

Out of Mind, But Not Out of Sight: Intentional Control of Visual Memory

Yuko Yotsumoto and Robert Sekuler
Volen Center for Complex Systems, Brandeis University

Does visual information enjoy automatic, obligate entry into memory or, after such information has been seen, can it still be actively excluded? To characterize the process by which visual information could be excluded from memory, we used Sternberg's recognition paradigm, measuring visual episodic memory for compound grating stimuli. Because recognition declines as additional study items enter memory, episodic recognition performance provides a sensitive index of memory's contents. Three experiments showed that an item occupying a fixed serial position in a series of study items could be intentionally excluded from memory. Additionally, exclusion does not depend upon low-level information such as stimulus spatial location, orientation or spatial frequency, and does not depend upon the precise timing of irrelevant information, which suggests that the exclusion process is triggered by some event during a trial. The results, interpreted within the framework of a summed similarity model for visual recognition, suggest that exclusion operates after considerable visual processing of the to-be-excluded item.

Keywords: Recognition memory, intentional ignoring, vision

The visual world teems with countless objects and events. To restrict incoming information to a manageable flow, the visual system selectively filters that information, prioritizing on the basis of attributes such as spatial location, color, direction of motion, and on the basis of behavioral relevance. The consequences of such filtering have been well documented in psychophysical and in physiological studies (for example, Hopfinger, Buonocore, & Mangun, 2000). An analogue to such visual filtering has been demonstrated in memory research, particularly in studies of intentional forgetting as realized in the directed forgetting paradigm (for a review, see MacLeod, 1998). In a typical directed forgetting experiment, subjects are instructed to forget some subset of the items they have studied. Carrying out this instruction improves recall of the remaining items relative to conditions in which all studied items must be remembered. In addition to this benefit from directed forgetting, some studies have revealed a cost: the presence of items that directed forgetting has rendered task-irrelevant can impair recall of task-relevant items. Although directed forgetting has been well documented for verbal materials, evidence is scant that such intentional filtering can alter the operation of non-verbal memory. However, costs associated with task-irrelevant items have been demonstrated

in auditory masking tasks and in visuospatial memory tasks (Crowder, 1978; Parmentier, Tremblay, & Jones, 2004). As for visual memory, under some circumstances behaviorally-irrelevant visual inputs seem to enjoy automatic, obligatory entry to visual memory, which allows behaviorally-irrelevant visual material to interfere with retrieval for task-relevant information. Such interference has been seen in visual memory for various spatial attributes, such as the distance separating dots (Hole, 1996) and the locations of cells in two-dimensional matrices (Toms, Morris, & Foley, 1994; Washburn & Astur, 1998). However, other studies suggest that interference is not inevitable, but depends upon the stimulus dimension and task details (for example, Lalonde & Chaudhuri, 2002; Ostendorf, Finke, & Ploner, 2004).

For a highly-sensitive assay of voluntary, active control over items' entry into visual memory, we adapted Sternberg's (1966, 1975) recognition memory paradigm, examining episodic recognition for series of compound gratings. Previously, Kahana and Sekuler (2002) showed that such stimuli afford many advantages when used as probes of memory and in computational modeling. Because the set of memoranda varies from one trial to the next, successful recognition performance on any trial requires access to episodic information, that is, information about the items seen on that trial. Systematic variation of perceptual differences among stimuli produces highly predictable changes in recognition performance (Zhou, Kahana, & Sekuler, 2004; Sekuler, Kahana, McLaughlin, Golomb, & Wingfield, 2005; Kahana, Zhou, Geller, & Sekuler, in press). These changes, in turn, support detailed quantitative accounts of recognition memory and intentional ignoring. Also, we propose that

Supported by Air Force Office of Scientific Research grant F49620-03-1-0376 and National Institutes of Health grant MH-55687. We thank Michael J. Kahana and Allison B. Sekuler for many useful suggestions. Correspondence concerning this article may be addressed to Robert Sekuler; e-mail may be sent to vision@brandeis.edu.

compared to stimuli such as nameable word or pictures of common objects, compound gratings presented briefly and in rapid succession resist consistent verbal description and rehearsal. For one thing, as can be seen from examples in Figure 1, compound gratings' perceptual homogeneity and confusability make it difficult to assign consistent verbal labels to such stimuli, particularly when they are presented briefly and in rapid succession. Additionally, when subjects are tested with compound gratings rather than words or picture, their encephalographic records show very much weaker oscillatory correlates of sub-vocal rehearsal (Hwang et al., 2005). Finally, recognition tests with compound gratings show no primacy effect such as might be expected had the memoranda been verbally labeled and rehearsed (Kahana & Sekuler, 2002).

Our experiments exploited the fact that episodic recognition memory for compound grating stimuli declines as additional items enter into memory, and that the decline depends upon the perceptual relationships among the items in memory (Kahana & Sekuler, 2002; Kahana et al., in press). Because of memory's dependence on the number of study items, recognition performance provides a sensitive index of memory's contents. If subjects could voluntarily completely exclude from memory one study item, recognition with $n + 1$ items would be comparable to recognition achieved with n items. If exclusion were only partially successful, recognition would fall somewhere between that expected for n and for $n + 1$ items. Such comparisons provide the foundation of our analysis.

The basic manipulation in one condition of our experiments encouraged subjects to view, but filter out or ignore a stimulus presented in one particular serial position. In designing our experiments, we were aware that instructions like ours can backfire, producing an effect opposite to what was intended. Such ironic effects have been well-studied by Wegner (1994), and include memory-related effects (Wenzlaff & Wegner, 2000; Macrae, Bodenhausen, Milne, & Ford, 1997). If our instructions to ignore generated an ironic effect, our manipulation might generate heightened memory for the to-be-ignored stimulus.

Experiment 1

This experiment determined whether subjects could intentionally exclude from memory one of a series of study items. On each trial, subjects first saw either one or two briefly-presented compound gratings, which comprised the *study series* for that trial. The study series was followed by a probe stimulus, hereafter, **p**, a compound grating that either replicated one of the study stimuli, or, with equal probability, was novel for that trial. The subject made a *yes-no* recognition response, judging whether **p** replicated a study item. We measured recognition under three experimental conditions, including one in which subjects attempted to ignore the study item in the second serial position of each two-item study series. Performance in this condition was compared to performance when either both items or just one had to be remembered.

