

A task-irrelevant stimulus attribute affects perception and short-term memory

JIE HUANG

Brandeis University, Waltham, Massachusetts

MICHAEL J. KAHANA

University of Pennsylvania, Philadelphia, Pennsylvania

AND

ROBERT SEKULER

Brandeis University, Waltham, Massachusetts

Selective attention protects cognition against intrusions of task-irrelevant stimulus attributes. This protective function was tested in coordinated psychophysical and memory experiments. Stimuli were superimposed, horizontally and vertically oriented gratings of varying spatial frequency; only one orientation was task relevant. Experiment 1 demonstrated that a task-irrelevant spatial frequency interfered with visual discrimination of the task-relevant spatial frequency. Experiment 2 adopted a two-item Sternberg task, using stimuli that had been scaled to neutralize interference at the level of vision. Despite being visually neutralized, the task-irrelevant attribute strongly influenced recognition accuracy and associated reaction times (RTs). This effect was sharply tuned, with the task-irrelevant spatial frequency having an impact only when the task-relevant spatial frequencies of the probe and study items were highly similar to one another. Model-based analyses of judgment accuracy and RT distributional properties converged on the point that the irrelevant orientation operates at an early stage in memory processing, not at a later one that supports decision making.

The ever-changing multitude of objects and events around us generates a torrent of sensory information. The pressure of this torrent puts a premium on selective mechanisms that can emphasize for further processing things that are task relevant and exclude or limit things that are not. Psychophysical and physiological studies have gone far toward characterizing the mechanisms by which attentional selectivity impacts visual perception (e.g., Boynton, 2009; Desimone & Duncan, 1995; Liu, Slotnick, Serences, & Yantis, 2003; Serences, Schwarzbach, Courtney, Golay, & Yantis, 2004). Attentional selectivity's impact on memory has been studied less than its impact on perception, but the impact on memory is arguably no less important. In fact, failures of attentional selectivity in memory are known to have important consequences, such as permitting task-irrelevant information to infiltrate and undermine memory (e.g., with aging; Hasher, Stoltzfus, Zacks, & Rypma, 1991; Rowe, Hasher, & Turcotte, 2008; Stevens, Hasher, Chiew, & Grady, 2008).

Recently, Yotsumoto and Sekuler (2006) attempted to probe the limits of attentional selectivity in visual short-term memory. Their study showed that when one briefly presented study stimulus in a sequence carried task-relevant information and the other study stimulus carried

task-irrelevant information, subjects' ability to exclude the task-irrelevant information from visual recognition memory was distinctly imperfect. They used Gabors as stimuli; a Gabor, sometimes called a *Gabor patch*, is a visual stimulus in which a sinusoidal luminance profile is windowed by a Gaussian (see Figure 1). In Yotsumoto and Sekuler's study, each trial's pair of study stimuli were Gabors, one oriented upright, the other oriented obliquely. Each stimulus was synthesized by superimposing two orthogonal sinusoidal components. After a stimulus pair had been seen, a subject's short-term memory for spatial frequency was probed with a Gabor whose orientation matched that of the task-relevant study stimulus, but whose spatial frequency either matched or did not match the frequency of the task-relevant study stimulus. Yotsumoto and Sekuler's three experiments revealed that the irrelevant stimulus's spatial frequency consistently influenced subjects' recognition judgments, although to a lesser degree than did the spatial frequency of the relevant stimulus.

Before concluding that Yotsumoto and Sekuler's (2006) results demonstrated a failure of attentional selectivity's influence on memory per se, one must entertain the plausible alternative that their results arose from a failure of selectivity in a precursor to memory (Vogel, Woodman,

J. Huang, jiehuang@brandeis.edu




Figure 1. Examples of 1-D Gabor stimuli (top row and left column) and 2-D Gabor stimuli. The 2-D Gabors comprise superimposed horizontally and vertically oriented 1-D Gabors. Across rows, the spatial frequency of the horizontal component increases in a ratio of 1:2:3; across columns, the spatial frequency of the vertical component does the same. For illustration purposes, the Michelson contrast of each sinusoidal component shown here is higher than the contrast used in the experiments.

