between alpha oscillations in left and right hemispheres. Therefore, we assessed such hemispheric differences in baseline alpha activity in each of the attentional conditions that we wanted to compare, Att₁ and Att₂. Differences in the temporal structures of trials in the two conditions necessitated that a different interval would be used to define "baseline alpha" in each condition. For Att₁, baseline alpha was defined by the mean alpha amplitude during -200 to 0 ms before cue onset; for Att₂, the absence of cue meant that baseline alpha had to be defined by the mean alpha amplitude during -200 to 0 ms before stimulus onset. As others have reported (Thut et al., 2006; Rihs et al., 2007), for each condition, baseline alpha at the left posterior ROI tended to exceed baseline alpha at the right posterior ROI, $F_{(1,10)} = 4.827$, $p < 0.05$ and $F_{(1,10)} = 6.577$, $p < 0.028$ for Att₁ and Att₂, respectively. To null this hemispheric asymmetry's influence on any analysis of attention-related effects, the alpha amplitude in each condition was baseline corrected by subtracting each electrode's baseline amplitude. We should note that correction for hemispheric differences in baseline alpha was not needed for our analysis of the prototype effect's association with alpha amplitude. Before they were binned, alpha amplitudes recorded from each hemisphere had been separately converted into z-scores. This fact exempted the analysis from likely influence of alpha hemispheric asymmetry and, hence, from the need for baseline correction.

After baseline correction, on Att₁ trials the average baseline-corrected alpha amplitude at posterior ROI during the prestimulus interval was significantly below zero, $t_{10} = -5.276$, $p < 0.001$ for Att_L and $t_{10} = -2.934$, $p < 0.015$ for Att_R. This reduction of alpha activity at posterior ROI following the presentation of the attention-directing cue suggests that the attention-directing cue altered the state of cortical regions that would be used for pro-

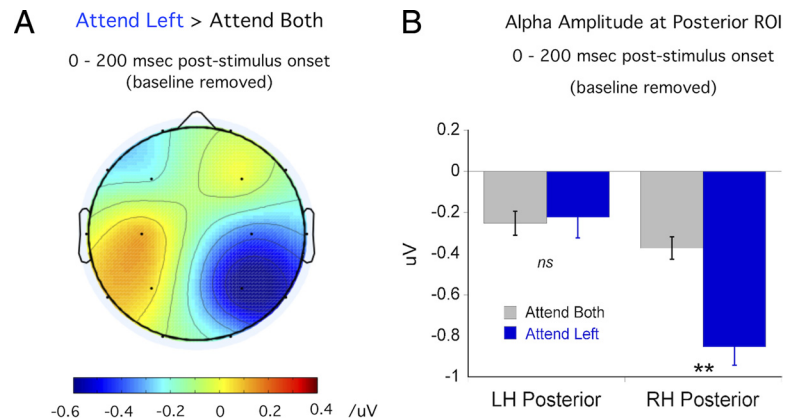


Figure 6. *A*, The difference between Att_L and Att₂ conditions' baseline-corrected alpha amplitude during 0–200 ms after stimulus onset. Results are averaged across the 11 subjects. *B*, Mean baseline-corrected alpha amplitude at posterior ROI during 0–200 ms after stimulus onset in Att_L (blue bars) and Att₂ (gray bars) conditions. The error bars are ± 1 within-subject standard errors of the mean. A nonsignificant difference between the two test conditions is shown as ns; a difference significant at $p < 0.01$ is shown as double asterisks (**). LH, Left hemisphere; RH, right hemisphere.

cessing the upcoming study stimuli. The prestimulus alpha reduction at posterior ROI did not differentiate between the hemisphere contralateral to the upcoming target and the hemisphere ipsilateral to the target.

Given a close link between alpha activity during visual encoding and the prototype effect, we aimed at characterizing the processes of top-down attentional modulation during that early period of visual processing in the following analyses. Specifically, we compared the topographical distribution of alpha amplitude in Att₁ and Att₂ conditions, as well as in Att_L and Att_R subconditions of Att₁. For individual subjects, data were averaged across trials in each condition.