To maximize the likelihood of demonstrating effective ignoring, different conditions were presented in blocks of trials during which subjects attempted exclusion or not. Also, study stimuli that were to-be-ignored differed in orientation by 45° from other stimuli, and to-be-ignored items always occurred at a fixed time within a trial. These conditions were meant to facilitate intentional ignoring.

Methods

Subjects. The subjects were eight paid volunteers whose ages ranged from 18 to 21 years. They had normal or corrected-to-normal visual acuity as measured with Snellen targets, and normal contrast sensitivity as measured with Pelli-Robson charts (Pelli, Robson, & Wilkins, 1988). Subjects were naive as to the experiment's purpose.

Stimuli. Stimuli were compound gratings, each synthesized by adding two orthogonal sinusoidal components of equal contrast. Stimuli were generated and presented using Matlab and extensions from the Psychophysics and Video Toolboxes (Brainard, 1997; Pelli, 1997). Gratings were presented in either of two orientations, upright and oblique. Upright gratings were synthesized by summing one vertical and one horizontal sinusoid, generating a luminance profile

$$L_{x,y} = L_{avg} \left[1 + 0.2 \left(\frac{\cos 2\pi f(x) + \cos 2\pi g(y)}{2} \right) \right] \quad (1)$$

where L_{avg} is mean luminance; f is the spatial frequency of the vertical component, in cycles per degree; g is the frequency of the horizontal component. Oblique gratings were generated by applying to Equation 1 a rotation matrix

$$L(x_o, y_o) = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} L(x_u, y_u) \quad (2)$$

where the subscripts o and u signify oblique and upright respectively. For our experiments $\theta = 45^\circ$.

Each stimulus comprised two sinusoidal components in orthogonal orientations. One orientation component was made critical to subjects' recognition judgments, while the other was irrelevant. The critical orientation component was the only component whose spatial frequency could vary among the stimuli on any trial. As a result of its variation, correct recognition judgments depended upon memory for this critical orientation component, whose spatial frequency covered the range of 0.5 to 3 cycles/deg visual angle. This variation eliminated any potential benefit from learning to recognize or ignore some particular spatial frequency. The critical orientation was either vertical, in which case we designate the compound grating as upright, or 45° , in which case we designate the compound grating as oblique. Irrelevant components were generated by adding a sinusoid orthogonal in orientation to the critical component: a horizontal sinusoidal component for each upright grating, and a 135° sinusoid for each oblique grating. The spatial frequency of an irrelevant component was chosen by a random selection

from a uniform distribution spanning 0.5 to 3 cycles/deg visual angle. This selection was constrained so that the same frequency was used for all stimuli in any one trial. For terminological convenience, throughout the rest of this paper we refer to the *critical component*'s spatial frequency as the *stimulus*' spatial frequency, without consideration of the irrelevant orientation.

Each compound grating subtended 5 deg visual angle in diameter, and was windowed with a bivariate Gaussian function. Stimulus contrast was 0.2, which was well above the detection threshold, but with a steady veiling luminance on the screen, produced minimal or no visible afterimage. To undermine the usefulness of possible comparisons between corresponding locations on successive gratings, each stimulus' orientation components were both randomly phase shifted by 0 to 0.25 cycle.

To reduce individual differences in recognition performance, we adjusted the stimuli with which each subject would be tested, tailoring differences between stimuli to each subject's own Weber fraction for spatial-frequency (Zhou et al., 2004). We used the method of constant stimuli to measure each subject's Weber fraction, presenting a pair of gratings successively on each trial. From one trial to the next, the first grating's vertical spatial frequency varied from 0.5 to 3.0 cycles/deg, while the vertical spatial frequency of the second stimulus was set to a fractional multiple of the standard. Subjects made *yes-no* judgments of whether the two stimuli did or did not match one another, and received feedback after each response. To mimic the conditions that would be used in our recognition experiments, we added a horizontal sinusoidal component to gratings. The frequency of the horizontal component was drawn from the same range as the vertical component's frequency, but because it was kept constant for both gratings on each trial, the horizontal component was rendered irrelevant to the *yes-no* judgment. With these stimuli and the same timing parameters that would later be used in the equivalent experimental condition, we estimated the frequency difference between standard and comparison gratings that produced a mean proportion correct of 75%. Across subjects, the resulting Weber fractions ranged from 24% to 39%.




Figure 1. Samples of compound grating stimuli used in all three experiments. The vertical components of the stimuli at the extreme left and right differ in spatial frequency from the center stimulus by ± 2 Weber fraction units; for the remaining stimuli, the vertical spatial frequency differs from that of the center stimulus by ± 1 Weber fraction unit. The samples shown here are for upright stimuli; the corresponding oblique stimuli could be produced by rotating these samples about their center by 45° . For ease of visibility, all stimuli are shown here at higher contrast than was used in the experiments. Note that the irrelevant, horizontal spatial frequency is constant within a trial, as is the irrelevant, 135° component for s_2 and p_2 in Rem2.

Apparatus. Stimuli were presented on a 15-inch computer monitor with a refresh rate of 95 Hz, and a resolution of 800 x 600 pixels. Routines from the Video Toolbox calibrated and linearized the display luminance. Mean screen luminance was maintained at 30 cd/m²; because all stimuli had the same mean luminance as this steady background, a stimulus' onset or offset produced no overall luminance transient. Subjects' heads were stabilized with a head and chin rest; viewing was binocular from distance of 114 cm. To promote consistent, central viewing, a small, black fixation cross appeared in the center of the display at the beginning of each trial. Subjects were instructed to fixate this cross, which remained on the screen until 200 msec prior to the onset of the trial's first study stimulus. In all experiments, each compound grating stimulus, study items as well as probes, was presented for 700 msec.

Procedure. Subjects were tested in three conditions. Two of them we designate as Rem2 and Rem1, signifying, respectively, that on each trial two study items or just one were presented and had to be remembered. In the third condition, Ignore, subjects viewed two study items on each trial, but only had to remember the first one; they could ignore the second study item. Figure 2 schematizes the stimulus sequences for each of the three conditions. The following paragraphs describe each condition in more detail.