& Luck, 2005). In particular, if selective attention actually failed at an earlier, perceptual level, the resultant perceptual binding of relevant and irrelevant information in the raw material on which memory must operate would have produced an apparent failure of attentional selectivity downstream of that failure—namely, in memory. To assess this alternative, we drew on stimuli adapted from those in Yotsumoto and Sekuler but began by directly measuring attentional selectivity in the visual processing of those stimuli. Then, in a separate experiment on selectivity in recognition memory, we adjusted the stimuli to compensate for the measured imperfections in visual selectivity. As a result, any failure of selectivity revealed in the tests of recognition memory with these compensated stimuli could be more confidently attributed to an imperfection in the attentional selectivity of memory. Finally, a model-based treatment of the recognition memory results, supplemented by an analysis of the subjects' reaction times (RTs), made it possible to quantify deficiencies in selectivity of memory and also to characterize the locus at which task-irrelevant

information entered into the computations on which recognition was based.

EXPERIMENT 1 Perception

For use in our experiments, we generated a library of Gabor stimuli. As will be explained below, some of our stimuli were 1-D Gabors, whose 1-D sinusoidal luminance profiles were vertical or horizontal in orientation; other stimuli were 2-D Gabors, each comprising superimposed horizontally and vertically oriented sinusoidal luminance gratings. Our stimuli were adapted from the ones that had been used by Yotsumoto and Sekuler (2006), but our choice of these stimuli was motivated also by the particular advantages that would be afforded by the Gabors' metric properties. First, parametric variation of the spatial frequencies associated with attended and nonattended orientations might allow us to examine selectivity with varying degrees of featural separation between the two orientations. Second, stimulus metric

properties would make it possible to separate selectivity effects related to perception (Experiment 1) from selectivity effects associated with memory (Experiment 2). Third, the metric properties would allow us to adjust the stimuli to minimize individual differences in recognition performance that actually reflect differences in visual discrimination ability (Kahana, Zhou, Geller, & Sekuler, 2007; Zhou, Kahana, & Sekuler, 2004). Finally, the metric properties of our stimuli would make it possible to formulate and test within a model-based framework quantitative, theoretical accounts of interactions between task-relevant and task-irrelevant stimulus attributes.

Experiment 1 measured spatial frequency discrimination thresholds, using an adaptive psychophysical algorithm with two-alternative forced choice judgments. On each trial, a subject judged the relative spatial frequencies of two Gabors presented in rapid succession. When a Gabor comprised both horizontal and vertical sinusoidal components, only one of the orientations was relevant to the discrimination judgment. On some trials, the subjects were instructed to disregard the horizontal component, basing their frequency judgments only on the vertical component; on other trials, the subjects were instructed to disregard the vertical component, making their judgments only on the horizontal component. When a trial's Gabors comprised just one oriented component, either horizontal or vertical, spatial frequency discriminations were based on that component. Since the subjects' discrimination thresholds for 1-D Gabors could not be influenced by the absent orthogonal orientation, their 1-D thresholds provided a baseline measure of perceptual discriminability. Any systematic difference between subjects' Weber fractions for 1-D and 2-D Gabors would be a sign of perceptual interaction between the orthogonally oriented stimulus components.

Method

Subjects

Ten subjects, 3 of them male, whose ages ranged from 18 to 30 years, participated in both experiments. They had normal or corrected-to-normal vision as measured with Snellen targets and normal contrast sensitivity as measured with Pelli-Robson charts (Pelli, Robson, & Wilkins, 1988). The subjects were naive as to the purpose of the experiments, and all were paid for their participation.

Apparatus

Gabor stimuli were generated and displayed using MATLAB 7 and extensions from the Psychophysics Toolbox (Brainard, 1997). Stimuli were presented on a 17-in. LCD computer monitor with a refresh rate of 95 Hz and a screen resolution of $1,600 \times 1,200$ pixels. The screen luminance was linearized by means of software adjustments, and the mean luminance was maintained at 30 cd/m². During testing, a subject sat with head supported by a chinrest and viewed the computer display binocularly from a distance of 57 cm.