Att₂ versus Att₁

To compare the topography of alpha activity during the early visual processing of the stimulus, we first examined the differences in alpha activity in Att₂ and Att₁ conditions. It is noteworthy that this analysis also serves a theoretically important purpose to evaluate two alternative ways that goal-directed control might operate. In particular, recall that on Att₂ trials attentional resources were shared between target and nontarget; however, on Att₁ trials attention was devoted to the target stimulus alone. Therefore, the comparison between the two conditions could address the question whether facilitatory or suppressive mechanism of attention selectivity was at work in this task.

Figure 6A shows the topographical distribution of this difference during 0–200 ms after study stimulus onset between Att_L and Att₂ conditions. Note that during this 200 ms period, the topography of alpha amplitude difference is dominated by an appreciable amplitude reduction at right posterior ROI. This means that alpha amplitude at posterior ROI contralateral to the target in Att_L was significantly lower than that in Att₂: $F_{(1,10)} = 17.573$, $p < 0.002$. Whereas no comparable difference was observed at left posterior ROI, this means that posterior alpha contralateral to the ignored nontarget stimulus in Att_L did not differ from that in Att₂: $F_{(1,10)} = 0.044$, $p = 0.80$ (Fig. 6B). Further analysis on other time windows during a trial demonstrated that such attentional effects on posterior alpha were restricted only to the particular time window of visual encoding (0–200 ms after stimulus onset). This pattern suggests that visuospatial attention played a role in suppressing alpha oscillations at posterior ROI contralateral to the attended target during early encoding.

The same comparison of alpha amplitude during 0–200 ms after stimulus onset was performed between Att₂ and Att_R conditions.

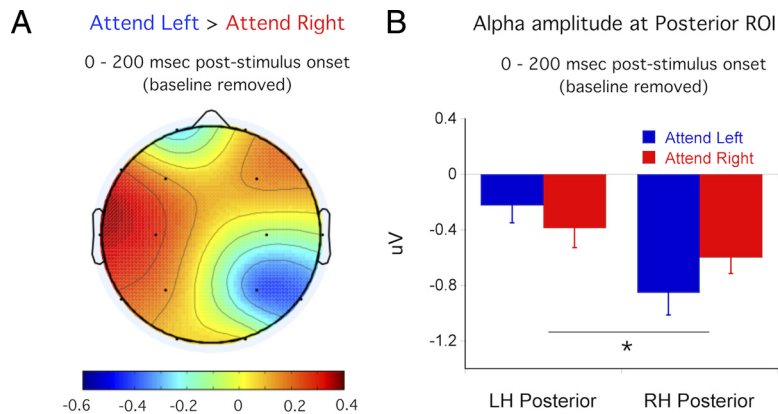


Figure 7. *A*, The difference between baseline-corrected alpha amplitude in Att_L and Att_R conditions during 0–200 ms after stimulus onset. Results are averaged across the 11 subjects. *B*, Mean baseline-corrected alpha amplitude at posterior ROI during 0–200 ms after stimulus onset in Att_L (blue bars) and Att_R (red bars) conditions. A difference significant at $p < 0.05$ is shown as an asterisk (*). LH, Left hemisphere; RH, right hemisphere.

However, alpha amplitude at posterior ROI in Att_R did not exhibit any significant difference from that in Att_L : all $p > 0.45$. Att_R condition's failure to show any difference in alpha amplitude from Att_L suggests an asymmetry in alpha-based expression of visuospatial attention that favors attending to the left hemifield. This asymmetry accords with the report by Thut et al. (2006) of an advantage for left hemifield visual processing with a visual detection task, which required subjects to shift attention to either a left or right peripheral location to detect an upcoming visual target.