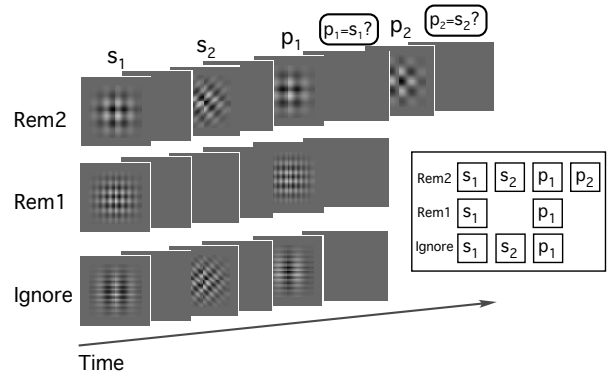


Figure 2. Examples of trial structure and tasks for three conditions. These are examples of stimuli for the upright group: s_1 and p_1 have an upright orientation, but s_2 and p_2 are oriented obliquely. For the oblique group, vertical stimuli and oblique stimuli were reversed: s_1 and p_1 are both oriented obliquely, but s_2 and p_2 have an upright orientation. The labels above the stimuli and questions shown for each condition were not actually present in the stimulus displays; they are given here only as reminders of the judgments that subjects had to make. The diagrams in the insert summarize the serial order of stimuli in each of the three conditions.

Rem2 Condition. On each trial, two study items, s_1 and s_2 , were presented in succession with an interstimulus interval of 1000 msec. After a 1000 msec delay, they were followed by a probe (p_1), which observers judged as matching or not matching s_1 . Then, 1000 ms after the subject's response, a second probe, p_2 , was shown, and subjects judged whether or not p_2 matched s_2 . By requiring separate judgments with respect to both s_1 and s_2 , subjects were encour-

aged to encode and remember both study stimuli. To help subjects keep the two judgments separate, s_1 and p_1 shared the same orientation, 45° from the orientation shared by s_2 and p_2 . For one group of subjects, s_1 and p_1 were upright, while s_2 and p_2 were oblique (the upright group); for the other group s_1 and p_1 were oblique, while s_2 and p_2 were upright (the oblique group).

Rem1 Condition. On each trial, just a single study stimulus (s_1) was presented. This was followed 2700 msec later by a single probe (p_1). Observers judged if p_1 was the same as s_1 , or not. s_1 and p_1 shared the same orientation, which was upright for half the subjects, and oblique for the others. Note that the 2700 msec delay between s_1 and p_1 matched the delay in Rem2 between s_1 and p_1 , which facilitated comparisons between these two conditions. For one group of subjects, s_1 and p_1 were upright; for the other group s_1 and p_1 were oblique.

Ignore Condition. On each trial, s_1 and s_2 , were followed by a single probe, p_1 . A 1000 msec interstimulus interval separated the items, as it did in Rem2. Subjects were instructed to view both study items, but to ignore s_2 , judging only whether p_1 matched the first study stimulus, s_1 , or not. To help subjects perform the task, the orientations of s_1 and s_2 always differed by 45° . For one group of subjects, s_1 and p_1 were upright, but s_2 was oblique (the upright group), and for the other group s_1 and p_1 were oblique, but s_2 was upright (the oblique group).

Previous work showed that recognition performance for stimuli like ours is strongly influenced by the similarity of probe and study items (Kahana & Sekuler, 2002). Therefore, for all conditions p_1 's spatial frequency could differ from that of s_1 by any of several different, fixed amounts: On half the trials in all conditions, p_1 was the same as s_1 , that is p_1 differed from s_1 by zero Weber fraction units; on remaining trials, p_1 was equally likely to differ from s_1 by ± 1 , ± 2 or ± 3 Weber fraction units; the same constraints were applied to p_2 and s_2 in Rem2, the only condition in which there was a second probe. In addition, we forced s_1 and s_2 to differ always by ± 2 Weber fraction units, which minimized confusions between those stimuli (Kahana & Sekuler, 2002).

In a block randomized design, for each condition, every subject was tested in a session with three separate blocks of trials. The first 10 trials of each 110-trial block were eliminated from data analysis. Three sessions were conducted per subject, producing a total of 300 trials per condition for each subject. Distinctive tones immediately after a response provided feedback about response correctness.

Results and Discussion

Figure 3 shows the mean proportion of correct responses for each condition. As differences between vertical and oblique groups were inconsequential, $F(1,6) = 0.182, p > .60$, results from the two groups have been combined in subsequent analyses. In Figure 3, the three gray bars represent mean recognition performance for p_1 in each of the three conditions; the striped bar shows performance for p_2

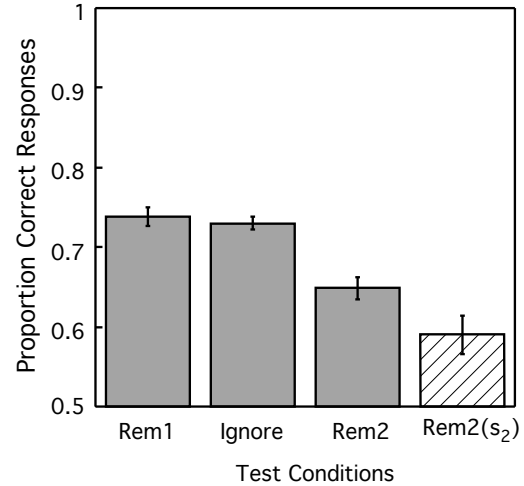


Figure 3. Proportion of correct responses for p_1 in each condition (gray bars), and for p_2 in Rem2 condition (striped bar). Error bars represent standard errors calculated for a within-subject, repeated measures design (Loftus & Masson, 1994).

in Rem2. Recognition performance differed significantly among conditions, $F(3,18) = 27.3, p < .001$. A priori comparisons showed that recognition for Ignore and Rem1 did not differ from one another, $p > 0.40$, whereas recognition for Rem2 was significantly worse than for the other two conditions, $p < 0.01$. Finally, in condition Rem2, p_2 was recognized less well than p_1 ($.01 < p < .02$).