Stimuli

Spatial frequency discrimination thresholds of 1-D and 2-D Gabors were measured in separate blocks of trials. The 2-D stimuli were generated by superimposing one horizontal sinusoidal grating and one vertical sinusoidal grating of equal contrast and then win-

dowing the result with a Gaussian. The resulting luminance profile $L_{x,y}$ is given by

$$L_{x,y} = L_{\text{avg}} \{1 + A[\cos(\pi fx) + \cos(\pi gy)]\},$$

where $L_{x,y}$ is the mean luminance; f is the spatial frequency of the stimulus's vertical component, in cycles per degree; g is the spatial frequency of the stimulus's horizontal component; and A is the Michelson contrast (Valberg, 2005, p. 80) for each component. A was fixed at 0.2, a value well above detection threshold. L_{avg} is the stimulus's mean luminance. The 1-D stimuli were 1-D Gabors, comprising just a vertical or just a horizontal component. Each stimulus subtended 6.59° of visual angle at a viewing distance of 57 cm and was windowed by a circular Gaussian function, with a space constant of 1.65° of visual angle. To keep the subjects from basing their judgments on local, retinotopic correspondence between stimuli, the absolute phases of each stimulus's oriented components were shifted by random values ranging from 0 to $\pi/2$.

When thresholds for 2-D Gabors were being measured, the subjects were instructed to attend to one orientation and make judgments on that orientation alone. Each trial started with a fixation point presented at the display's center for 300 msec, which was followed after 300 msec by two sequentially presented Gabors, each for 700 msec, with a 500-msec interstimulus interval. The subjects had to identify the grating, first or second, whose spatial frequency on the relevant orientation was higher; the spatial frequency on the other orientation was task irrelevant. The first and second Gabors were equally likely to have the higher relevant spatial frequency. After each trial, a distinctive tone provided knowledge of response correctness. The subjects' discrimination thresholds on vertical and horizontal orientations in 2-D Gabors were measured in separate blocks of trials. In some blocks, the subjects based their judgments on the vertical spatial frequencies of the stimuli; in other blocks, the task-relevant orientation was horizontal. From trial to trial, the stimulus that had the lower spatial frequency of the two on the relevant orientation was chosen from a uniform random distribution ranging from 0.5 to 5 cycles/deg. This distribution spanned the range of stimulus spatial frequencies that would be used later, in Experiment 2. The difference in spatial frequency between the relevant orientation of the two Gabors on each trial was controlled by an adaptive psychophysical algorithm, QUEST (Watson & Pelli, 1983). The spatial frequencies on the task-irrelevant orientation of two Gabors were always different from one another and varied randomly between 0.5 and 6 cycles/deg. In tests with 1-D Gabors, 5 subjects' thresholds were measured with the vertical orientation being task relevant; for the other 5 subjects, the horizontal orientation was task relevant.

Procedure

Three discrimination thresholds were measured for each subject: (1) a threshold for 1-D Gabors, (2) a threshold on the vertical orientation in 2-D Gabors, and (3) a threshold on the horizontal orientation in 2-D Gabors. For half the subjects, their thresholds for 1-D Gabors were measured first; for the others, their thresholds for 2-D Gabors were measured first, during which the thresholds on vertical orientations in 2-D Gabors were measured first for half of them. Each discrimination threshold was estimated several times, in independent runs of the QUEST algorithm. In each run of the QUEST algorithm, 150 trials were used to estimate the difference between spatial frequencies that produced correct judgments 79% of the time. Independent QUEST runs were carried out until the thresholds from two consecutive runs differed from one another by less than 2%. The number of runs needed to measure a single subject's threshold for 1-D or for 2-D Gabors ranged from two to four across all 10 subjects.