Att_L versus Att_R

We next assessed attention's impact on alpha amplitude in the two subconditions of Att_L . For this, we compared alpha topography in Att_L and Att_R . Figure 7*A* depicts the topographical distribution of the difference in baseline-corrected alpha amplitude between Att_L and Att_R conditions during 0–200 ms after stimulus onset. As can be appreciated in the figure, Att_R had reduced alpha amplitude at left posterior ROI, while Att_L had reduced alpha amplitude at right posterior ROI; the two-way interaction of condition \times hemisphere is significant: $F_{(1,10)} = 7.434$, $p < 0.021$. This significant interaction indicates that during 0–200 ms after stimulus onset, alpha amplitude was always reduced to a greater extent at the posterior ROI contralateral to the target. Further analysis on other time windows did not show such an attentional effect at posterior ROI: all $p > 0.10$.

Additionally, Figure 7*A* also suggests some attentional modulation at central ROI, as Att_L and Att_R exhibited opposite patterns of alpha reduction at central ROI in the two hemispheres. To examine such a possibility, we compared baseline-corrected alpha amplitude at central ROI between Att_L and Att_R conditions. We found significant interaction of condition \times hemisphere for central ROI over the entire prestimulus interval (–1200 to 0 ms) and the entire post-stimulus interval (0–1600 ms): $F_{(1,10)} = 6.286$, $p < 0.031$ and $F_{(1,10)} = 4.926$, $p < 0.05$, respectively.

Before concluding that selective attention exerts an effect on alpha amplitude at central ROI, we should consider the possibility that alpha amplitude changes at those sensor locations could be unrelated to attention per se. Recall that the matching task used here required subjects to manipulate the computer mouse, executing a series of hand/finger movements. Unilateral voluntary movement or even motor imagery can induce a suppression of the mu rhythm, whose frequency band (8–13 Hz) overlaps that of the alpha band over central scalp (Pfurtscheller, 1992;

Pfurtscheller and Lopes da Silva, 1999). As a result, it is possible that the alpha amplitude changes seen over central regions could actually have been a suppression of mu oscillations triggered by the subject's anticipation and/or imagery associated with the matching task. However, as subjects always used their right hands to perform the matching task, reductions in mu rhythm induced by motor components should be seen only contralaterally, that is, always over the left central regions in either test condition. As this prediction is disconfirmed by our findings, we believe that changes in alpha amplitude at central ROI reflect some persistent attention-related effect that can be seen throughout most of a trial.

In summary, top-down visuospatial attention modulated alpha amplitude at posterior ROI during the initial period of visual processing. This is the interval during which alpha amplitude at the posterior ROI contralateral to target was strongly associated with the prototypical stimulus' effect on recall. Furthermore, it appears that in this case selective attention operates via some facilitatory mechanism that potentiates visual encoding of the target stimulus. Modulation of attention was observed also at central ROI, although this modulation was more distributed over time, before and after stimulus onset, and was unrelated to errors in subsequent recall.

Discussion

Alpha oscillations and attention

When a trial-specific cue in our experiment directed attention toward one visual hemifield, the cue differentially affected the hemispheres' alpha amplitudes at posterior sites immediately after stimulus onset. This observation is consistent with the use of alpha oscillations as a marker of visuospatial attention (Kastner et al., 1999; Worden et al., 2000; Sauseng et al., 2005; Kelly et al., 2006), and it accords well with previous reports that neural signals of selective attention can be detected almost immediately upon the onset of a visual stimulus presented at a to-be-attended location (Mangun et al., 1997; Hillyard et al., 1998; Hopfinger et al., 2000). Several researchers have proposed that the amplitude of alpha oscillations over posterior regions of the brain reflects the momentary state of excitability of the visual cortex (Pfurtscheller, 2001; Thut et al., 2006; Romei et al., 2008). In that view, when alpha amplitude is low the visual cortex should be more responsive to incoming stimulation (Jokisch and Jensen, 2007; Dijk et al., 2008). Extrapolating that proposal to our own recall task, the amplitude of alpha oscillations around the time of stimulus presentation would be related, although inversely, to the reliability with which a visual stimulus is encoded. Moreover, the strong positive relationship between the prototypic effect and posterior alpha amplitude seen in our study suggests that when the encoding of a target is less robust, the prototypic stimulus can gain in influence.