The similarity of mean recognition scores for Rem1 and Ignore suggests that the behaviorally irrelevant stimulus, s_2 , was successfully kept out of memory in Ignore. However, previous work showed that the proportion of correct recognition responses depends strongly on the spatial frequency difference between the probe and study items (Kahana & Sekuler, 2002; Zhou et al., 2004). So, for a more sensitive index of differences in memory among experimental conditions, we broke down Figure 3's results according to the spatial frequency difference between study and probe items. The outcome of this analysis is shown in Figure 4.

When different p_1 - s_1 or p_2 - s_2 similarities were taken into account, the proportion of yes responses varied not only across conditions ($F(3,21) = 38.4, p < 0.001$), but also with the p - s similarities ($F(3,21) = 92.9, p < .001$). In particular, *a priori* comparisons using the Bonferroni adjustment showed significant differences between Rem1 and Ignore ($p < 0.01$), and also between Ignore and Rem2 p_1 , $p < 0.01$. There was also a significant interaction between conditions and the p - s spatial frequency differences, $F(9,63) = 6.5, p < .001$. When the p and s items differed either by three Weber fraction units or by none, the conditions did not differ among themselves after Bonferroni correction. When s and p differed by just one Weber unit, Ignore produced significantly lower proportion yes responses for p_1 than did Rem2 ($p < 0.001$). The corresponding difference between Ignore and Rem1 just failed to reach statistical significance ($.05 < p < .06$). When the spatial frequency difference be-

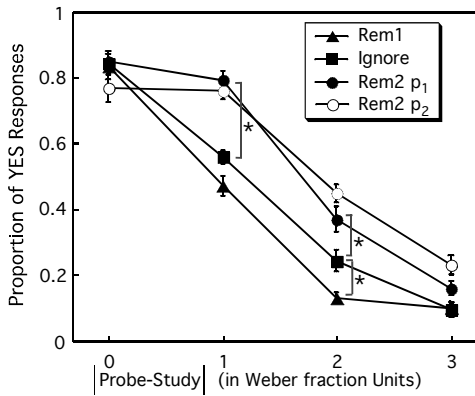


Figure 4. Proportion of *yes* responses as a function of the spatial frequency difference between s_1 and p_1 , or s_2 and p_2 . Values on the x-axis are absolute values of the difference, in Weber fraction units. The figure shows the proportion of *yes* responses in Rem1 condition with filled triangle, Ignore condition with filled square; the proportion correct for p_1 in Rem2 is represented by filled circle; the corresponding values for p_2 in that same condition are represented by open circles. Significant differences ($p < .05$) within each Weber fraction unit are indicated by asterisks. Error bars represent within-subject standard errors.

tween study stimulus and probe was two Weber units, Ignore was significantly different from both Rem1 ($p < .05$) and p_1 for Rem2 ($p < .05$). Significant differences within each p - s Weber unit distance are shown with asterisks in Figure 4.

Our results show that subjects had considerable success in excluding from memory an irrelevant stimulus that had been interposed between a study item and the probe. As Figure 4 demonstrated, this success varied with the spatial frequency difference between p and matching study item. In particular, when p was not identical to, but also not greatly different from the study item, performance was degraded by the interposition of a behaviorally-irrelevant stimulus, s_2 . Note that trialwise variation in the spatial frequency of the memoranda means that any observed ignoring could not have resulted from low-level learning to ignore a specific spatial frequency or narrow range of spatial frequencies.

Finally, note that recognition of p_2 in Rem2 was slightly poorer than the recognition of p_1 in that same condition. Because p_2 's presentation was delayed until the subject made a judgment for p_1 , the difference in performance for the two probes might have reflected a time-dependent degradation in memory for s_2 . On average the interval between p_2 and s_2 was only ≈ 0.75 seconds longer than the interval between p_1 and s_1 . We suspected that this slight increase in delay could not explain the loss in recognition of compound gratings, which is essentially invariant with delays up to ≈ 10 times as long (Magnussen, Greenlee, Asplund, & Dymes, 1991; Magnussen, Greenlee, & Thomas, 1996; Kahana & Sekuler, 2002). However, we decided to assess directly the possible contribution of pre-probe delay. Trials were sorted into five, equal width bins, based on the interval between p_2 and s_2 . A repeated measures ANOVA showed no effect of the inter-

val between s_2 and p_2 on the recognition performance for p_2 ($F(4,28) = 1.36$, $p > .25$). This suggests that the small, but significant difference between recognition p_1 and p_2 was not the consequence of a delay in testing p_2 .

This experiment produced strong evidence that s_2 could be considerably excluded from memory. Before considering the origin of this effect, we had to examine the possibility that subjects had exploited a simple orienting response in order to exclude s_2 . In particular, we wondered if subjects might have closed their eyes or averted their gaze around the time of s_2 's presentation. To test this idea, we made a video of three subjects' eyes during the time they were being tested in the Ignore condition. A light emitting diode (LED) was mounted on the forehead support, just out of subjects field of view. During the presentation of s_2 , the LED was illuminated. A webcam generated a video image in which the LED and the subject's eyes were clearly visible. The video was saved to disk for analysis offline. Later, three naive volunteers viewed the videos in order to judge whether the subjects had been looking directly at the computer display during the periods of LED illumination.

All three judges expressed confidence that subjects had been looking at the display during nearly all of the trials, on average, 97.8%. This result suggests that s_2 's voluntary, substantial exclusion from memory probably did not result from subjects' shifting their gaze to avoid viewing s_2 . We followed up by calibrating the validity of these judgments against a new, staged video in which we controlled whether subjects were or were not looking at the monitor. With this calibration video, judges correctly distinguished the gaze position shown in the video some 95.5% of the time. This supports the idea that during the experiment subjects had not averted their gaze or closed their eyes in order to exclude s_2 .