Results and Discussion

Figure 2A summarizes the results of spatial frequency discrimination with 1-D and 2-D Gabors. Because horizontal and vertical 1-D measurements were made with




Figure 2. (A) Experiment 1: Box-and-whisker plots showing the Weber fractions for spatial frequency achieved with 1-D Gabors and with 2-D Gabors. Results for 2-D Gabors are separated according to whether the horizontal or the vertical orientation was task relevant. (B) Experiment 2: Box-and-whisker plots of $P(\text{correct})$ recognition for 1-D and 2-D Gabors. Note that the lower whisker of the 1-D box is difficult to see, since its end overlaps with the lower bound of the box.

just 5 subjects each, the figure aggregates results across those two conditions. The vertical-relevant and horizontal-relevant 2-D Weber fractions were based on 10 subjects each, and results are shown separately for the two. Weber fractions for 1-D Gabors ranged from 9.1% to 22.5% ($M = 14.44\%$, $SD = 5.25\%$). Weber fractions for the vertical component of a 2-D Gabor ranged from 10.6% to 35% ($M = 19.01\%$, $SD = 7.75\%$), and the Weber fractions for the horizontal component of a 2-D Gabor ranged from 12.7% to 29.9% ($M = 19.32\%$, $SD = 5.99\%$). Thresholds measured with 2-D Gabors were significantly higher than thresholds with 1-D Gabors, despite the fact that, nominally, the 2-D measures were based on only one orientation [$t(9) = 2.308$, $p < .05$, and $t(9) = 3.652$, $p < .01$, for the vertical and horizontal components, respectively]. Finally, thresholds of vertical and horizontal components of 2-D Gabors did not differ significantly from one another [$t(9) = -0.183$, $p = .859$]. The significant difference between 1-D and 2-D thresholds demonstrates that the presence of an irrelevant orientation, despite the fact that it was to be ignored, substantially interfered with the visual discrimination of the relevant orientation's spatial frequency.

The experiment produced discrimination thresholds taken with two different kinds of stimuli, 1-D and 2-D Gabors. We wondered about the relationship, if any, between the two sets of measurements. In particular, would subjects who had relatively high thresholds with one kind of stimulus have relatively high thresholds for the other? To determine whether the differences between 1-D and 2-D thresholds were random or reflected some stable difference among individual subjects, we computed the pairwise correlation between the subjects' two sets of threshold measurements. This correlation proved to be highly reliable (thresholds for the horizontal component of a 2-D Gabor and thresholds for a 1-D Gabor, $r = .725$, $p < .009$; thresholds for the vertical component of a 2-D Gabor and thresholds of a 1-D Gabor, $r = .595$, $p < .035$;

thresholds for the horizontal component of a 2-D Gabor and thresholds for the vertical component of a 2-D Gabor, $r = .732$, $p < .008$). These substantial correlations point to consistent individual differences for visual discrimination and reinforce the importance of taking account of such differences when stimuli are generated for tests of recognition memory. As a result, the spatial frequencies of Gabors used in Experiment 2 were tailored to each individual subject's Weber fractions, one for 1-D stimuli and one for each of the horizontal and vertical components of 2-D stimuli.

EXPERIMENT 2 Memory

Experiment 2 assessed short-term visual memory with Sternberg's (1966) recognition paradigm. On each trial, two study Gabors were presented in rapid succession. These were followed by a probe Gabor (**p**), which subjects judged as having been among the study Gabors (*yes*) or not (*no*). Tests were made on both 1-D and 2-D Gabors, in a block design. Although 2-D Gabors comprised both horizontal and vertical sinusoidal components, for each subject only one of the two orientations was relevant to the *yes-no* judgments. All of Experiment 1's subjects also served here. For the 5 subjects whose 1-D horizontal thresholds were measured in Experiment 1, the horizontal component of the 2-D stimuli was designated task relevant in Experiment 2; for the 5 subjects whose 1-D vertical thresholds were measured in Experiment 1, the vertical component was task relevant for the 2-D stimuli here.