So far, we have emphasized the functional significance of alpha oscillations immediately after stimulus onset. These oscillations reflect effects of visuospatial attention early on in stimulus encoding, and they are strongly related to subsequent recall. An important question concerns how early sensory-evoked brain potentials [evoked potentials (EPs) or event-related potentials (ERPs)] might be involved in these processes. A number of studies have shown that early sensory components of ERPs, e.g., P1, the earliest positive de-

flexion, covary with alpha oscillations (Makeig et al., 2002; Gruber et al., 2005). P1 component of ERP, which peaks ~ 100 ms after visual stimulus onset, was found to be the earliest manifestation of top-down attentional control (Luck et al., 1990; Hillyard et al., 1998) and is suggested to be influenced by alpha oscillations in two major ways: by the degree of alpha phase locking (Klimesch et al., 2004; Freunberger et al., 2009) and by a change in the amplitude of alpha oscillatory activity (Freunberger et al., 2008). With a spatial cuing task similar to ours, Freunberger et al. (2008) reported an association between higher P1 amplitudes and higher alpha power; in both cases, an increase in amplitude is functionally related to voluntary inhibition. In their study, subjects had to direct attention to some cued location, in either left or right hemifield. Based on those previous findings, it may be that P1 reflects, at least in part, the alpha oscillations seen here during the initial phase of stimulus encoding.

In our study, shifts of attention modulated alpha oscillations not only at posterior recording sites, but at central recording sites as well. In contrast to modulation observed at posterior sites, modulation at central sites was both less phasic and less stimulus bound. In fact, this central modulation persisted throughout most of a trial, all the way from the cue onset to the onset of the comparison Gabor. Source analysis, combined EEG/functional magnetic resonance imaging studies, and recordings of local field potentials have identified generators of alpha oscillations in several nodes of the attentional control network, including regions of frontal, parietal, and inferotemporal cortex (Laufs et al., 2003; Moosmann et al., 2003; De Munck et al., 2007; Dockree et al., 2007; Bollimunta et al., 2008). Using a memory task, Jensen et al. (2002) found load-dependent changes in alpha power at both central and posterior regions.

Their findings suggest that central regions may be involved in producing attention-related and/or memory-dependent alpha activity. We should clarify that, although alpha oscillations at central regions were distinctly modulated by shifts of attention in our study, the amplitude of such oscillations was not related to our main behavioral measure, the magnitude of the prototype effect. Together with the observation of Bollimunta et al. (2008), the lack of relationship between central alpha oscillations and behavior that we saw suggests an important distinction between the functional significance of alpha at central sites and alpha at posterior sites. Alpha oscillations at central regions are likely to reflect general shifts of attention, while stimulus-bound alpha effects at posterior regions are more likely related to the robustness of target encoding and, consequently, are able to predict the fidelity with which a remembered target is recalled.

Origin of the prototype effect

Wilken and Ma (2004) observed that the recalled value of a stimulus tends to be shifted toward the mean value of stimuli seen on some previous trials. They proposed that this prototype effect comprises an adaptive response to trial-to-trial fluctuations in the fidelity of memory, a proposal subsequently supported by results of Huang and Sekuler (2010), with recall of stimulus' spatial frequency. Wilken and Ma's proposal receives additional support here from the strong association between alpha amplitude, our marker of selective attention, and the influence of the prototypical stimulus. In particular, the magnitude of the prototype effect on any trial was related to the alpha amplitude on that trial. Moreover, this strong association was restricted to posterior regions and was seen only during the very earliest phase of visual processing. When a fluctuation in attention causes a stimulus to be encoded with less precision, subsequent recall of that stimulus is more likely to be supplemented by the prototypical stimulus. This interpretation of our results is entirely compatible with the view that the prototypical stimulus operates as a "prior" in

the Bayesian sense (Ma et al., 2006, 2008). We believe that future research could profitably test a Bayesian account of the prototype effect at a more detailed level. Such a test might build on demonstrations that perceptual responses can be strongly influenced by a subject's current prior (Landy and Kojima, 2001; Ernst and Banks, 2002; Körding and Wolpert, 2004; Körding et al., 2004; Stocker and Simoncelli, 2006). Applied to our memory paradigm, such a test might manipulate the physical reliability of the target stimulus by introducing varying amounts of visual noise and measuring the resulting prototype effect. Clearly, though, this approach would require far more data and a broader range of conditions than those in our study.