Experiment 2

Purpose

In Experiment 1, the behaviorally-irrelevant stimulus, s_2 , differed in orientation from s_1 and p_1 . This orientation difference was meant to aid subjects' efforts to exclude the irrelevant stimulus, but is such a difference necessary for memory exclusion? The answer is important theoretically because it could offer insight into the process or processes that control ignoring. If intentional ignoring persisted in the absence of orientation differences, it would suggest that exclusion does not depend on low level spatial differences, of which orientation is one example. Specifically, with an orientation difference between the to-be-ignored item and the other list items, a subject could perhaps filter out all occurrences of the irrelevant item's orientation throughout an entire block of trials; whereas if the orientation were held constant and shared by all list items, subjects would be deprived of that potential cue. We therefore repeated the Ignore condition, with the orientation difference eliminated.

Methods

Stimuli and procedures were the same as those in Experiment 1's Ignore condition, except that s_1 , s_2 , and p_1 all shared the same orientation. The eight subjects from Experiment 1 served in 310 trials, of which the first 10 were treated as practice. Four subjects comprised the oblique group, four the upright group. Subjects in the upright group were shown only upright gratings; subjects in the oblique group were shown only oblique gratings.

Results and Discussion

As oblique and upright groups' results were not noticeably different ($p > .50$), we ignored group membership in all analyses. Figure 5A shows the proportion of *yes* responses as a function of spatial frequency differences between study stimulus, s_1 , and probe, p_1 . For comparison, the range bracketed by the standard errors in Experiment 1's Ignore condition is shown as a gray band; in addition, the proportion of *yes* responses from Experiment 1's Rem2 condition is shown as open circles. An ANOVA confirmed what can be seen in Figure 5A, namely that the proportion of *yes* recognition responses in Experiment 1 and Experiment 2 did not differ ($F(1, 7) = 0.901, p = .37$). This experiment tested only the Ignore condition, which was then compared to results from Experiment 1, with the same subjects. Such comparisons between experiments must be made with caution because with shared testing conditions, the two sets of results could be differentially influenced by the operation of any practice effect. However, a pilot study showed no change in ignoring even with extensive practice, which strengthens our belief that results from the two experiments can in fact be compared.

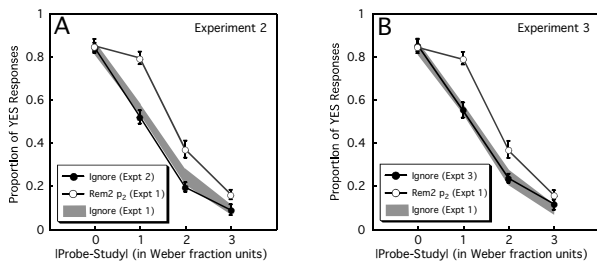


Figure 5. Proportion of *yes* responses as a function of the absolute value of the difference in spatial frequency between study stimulus and probe, in Weber fraction units. Left: Experiment 2, all stimuli shared the same orientation. Within-subject standard errors are shown by the vertical bars around each data point. Right: Experiment 3, the timing of s_2 was unpredictable. To facilitate comparison with results from Experiment 1, in both panels, the range bracketed by standard errors in Experiment 1's Ignore condition is shown as a gray region; additionally, in each panel, the open circles represent the results from Experiment 1's Rem2 condition.

Numerous psychophysical and physiological studies have shown that attention can be allocated to some specific low-level feature, either diminishing or enhancing that feature (for example, Carrasco, Ling, & Read, 2004; Watanabe et al., 1998; Luck, Chelazzi, Hillyard, & Desimone, 1997; Treue,

2001; Somers, Dale, Seiffert, & Tootell, 1999; Slotnick, Schwarzbach, & Yantis, 2003). If the processing and storage of visual information for memory were dependent upon orientation, a low-level stimulus characteristic, Experiments 1 and 2 would have yielded differential success in exclusion from memory. Because the behaviorally-irrelevant items in Experiment 1 differed in orientation from other items on each trial, orientation-dependent processing and storage could have aided exclusion in that experiment. However, we found no difference between the Ignore condition in Experiment 1 and in Experiment 2, which supports the idea that the success of intentional ignoring does not depend crucially upon low-level differences between stimuli.

Experiment 3

Experiment 1 showed that subjects could, at least partially, exclude from memory a behaviorally-irrelevant stimulus. Experiment 2 showed that this effect did not require an orientation difference between relevant and irrelevant stimuli, but left open the question of what actually triggers the process of exclusion. One possibility emphasizes the timing of the irrelevant stimulus. In both experiments, the ISI and stimulus durations were constant from trial to trial, with s_2 always displayed exactly 1700 ms after the trials start. As a result, subjects could have learned to exclude visual input during a particular interval within a trial. This would resemble a recently demonstrated, time-dependent neural inhibition of motion processing (Bisley, Zaksas, Droll, & Pasternak, 2004). Alternatively, in our experiments, ignoring might have been triggered by some event-based process, either by registering the offset of s_1 and/or the onset of s_2 . Such processes would resemble the active gating mechanisms that filter out irrelevant or extraneous events that occur on trials in psychophysical and physiological studies of sensory processing. These extraneous events occur on every trial, before, during or after a trial's critical visual stimulus (see, for example Seidemann, Zohary, & Newsome, 1998). In order to distinguish between time-based and event-based accounts, we varied the ISI and the pre-probe delay, which made it difficult for subjects to initiate exclusion until they actually saw the onset of s_2 .

Methods

Stimuli and procedures were the same as those in Experiment 1's Ignore condition, except that ISI and the pre-probe delay varied randomly. The pair of values for each trial's ISI and pre-probe delay were drawn independently from the same uniform random distribution, which spanned a range of 300 to 1700 msec. Taking into account the duration (700 msec) of s_2 , these random delays caused the onset of p_1 to vary from 1300 to 4100 msec. Subjects were the eight individuals who had participated in the preceding experiments. Again, half the subjects comprised the upright group, the other half the oblique group.

Results and Discussion

As differences between upright and oblique groups were minimal ($p > .40$), we collapsed results from the two groups. Figure 5B shows the proportion of *yes* responses as a function of spatial frequency differences between s_1 and p_1 , measured with unpredictable onset of s_2 . For comparison, the range of standard errors of Experiment 1's Ignore condition is represented by the gray band, and Rem2 condition is represented by open circle in Figure 5B. The proportion of *yes* responses for Experiment 2 was essentially the same as the corresponding values from Experiment 1 ($F(1, 7) = .01, p > 0.50$). So, even when the timing of s_2 's onset was unpredictable, subjects still successfully ignored s_2 , and did so in a way that not different what we saw with fixed, predictable timing of s_2 . This suggests that ignoring was triggered by events such as the onset of s_2 and/or the offset of s_1 , not by an exclusively time-based anticipatory process.