Experiment 1 demonstrated that the presence of an irrelevant orthogonal orientation influenced spatial frequency discrimination. In order to neutralize the irrelevant orientation's visual impact in Experiment 2, and to isolate effects associated with memory, the subjects' recognition memory was tested with stimuli whose spatial frequen-

cies were adjusted to compensate for the interference of the irrelevant orientation and the individual differences in spatial frequency discrimination that had been revealed in Experiment 1. For each subject, the spatial frequencies on each orientation were normalized to that subject's discrimination threshold. Specifically, when recognition memory was to be measured with a 1-D stimulus, the stimuli were adjusted for each subject's 1-D threshold from Experiment 1; when recognition was to be measured with a 2-D stimulus, the horizontal component's spatial frequency was scaled by each subject's threshold for the horizontal component of a 2-D stimulus in Experiment 1, and the vertical component's spatial frequency was scaled by each subject's threshold for the vertical component of a 2-D stimulus in Experiment 1.

To explain how stimuli were scaled, consider the scaling operation for some hypothetical subject. Imagine that in Experiment 1, this subject's spatial frequency discrimination had been tested with 1-D Gabors of horizontal orientation (as half the subjects were). Let x be this subject's Weber fraction for 1-D Gabors; let h be this subject's Weber fraction for 2-D Gabors whose horizontal component was task relevant; and let v be the subject's Weber fraction for 2-D Gabors whose vertical component was task relevant. When the design of Experiment 2 required a pair of 1-D stimuli separated by n just noticeable differences (JNDs) in spatial frequency, the ratio of the higher frequency stimulus to the lower frequency stimulus for this subject was set to $(1.0 + x)^n$. Similarly, if the experimental protocol required that a pair of 2-D stimuli be separated by n JNDs for this subject, the ratio of higher frequency to lower frequency for the horizontal component was set to $(1.0 + h)^n$, and the ratio of higher frequency to lower frequency for the vertical component was set to $(1.0 + v)^n$. Since the subjects' 1-D Weber fractions were smaller than the subjects' Weber fractions for either orientation in a 2-D stimulus, these separate adjustments equated 1-D and 2-D differences perceptually. In addition, because all the adjustments were done on an individual-subject basis, the adjustments tended to mitigate individual differences in visual discrimination as possible contributors to any effects that might be seen in the results from Experiment 2 (Kahana et al., 2007; Zhou et al., 2004).

Furthermore, in order to examine memory selectivity with a varying featural separation between the task-relevant and task-irrelevant attributes of study items and probes, the spatial frequency relationships on relevant and irrelevant orientations were varied parametrically and independently of one another. Specifically, on some trials, frequencies on the relevant orientation were highly similar to one another, as were frequencies on the irrelevant orientation; on other trials, frequencies on the relevant orientation were dissimilar to one another, as were frequencies on the irrelevant orientation. On other trials still, the frequencies on one orientation were similar to one another, but frequencies on the other orientation differed from one another. These diverse similarity relationships allowed us to test not only whether irrelevant orientation affected

recognition memory or not, but also how finely tuned the irrelevant information's effect was—for example, how the effect might depend on \mathbf{p} 's summed similarity to each study item.

Method

Apparatus and Stimuli

Except as specified below, the apparatus and stimuli were the same as those in Experiment 1. The subjects' recognition memory was tested with 1-D Gabors (oriented either horizontally or vertically) and 2-D Gabors (one horizontal and one vertical, superimposed). Tests with 1-D and 2-D stimuli were carried out in separate blocks of trials.

For both kinds of stimuli, each trial began with a fixation point presented for 300 msec at the display's center. Then, 500 msec later, two study Gabors, \mathbf{s}_1 and \mathbf{s}_2 , were presented sequentially for 700 msec each, with a 500-msec interstimulus interval. Finally, after an interval of 500 ± 50 msec, a probe Gabor (\mathbf{p}) followed. The subjects judged whether \mathbf{p} 's spatial frequency matched that of either \mathbf{s}_1 or \mathbf{s}_2 , by pressing keys on the keyboard representing *yes* and *no*. The probe stimulus remained visible until the subjects responded but disappeared 1,000 msec after its onset if no response had been made by then. After this 1,000-msec period, an additional 2 sec was allowed for a response to be made. After each trial, feedback was provided to the subjects via brief, distinctive tones signifying a correct response, an incorrect response, or a failure to respond within the time allowed. Two sample trials are shown schematically in Figure 3.