Mechanisms of attentional selectivity: facilitation versus inhibition

There are two competing views of the way that top-down attentional influences work. One account views attentional control as suppressing or inhibiting neural responses to task-irrelevant information. In fact, in some circumstances spatial attention does appear to operate by excluding task-irrelevant external noise or distracters (Watson and Humphreys, 1997; Cepeda et al., 1998; Worden et al., 2000; Zanto and Gazzaley, 2009). A competing view portrays attentional control as facilitatory, that is, as a process that boosts the gain of neural responses to task-relevant information. To clarify the roles of facilitatory and inhibitory processes in our task, we compared the neural signals produced by task-relevant and task-irrelevant stimuli in Att₁ and Att₂ conditions. We saw a marked reduction in alpha amplitude associated with the encoding of the target stimulus in Att₁ compared to what was seen in Att₂. So, when attention could be directed to the target, the stimulus-dependent neural signal was enhanced, that is, alpha oscillations were reduced. Importantly, no comparable difference was seen for the nontarget, which suggests that an inhibitory mechanism was not called into play. Therefore, the pattern of results hints at the influence of a facilitatory mechanism that enhances neural responses to the task-relevant target stimulus. This result would be consistent with the proposal of Hillyard et al. (1998) that a facilitatory mechanism is preferentially engaged when task-relevant stimuli must be processed in the context of difficult tasks (here, "direct matching" of remembered stimulus details). The intrinsic difficulty of such tasks can be contrasted to simple detection, or even familiarity-based recognition, as expressed in binary "yes-no" responses. Moreover, when, as in our study, a trial's stimulus set does not exceed the capacity of short-term memory, the inhibition of task-irrelevant information might be relegated to a secondary role.

In summary, visuospatial attention in a short-term memory task modulates the amplitude of alpha oscillations over posterior regions of the brain. In turn, right after target onset this modulation of alpha is a strong predictor of the degree to which recall will be distorted by task-irrelevant information, namely, the value of the previously seen prototypical stimulus. Finally, deployment of selective attention toward one of the two study stimuli considerably improved the precision with which that stimulus was subsequently recalled. This effect of selective attention might be added to the short list of phenomena in which attention alters not only the detectability of a visual stimulus, but also its appearance (Carrasco et al., 2004). The origin of this novel sharpening effect of selective attention is uncertain, but it might entail a frequency-selective counterpart to the attention-driven increase in directional selectivity seen in neurons in the middle temporal area of the primate brain (Martinez-Trujillo and Treue, 2004).

References

- Berry AS, Zanto TP, Rutman AM, Clapp WC, Gazzaley A (2009) Practice-related improvement in working memory is modulated by changes in processing external interference. *J Neurophysiol* 102:1779–1789.