Note that the random variation in each trial's ISI and pre-probe delay introduced a random delay not only into the timing of s_2 , but also into the timing of p_1 . As a result, from trial to trial, subjects' memory of s_1 had to be retained for varying durations. In particular, when added to the other, fixed intervals, random variation in ISI and in pre-probe delay, combined to produce retention intervals ranging from 1300 to 4100 msec. To examine the effect of retention interval length, for each subject the intervals separating s_1 and p_1 were sorted into five bins of equal width. Then the proportion correct was calculated for trials in each bin, and plotted in Figure 6.




Figure 6. Proportion correct recognitions as a function of the interval between s_1 and p_1 , plotted with a linear regression line. Values on x-axis represent the median of each bin. Error bars represent within-subject standard errors.

We then used least squares to find the best-fitting regression line for each subject's results. A t-test showed that the mean slope of these regression lines was not different from zero ($t(7) = .72, p > .50$), which confirms that memory for spatial frequency is robust, at least over several seconds (Kahana & Sekuler, 2002). Moreover, it shows that memory is

robust even when an irrelevant stimulus is inserted at a random time into the retention interval.

General Discussion

The incomplete, but substantial exclusion observed in this study varied with the perceptual difference between study items and the p . Moreover, results showed that intentional ignoring was triggered not by time-based anticipation of stimulus onset, but by some event, either the onset of the to-be-ignored stimulus, or the offset of the item preceding that to-be-ignored stimulus.

Summed similarity framework for ignoring

Our results do not unequivocally identify a particular stage or stages of recognition memory—encoding, retention, retrieval, matching or decision making—at which active control of visual memory operates. For example, the exclusion process might have prevented the stimuli from being encoded, or the behaviorally-irrelevant stimulus might have been fully processed in early vision, but somehow kept from interfering with the memory of the other study item, either during retention or comparison. To narrow down the stages at which intentional ignoring might possibly have operated, we cast some of our results into the framework of a summed similarity visual recognition memory model. Summed similarity models for visual recognition memory (Nosofsky, 1986; Kahana & Sekuler, 2002) base *yes-no* recognition judgments on the summed pairwise similarities between p and each of the study items in turn. When this summed similarity exceeds some criterion value, the model responds *yes*, the p does match one of the study items. The experiments reported here did not yield sufficient data to support detailed modeling like that in Kahana and Sekuler (2002), but it is worthwhile to examine at least some of our results in a summed-similarity framework.

To begin the model-based analysis, we identified trials on which p_1 and s_1 were separated by one Weber fraction unit. (Note that in every condition the study items, s_1 and s_2 , differed by two Weber fraction units in spatial frequency.) Under these circumstances, as Figure 7 illustrates, s_2 could either be similar (Case 1) or dissimilar (Case 2) in spatial frequency to p_1 . Because of the differential similarity of p_1 and s_2 in the two cases, the summed similarity associated with Case 1 would be larger than the summed similarity with Case 2. In the model framework, this difference in summed similarity means that Case 1 would produce more *yes* judgments than Case 2. Of course, if s_2 had no influence whatever on the summed similarity computation, the resulting judgments should be uninfluenced by s_2 's similarity or dissimilarity to p_1 . Because p_1 actually matched neither s_1 nor s_2 , these *yes* judgments would comprise false alarms (erroneous recognitions).

Figure 8 shows the proportion of false alarms (erroneous *yes* responses) produced in various conditions. The vertical gray bars represent false alarms for Case 1; the dotted bars represent false alarms for Case 2. From left to right, successive pairs of vertical bars are for Rem2, and Ignore for

Experiment 1, Experiment 2 and Experiment 3, respectively. Note that within each pair of bars, the false alarm rate was higher with s_2 similar to p_1 (Case 1) than with s_2 dissimilar to p_1 (Case 2), an outcome consistent with the summed-similarity framework. The dark gray horizontal bar shows ± 1 standard error about the mean false alarm rate in Rem1. As s_2 was absent entirely from Rem1, this horizontal bar represents the false alarm rate generated exclusively by the presence of s_1 . Because the similarity of both p_1 and s_1 , as well as s_1 and s_2 were controlled in this analysis, differences in false alarm rates in Figure 8 could have come only from differences in similarity of s_2 and p_1 . Significant deviations upward from the horizontal bar represent significant increases in false alarms, which presumably arose because s_2 made some significant contribution to summed similarity, and hence to erroneous recognitions.




Figure 7. Diagrams illustrating the spatial frequency relationships among s_1 , s_2 , and p_1 in Rem2 and Ignore conditions. In all cases shown, s_1 and p_1 differed by just one Weber unit. The top row illustrates the case in which the spatial frequencies of s_2 and p_1 were similar to one another, differing by just one Weber fraction unit. The bottom row illustrates the case in which the spatial frequencies of s_2 and p_1 were dissimilar, differing by three Weber fraction units. Within each row, at the left, s_1 's frequency is lower than s_2 's, and at the right, s_1 's frequency is higher than s_2 's.

Consider now the results from the Ignore condition in the three experiments. Comparing the gray vertical bars in Figure 8 shows that when subjects attempt to ignore s_2 , false alarms occur at a significantly lower rate than in Rem2, when subjects are attempting to remember both study items, ($p < .001$ for Experiments 1, 2, and 3). Furthermore, comparing the rightmost three gray bars against the control value from Rem1 shows that false alarms occur at a higher rate in Ignore than in Rem1 ($p < .02$ for Experiments 1 and 3, $p = .06$ for Experiment 2).