When recognition was being tested with 2-D stimuli, only one orientation of each 2-D stimulus was relevant to the task; the subjects were instructed to attend to that relevant orientation and base their recognition judgments on its spatial frequency alone; the other orientation was irrelevant to the task and was to be ignored. Throughout testing, the vertical spatial frequency was relevant for 5 subjects; for the other 5, the horizontal spatial frequency was relevant. As was mentioned above, the assignment of a task-relevant orientation to each subject preserved the orientation assignment with which that subject's 1-D threshold had been measured in Experiment 1.

In order to force the subjects to base recognition judgments on the most recently seen stimuli, the spatial frequencies of all stimuli— \mathbf{s}_1 , \mathbf{s}_2 , and \mathbf{p} —varied from trial to trial, subject to the design constraints described in the next two paragraphs. Furthermore, the relevant spatial frequencies of \mathbf{s}_1 and \mathbf{s}_2 differed always by four JNDs, with \mathbf{s}_1 's frequency higher than \mathbf{s}_2 's on half the trials, and the relationship reversed on the remaining trials. Since the subjects' thresholds differed from one another (see Figure 2), the spatial frequency values resulting from the JND scaling differed accordingly. By definition, on *target* trials \mathbf{p} 's relevant spatial frequency replicated the relevant frequency of either \mathbf{s}_1 or \mathbf{s}_2 , with the two alternatives being equally likely. On *lure* trials, \mathbf{p} 's relevant frequency could assume one of three values relative to \mathbf{s}_1 and \mathbf{s}_2 : It could be two JNDs below the lower frequency study item, two JNDs above the higher frequency study item, or midway between the study items' frequencies in JND units. On lure trials, all three possible values of \mathbf{p} 's relevant spatial frequency (relative to the study items' frequencies) were equally likely. These constraints were imposed so that the proportion of false recognitions could be assessed for different metric relationships between \mathbf{p} and study items (Sekuler & Kahana, 2007). To prevent identical stimuli from appearing trial after trial, the constrained selection of spatial frequencies for each trial was perturbed by a random base frequency drawn afresh on each trial. Details of this perturbation will be given in the next two paragraphs.

Design for testing memory with 1-D stimuli. As was mentioned above, the study items, \mathbf{s}_1 and \mathbf{s}_2 , always differed by four JNDs in spatial frequency. On half the trials, \mathbf{s}_1 's frequency was




Figure 3. A schematic example showing stimuli on an example target trial (upper panel) and on an example lure trial (lower panel). In both examples, vertical orientation was relevant; that is, the subject's judgment was to be based on the vertical spatial frequencies. Horizontal orientation was irrelevant. In the example target trial, p matched s_1 on its relevant orientation, but not on its irrelevant one. The durations of various events are indicated on the timeline at the bottom of each panel.

lower than s_2 's frequency; on the remaining trials, this relationship was reversed. The spatial frequency of the lower frequency study item was set to three JNDs above a random base frequency drawn from a continuous uniform distribution whose range was 0.6–1.0 cycles/deg. Since each stimulus comprised just a single orientation, that orientation—horizontal for half the subjects, vertical for the others—was always relevant. When p 's spatial frequency matched that of either s_1 or s_2 , p qualified as a target; if it matched neither, p qualified as a lure.