- Bollimunta A, Chen Y, Schroeder CE, Ding M (2008) Neuronal mechanisms of cortical alpha oscillations in awake-behaving macaques. *J Neurosci* 28:9976–9988.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Busch NA, Dubois J, VanRullen R (2009) The phase of ongoing EEG oscillations predicts visual perception. *J Neurosci* 29:7869–7876.
- Busemeyer JR, Myung IJ (1988) A new method for investigating prototype learning. *J Exp Psychol Learn Mem Cogn* 14:3–11.
- Carrasco M, Ling S, Read S (2004) Attention alters appearance. *Nat Neurosci* 7:308–313.
- Cepeda NJ, Cave KR, Bichot NP, Kim MS (1998) Spatial selection via feature-driven inhibition of distractor locations. *Percept Psychophys* 60:727–746.
- Delorme A, Makeig S (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods* 134:9–21.
- de Munck JC, Gonçalves SI, Huijboom L, Kuijter JP, Pouwels PJ, Heethaar RM, Lopes da Silva FH (2007) The hemodynamic response of the alpha rhythm: an EEG/fMRI study. *Neuroimage* 35:1142–1151.
- van Dijk H, Schoffelen JM, Oostenveld R, Jensen O (2008) Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. *J Neurosci* 28:1816–1823.
- Dockree PM, Kelly SP, Foxe JJ, Reilly RB, Robertson IH (2007) Optimal sustained attention is linked to the spectral content of background EEG activity: greater ongoing tonic alpha (approximately 10 Hz) power supports successful phasic goal activation. *Eur J Neurosci* 25:900–907.
- Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415:429–433.
- Freunberger R, Höller Y, Griesmayr B, Gruber W, Sauseng P, Klimesch W (2008) Functional similarities between the P1 component and alpha oscillations. *Eur J Neurosci* 27:2330–2340.
- Freunberger R, Fellinger R, Sauseng P, Gruber W, Klimesch W (2009) Dissociation between phase-locked and nonphase-locked alpha oscillations in a working memory task. *Hum Brain Mapp* 30:3417–3425.
- Gazzaley A, Clapp W, Kelley J, McEvoy K, Knight RT, D'Esposito M (2008) Age-related top-down suppression deficit in the early stages of cortical visual memory processing. *Proc Natl Acad Sci U S A* 105:13122–13126.
- Gruber WR, Klimesch W, Sauseng P, Doppelmayr M (2005) Alpha phase synchronization predicts P1 and N1 latency and amplitude size. *Cereb Cortex* 15:371–377.
- Hillyard SA, Vogel EK, Luck SJ (1998) Sensory gain control (amplification) as a mechanism of selective attention: electrophysiological and neuroimaging evidence. *Philos Trans R Soc Lond B Biol Sci* 353:1257–1270.
- Hopfinger JB, Buonocore MH, Mangun GR (2000) The neural mechanism of top-down attentional control. *Nat Neurosci* 3:284–291.
- Huang J, Sekuler R (2010) Distortions in recall from visual memory: two classes of attractors at work. *J Vis* 10:24.1–27.
- Jensen O, Gelfand J, Kounios J, Lisman JE (2002) Oscillations in the alpha band (9–12 Hz) increase with memory load during retention in a short-term memory task. *Cereb Cortex* 12:877–882.
- Jokisch D, Jensen O (2007) Modulation of gamma and alpha activity during a working memory task engaging the dorsal or ventral stream. *J Neurosci* 27:3244–3251.
- Kastner S, Pinsk MA, De Weerd P, Desimone R, Ungerleider LG (1999) Increased activity in human visual cortex during directed attention in absence of visual stimulation. *Neuron* 22:751–761.
- Kelly SP, Lalor EC, Reilly RB, Foxe JJ (2006) Increases in alpha oscillatory power reflect an active retinotopic mechanism for distracter suppression during sustained visuospatial attention. *J Neurophysiol* 95:3844–3851.
- Klimesch W (1999) EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Res Rev* 29:169–195.
- Klimesch W, Schack B, Schabus M, Doppelmayr M, Gruber W, Sauseng P (2004) Phase-locked alpha and theta oscillations generate P1–N1 complex and are related to memory performance. *Brain Res Cogn Brain Res* 19:302–316.
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247.
- Körding KP, Ku SP, Wolpert DM (2004) Bayesian integration in force estimation. *J Neurophysiol* 92:3161–3165.