This pattern of results shows that when s_2 is similar to p_1 , subjects' intention to ignore s_2 successfully reduces its impact, but leaves a small, significant residual effect. In other words, ignoring appears to have been incomplete. Finally, the dotted bars show that when s_2 was dissimilar to p_1 , false alarm rates in Ignore did not differ from the false alarm rate in Rem1 ($p > .45$ for Experiments 1-3). A summed similarity treatment of these conditions suggests that spatial characteristics of s_2 enters partially into memory even when subjects were attempting to ignore it, although this effect is small compared to the impact of that same stimulus when s_2 must be remembered.

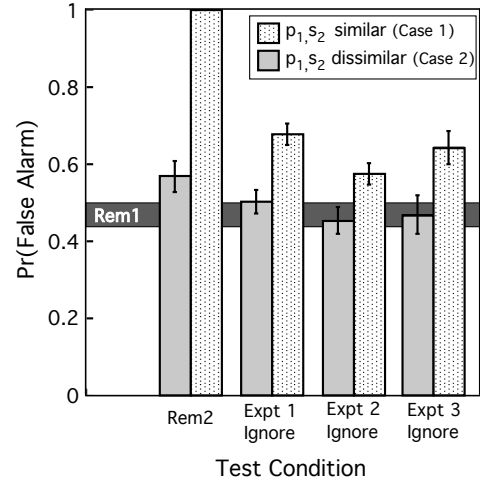


Figure 8. Proportion of false alarms to p_1 in two conditions of Experiment 1 (Ignore and Rem2), and in Experiment 2 and 3's Ignore condition. Data are only for trials on which p_1 differs by one Weber fraction unit from s_1 . Gray bars represent results when s_2 was perceptually dissimilar to p_1 . Dotted bars represent results when s_2 was perceptually similar to p_1 . The horizontal bar running across the figure shows mean false alarms, ± 1 standard error, for p_1 in Rem1. Error bars show within subject standard errors. Note that unlike other figures in the manuscript, poorer performance (more false recognition) is shown as increasing on the vertical axis.

On and Off Gating of Ignoring

Because the display's mean luminance was held constant throughout each experiment, and because the spatial frequency of to-be-ignored items varied substantially trial by trial, any event-based intentional ignoring could not have been triggered by low-level features such as luminance transients or the onset of a particular spatial frequency. And Experiment 2 demonstrated that another low-level feature, orientation was not the trigger either. Additionally, because all stimuli, behaviorally-relevant as well as behaviorally-irrelevant, occupied the same region in the visual field, ignoring could not have arisen from some location-dependent attentional process (for example, Mangun, 1995; Slotnick et al., 2003). Instead, it is likely that higher-order processes are involved.

In Experiment 1, we observed both substantial intentional ignoring (the difference between Ignore and Rem2 p_1 in Figure 4), as well as a small failure of ignoring (the difference between Rem1 and Ignore in Figure 4). We also observed that the characteristics of to-be-ignored stimuli were sufficiently deeply processed as to affect the proportion of false alarms (see Figure 8). In addition, Experiment 3 indicated that ignoring is most likely triggered by some event-based process, not by anticipation of the time at which s_2 would occur. Experiment 3 was not able to distinguish between the possibility that ignoring was triggered by the onset of the to-be-ignored stimulus or by the offset of its preceding stimulus. However, we believe that analysis of false alarms (Figure 8) makes it likely that ignoring is triggered by stimulus onset.

That analysis revealed that a task-irrelevant item was at least partially processed despite the fact that it was supposed to have been ignored. Apparently, before the exclusion process could be fully engaged, the irrelevant stimulus was processed to some degree. In this view, visual registration of the irrelevant stimulus is required for the initiation of its exclusion. As a result, the initial portion of the irrelevant stimulus could enter memory before exclusion took hold fully. This initial, partial leakage of information into memory could explain why exclusion of the stimulus was only partially successful. Therefore, the failure to achieve complete ignoring might reflect the latency of the exclusion process rather than any actual imperfection of that process.

Although intentional ignoring might have been gated on in part by the onset of the behaviorally-irrelevant stimulus, this does not explain how gating might be turned off, as it must be in order for **p** to be processed. We offer two speculative hypotheses about the mechanisms that gate off intentional ignoring. The first hypothesis treats termination of intentional ignoring as event-based, that is, intentional ignoring is turned off by the end of the irrelevant stimulus (e.g. Seidemann et al., 1998), or by the onset of the next event. An alternative hypothesis characterizes the process as time-based. Because s_2 was always presented for 700 msec, it was possible for subjects to learn the timing, and terminate intentional ignoring when that interval had passed. The present experiments cannot distinguish between these alternatives.

Next steps

To maximize subjects' ability to control memory, the experiments reported here used a block design. As a result, within each block of trials subjects could adopt and maintain a consistent strategy. We do not know whether the block design was crucial to success: Would comparable (or any) exclusion be observed if conditions were randomized, with an instructional cue provided just ahead of each trial? And how far in advance would such a cue have to be presented? When trial-specific cues have been used to direct visual attention (e.g. Ball & Sekuler, 1981), the cues require ≈ 400 or more milliseconds lead-time for full effect. If active control of memory operated on a comparably slow time-scale it might be difficult to detect clear signs that memory had been influenced, at least with stimuli presented at or near the rapid timing used here.

Our study revealed substantial active control over people's recognition of what they have just seen, and suggested that this control has its effect after significant visual processing of the to-be-ignored stimulus. Without such control over visual memory for behaviorally-irrelevant stimuli, the burden of processing visual information under ordinary everyday conditions could substantially subvert memory for items and events that were behaviorally relevant. Such subversion could be especially troubling for many elderly people. As Hasher and Zacks (1988) showed, some cognitive deficits exhibited by older people arise from diminished inhibitory control over memory (see also Hasher, Tonev, Lustig, & Zacks, 2001). It would therefore be valuable to determine whether

older subjects are disadvantaged in the exclusion paradigm used here.

Based on recognition performance averaged over trials, the intentional ignoring observed in the present study might be described as "imperfect." As this characterization depends upon comparisons between groups of trials, the imperfection could have arisen in two distinctly different ways. First, subjects might have been entirely successful in ignoring the item, but only on a fraction of all trials; on other trials, ignoring would have been ineffective. Alternatively, ignoring could have been fairly consistent over trials, at about the level represented by our mean data. Unfortunately, the present experiments cannot distinguish between these two alternatives, which necessarily limits our understanding of how the to-be-ignored items are processed. Further research, such as event-related studies of trial-wise brain activation, could shed significant light on this important question.