Design for testing memory with 2-D stimuli. For the 2-D recognition task, spatial frequencies on the relevant orientation were generated as they were in the 1-D recognition task, except that, now, each trial's random base frequency was drawn from a continuous uniform distribution whose range was 0.4–0.8 cycles/deg. As defined above, when p 's spatial frequency on the relevant orientation matched that of s_1 or s_2 , p qualified as a target; if it matched neither, p qualified as a lure. These definitions are consistent with those for 1-D stimuli. But, as the following explains, these two categories (*target* and *lure*) could be subdivided for purposes of analysis into theoretically useful subcategories. These subcategories were defined by the relationship of p 's irrelevant spatial frequency to the irrelevant spatial frequencies of the study items. For the irrelevant orientation, s_1 's spatial frequency and s_2 's spatial frequency were sampled independently and without replacement from the set $\{+1, +3, +5, +7, +9\}$ JNDs, relative to that trial's random base frequency. As a result, s_1 and s_2 could differ on the irrelevant orientation by two, four, six, or eight JNDs. In order to quantify the influence of the irrelevant orientation on recognition of the relevant orientation, we manipulated the spatial frequencies associated with the two orientations, which produced several different quantitative relationships between the two. The possible relationships between relevant and irrelevant spatial frequencies are illustrated in Figure 4, where various subtypes of p are depicted as occupying particular locations in a 2-D space whose horizontal axis is the relevant orientation's spatial frequency and whose vertical axis is the irrelevant orientation's spatial frequency.

When p 's spatial frequency on the relevant orientation matched that of s_1 or s_2 , p qualified as a target; if it matched neither, p qualified as a lure. When a target p matched one study item not only on the relevant orientation, but also on the irrelevant orientation, p qualified as a *supertarget*. Conditions producing such supertargets are represented by stars in Figure 4. When a target p matched one study item on the relevant orientation, but matched neither s_1 nor s_2 on the irrelevant orientation, the p qualified as a *subtarget*. Subtargets could be further divided into two subtypes: For one type of subtarget, if p 's spatial frequency on the relevant orientation matched that of s_1 , p matched s_2 (the other study item) on the irrelevant orientation, and vice versa; for another type of subtarget, p matched neither s_1 nor s_2 on the irrelevant orientation. Analogous constraints distinguished a *superlure* p from a *sublure* p . Specifically, a lure p was a superlure if its spatial frequencies on both the relevant and irrelevant orientations differed from the corresponding frequencies of both s_1 and s_2 . A lure p qualified as a sublure if its spatial frequency on the irrelevant orientation matched that of s_1 or s_2 (but by definition, there was no match on the relevant orientation). The introduction of these various p subtypes made possible a fine-grained examination of the irrelevant orientation's effect on recognition. In particular, we were interested in comparing the recognition performances for the supertarget and subtarget and the false recognition rates for the superlure and sublure.

Procedure

Each subject served in four sessions, one devoted to measurements of recognition with 1-D stimuli and three devoted to measurements with 2-D stimuli. The order of the four sessions was randomized anew for each subject. In 1-D recognition testing, a subject was tested with a 1-D grating whose orientation corresponded to that subject's designated relevant orientation. For such testing, 212 trials were run in a single session, with the first 20 trials being practice. Each session with 2-D stimuli comprised 380 trials, the first 20 trials of each being practice. Over the three sessions with 2-D stimuli, a subject was tested on a total of 1,080 trials. For each subject, the




Figure 4. Diagram illustrating various subtypes of **p** stimulus. The horizontal axis represents stimulus spatial frequency in just noticeable difference (JND) units on the relevant orientation; the vertical axis represents stimulus spatial frequency in JND units on the irrelevant orientation. In the diagram's example, s_1 's spatial frequencies are three and seven JNDs on relevant and irrelevant orientations, respectively; s_2 's spatial frequencies are seven JNDs on the relevant orientation and one JND on the irrelevant orientation. A supertarget **p** matches either s_1 or s_2 and does so on both of its orientations; the locations of supertarget **p**s for this example are represented by stars. A subtarget **p** matches either s_1 or s_2 on the relevant orientation but does not match the same study item on the irrelevant orientation; these subtarget **p**s are represented in the diagram by squares. A superlure **p** does not match either the relevant or the irrelevant orientation of s_1 or s_2 ; these **p**s are represented by filled triangles. A sublure **p** matches s_1 or s_2 on the irrelevant orientation, but not on the relevant one; such **p**s are represented by empty triangles.

same orientation was designated as relevant throughout. As was mentioned earlier, for each subject, this orientation—horizontal for half the subjects and vertical for the others—was the same as the orientation with which the subject's 1-D threshold had been measured in Experiment 1.