- Landy MS, Kojima H (2001) Ideal cue combination for localizing texture-defined edges. *J Opt Soc Am A Opt Image Sci Vis* 18:2307–2320.
- Laufs H, Kleinschmidt A, Beyerle A, Eger E, Salek-Haddadi A, Preibisch C, Krakow K (2003) EEG-correlated fMRI of human alpha activity. *Neuroimage* 19:1463–1476.
- Luck SJ, Heinze HJ, Mangun GR, Hillyard SA (1990) Visual event-related potentials index focused attention within bilateral stimulus arrays. II. Functional dissociation of P1 and N1 components. *Electroencephalogr Clin Neurophysiol* 75:528–542.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Ma WJ, Beck JM, Pouget A (2008) Spiking networks for Bayesian inference and choice. *Curr Opin Neurobiol* 18:217–222.
- Makeig S, Westerfield M, Jung TP, Enghoff S, Townsend J, Courchesne E, Sejnowski TJ (2002) Dynamic brain sources of visual evoked responses. *Science* 295:690–694.
- Mangun GR, Hopfinger JB, Kussmaul CL, Fletcher EM, Heinze HJ (1997) Covariations in ERP and PET measures of spatial selective attention in human extrastriate visual cortex. *Hum Brain Mapp* 5:273–279.
- Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14:744–751.
- Moosmann M, Ritter P, Krastel I, Brink A, Thees S, Blankenburg F, Taskin B, Obrig H, Villringer A (2003) Correlates of alpha rhythm in functional magnetic resonance imaging and near infrared spectroscopy. *Neuroimage* 20:145–158.
- Pelli DG, Robson JG, Wilkins AJ (1988) The design of a new letter chart for measuring contrast sensitivity. *Clin Vis Sci* 2:187–199.
- Pestilli F, Ling S, Carrasco M (2009) A population-coding model of attention's influence on contrast response: estimating neural effects from psychophysical data. *Vis Res* 49:1144–1153.
- Pfurtscheller G (1992) Event-related synchronization (ERS): an electrophysiological correlate of cortical areas at rest. *Electroencephalogr Clin Neurophysiol* 83:62–69.
- Pfurtscheller G (2001) Functional brain imaging based on ERD/ERS. *Vis Res* 41:1257–1260.
- Pfurtscheller G, Lopes da Silva FH (1999) Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clin Neurophysiol* 110:1842–1857.
- Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. *Annu Rev Neurosci* 27:611–647.
- Rihs TA, Michel CM, Thut G (2007) Mechanisms of selective inhibition in visual spatial attention are indexed by alpha-band EEG synchronization. *Eur J Neurosci* 25:603–610.
- Romeo V, Brodbeck V, Michel C, Amedi A, Pascual-Leone A, Thut G (2008) Spontaneous fluctuations in posterior alpha-band EEG activity reflect variability in excitability of human visual areas. *Cereb Cortex* 18:2010–2018.
- Sauseng P, Klimesch W, Doppelmayr M, Pecherstorfer T, Freunberger R, Hanslmayr S (2005) EEG alpha synchronization and functional coupling during top-down processing in a working memory task. *Hum Brain Mapp* 26:148–155.
- Smith BG, Thomas JP (1989) Why are some spatial discriminations independent of contrast? *J Opt Soc Am A* 6:713–724.
- Stocker AA, Simoncelli EP (2006) Noise characteristics and prior expectations in human visual speed perception. *Nat Neurosci* 9:578–585.
- Thut G, Nietzel A, Brandt SA, Pascual-Leone A (2006) Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *J Neurosci* 26:9494–9502.
- Tuladhar AM, ter Huurne N, Schoffelen JM, Maris E, Oostenveld R, Jensen O (2007) Parietal-occipital sources account for the increase in alpha activity with working memory load. *Hum Brain Mapp* 28:785–792.
- Watson AB, Pelli DG (1983) QUEST: a Bayesian adaptive psychometric method. *Percept Psychophys* 33:113–120.
- Watson DG, Humphreys GW (1997) Visual marking: prioritizing selection for new objects by top-down attentional inhibition of old objects. *Psychol Rev* 104:90–122.
- Wilken P, Ma WJ (2004) A detection theory account of change detection. *J Vis* 4:1120–1135.
- Worden MS, Foxe JJ, Wang N, Simpson GV (2000) Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *J Neurosci* 20:RC63.
- Zanto TP, Gazzaley A (2009) Neural suppression of irrelevant information underlies optimal working memory performance. *J Neurosci* 29:3059–3066.