References

- Ball, K., & Sekuler, R. (1981). Cues reduce direction uncertainty and enhance motion detection. *Perception & Psychophysics*, *30*, 119-28.
- Bisley, J., Zaksas, D., Droll, J., & Pasternak, T. (2004). Activity of neurons in cortical area MT during a memory for motion task. *Journal of Neurophysiology*, *91*, 286-300.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 443-446.
- Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, *7*, 308-313.
- Crowder, R. (1978). Mechanisms of auditory backward masking in the stimulus suffix effect. *Psychological Review*, *85*, 502-524.
- Hasher, L., Tonev, S., Lustig, C., & Zacks, R. (2001). Inhibitory control, environmental support, and self-initiated processing in aging. In M. Naveh-Benjamin, M. Moscovitch, & I. Roediger, R L (Eds.), *Perspectives on human memory and cognitive aging: Essays in Honour of Fergus Craik* (p. 286-297). East Sussex, UK: Psychology Press.
- Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: A review and a new view. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, p. 193-225). San Diego: Academic Press.
- Hole, G. J. (1996). Decay and interference effects in visuospatial short-term memory. *Perception*, *25*, 53-64.
- Hopfinger, J. B., Buonocore, M. H., & Mangun, G. R. (2000). The neural mechanisms of top-down attentional control. *Nature Neuroscience*, *3*, 284-91.
- Hwang, G., Jacobs, J., Geller, A., Danker, J., Sekuler, R., & Kahana, M. J. (2005). Eeg correlates of verbal and nonverbal working memory. *Behavioral and Brain Functions*, *1*, 1-13.
- Kahana, M. J., & Sekuler, R. (2002). Recognizing spatial patterns: A noisy exemplar approach. *Vision Research*, *42*, 2177-2192.
- Kahana, M. J., Zhou, F., Geller, A., & Sekuler, R. (in press). Lure-similarity affects visual episodic recognition: Detailed tests of a noisy exemplar model. *Memory & Cognition*.
- Lalonde, J., & Chaudhuri, A. (2002). Task-dependent transfer of perceptual to memory representations during delayed spatial frequency discrimination. *Vision Research*, *42*, 1759-1769.

- Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within subject designs. *Psychonomic Bulletin & Review*, *1*, 476-490.
- Luck, S., Chelazzi, L., Hillyard, S., & 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.
- MacLeod, C. M. (1998). Directed forgetting. In J. M. Golding & C. M. MacLeod (Eds.), *Intentional forgetting: Interdisciplinary approaches*. (p. 1-57). Mahwah, NJ: Lawrence Erlbaum Associates.
- Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Ford, R. L. (1997). On the regulation of recollection: The intentional forgetting of stereotypical memories. *Journal of Personality & Social Psychology*, *72*, 709-719.
- Magnussen, S., Greenlee, M. W., Asplund, R., & Dymes, S. (1991). Stimulus specific mechanism of visual short-term memory. *Vision Research*, *31*, 1213-1219.
- Magnussen, S., Greenlee, M. W., & Thomas, J. P. (1996). Parallel processing in visual short-term memory. *Journal of Experimental Psychology: Human Perception & Performance*, *22*, 202-212.
- Mangun, G. R. (1995). Neural mechanisms of visual selective attention. *Psychophysiology*, *32*, 4-18.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39-57.
- Ostendorf, F., Finke, C., & Ploner, C. (2004). Inhibition of visual discrimination during a memory-guided saccade task. *Journal of Neurophysiology*, *92*, 660-664.
- Parmentier, F., Tremblay, S., & Jones, D. (2004). Exploring the suffix effect in serial visuospatial short-term memory. *Psychonomics Bulletin & Review*, *11*, 289-295.
- Pelli, D. G. (1997). The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437-442.
- Pelli, D. G., Robson, J. G., & Wilkins, A. J. (1988). Designing a new letter chart for measuring contrast sensitivity. *Clinical Vision Sciences*, *2*, 187-199.
- Seidemann, E., Zohary, E., & Newsome, W. (1998). Temporal gating of neural signals during performance of a visual discrimination task. *Nature*, *394*, 72-75.
- Sekuler, R., Kahana, M. J., McLaughlin, C., Golomb, J., & Wingfield, A. (2005). Preservation of episodic visual recognition memory in aging. *Experimental Aging Research*, *31*, 1-13.
- Slotnick, S. D., Schwarzbach, J., & Yantis, S. (2003). Attentional inhibition of visual processing in human striate and extrastriate cortex. *Neuroimage*, *19*, 1602-1611.
- Somers, D. C., Dale, A. M., Seiffert, A. E., & Tootell, R. B. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proceedings of the National Academy of Sciences, U S A*, *96*, 1663-1668.
- Sternberg, S. (1966). High speed scanning in human memory. *Science*, *153*, 652-654.
- Sternberg, S. (1975). Memory scanning: New findings and current controversies. *Quarterly Journal of Experimental Psychology*, *27*, 1-32.
- Toms, M., Morris, N., & Foley, P. (1994). Characteristics of visual interference with visuospatial working memory. *The British Journal of Psychology*, *85*, 131-144.
- Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neuroscience*, *24*, 295-300.
- Washburn, D. A., & Astur, R. (1998). Nonverbal working memory of humans and monkeys: rehearsal in the sketchpad? *Memory & Cognition*, *26*, 277-86.
- Watanabe, T., Harner, A., Miyauchi, S., Sasaki, Y., Nielsen, M., Palomo, D., & Mukai, I. (1998). Task-dependent influences of attention on the activation of human primary visual cortex. *Proceedings of the National Academy of Sciences, U S A*, *95*, 11489-11492.
- Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review*, *101*, 34-52.
- Wenzlaff, R., & Wegner, D. (2000). Thought suppression. *Annual Review of Psychology*, *51*, 59-91.
- Zhou, F., Kahana, M. J., & Sekuler, R. (2004). Short-term episodic memory for visual textures: A roving probe gathers some memory. *Psychological Science*, *153*, 112-118.