Results and Discussion

Performance with 2-D stimuli was invariant over the three sessions devoted to such stimuli [$F(2,18) = 0.027$, $p = .973$; one-way ANOVA with the factor of experimental session]. Therefore, subsequent analyses of recognition with 2-D stimuli were based on data averaged over the three sessions. To determine whether the irrelevant orientation affected recognition, we compared performance with 2-D stimuli with performance with 1-D stimuli. Figure 2B shows that results differed between the two conditions: Recognition with 2-D stimuli, although nominally based only on the relevant orientation, was poorer than recognition with 1-D stimuli [$F(1,9) = 7.196$, $p < .025$; one-way ANOVA with the factor of dimension]. These results demonstrate that recognition memory for the relevant orientation's spatial frequency was significantly affected by the presence of the irrelevant orientation. Note that this difference in recognition memory manifested itself de-

spite the fact that the stimulus space used for 2-D stimuli had been expanded, relative to that used for 1-D stimuli, reflecting Experiment 1's demonstrated difference in the visual processing of 1-D and 2-D stimuli.

As a prelude to a more fine-grained analysis of how the irrelevant orientation impacted recognition memory, we compared performance with 1-D stimuli against performance with 2-D stimuli, doing so separately for all target trials and for all lure trials; that is, we ignored differences among the subtypes of each. As might be expected, these overall comparisons produced ambiguous results. For all target trials and for all lure trials, mean recognition performance was marginally poorer with 2-D stimuli than with 1-D stimuli, although each difference was small [$t(9) = 2.115$, $p < .064$, and $t(9) = 1.797$, $p = .106$, for target and lure trials, respectively]. However, the existence of various subtypes of target and lure trials suggests the utility of a more fine-grained analysis of how the irrelevant orientation influenced recognition. To understand the basis of this analysis, recall that the relevant spatial frequency of a nonmatching (lure) **p** stimulus was constrained to take particular values, relative to the study items' spatial frequencies. The possible relevant frequencies for a nonmatching **p** could be (1) two JNDs below the lower frequency study item's relevant frequency, (2) two JNDs above the higher frequency study item's relevant frequency, or (3) midway, in JND units, between the relevant frequencies of the two study items. This variation in **p** made it possible to express Experiment 2's results as *mnemonic functions*, where $P(\text{Yes})$ is plotted against **p**'s relationship to the study items' spatial frequencies. The theoretical value of such functions has been discussed elsewhere (Sekuler & Kahana, 2007).

Figure 5A shows the results as mnemonic functions in which $P(\text{Yes})$ is plotted against the **p**'s frequency on the relevant orientation. Since the actual spatial frequencies of the stimuli varied from trial to trial, the values shown on the x -axis of Figure 5A correspond to those on Figure 4's x -axis, which are the normalized spatial frequencies in JND units, with a value of 1.0 assigned to each trial's lowest possible **p** frequency. To simplify the graph, trials with s_1 higher and s_2 higher have been averaged. Thus, values on the x -axis correspond to the absolute difference between s_1 and s_2 , not their algebraic difference. Since s_1 and s_2 are labeled 3 and 7 JNDs, JND values 3 and 7 on the x -axis are targets, and JND values 1, 5, and 9 are lures. The three curves represent different relationships between **p**'s frequency and study items' frequencies on the irrelevant orientation. More specifically, the parameter of the family of curves is the match between the irrelevant frequencies of **p** and a study item. At each value on the x -axis, the differences among the curves represent the differential influence of the irrelevant orientation. For target probes, $P(\text{Yes})$ was strongly affected by the match or nonmatch in the irrelevant orientation's spatial frequency. As can be easily seen in Figure 5A, supertarget **p**s were recognized more readily than subtargets [$F(2,18) = 16.781$, $p < .001$, and $F(2,18) = 29.191$, $p < .001$, for target **p**s that matched s_1 or s_2 , respectively]. However, for lure probes, $P(\text{Yes})$ responses (false recognitions) did not significantly differ for superlure and for sublure **p**s

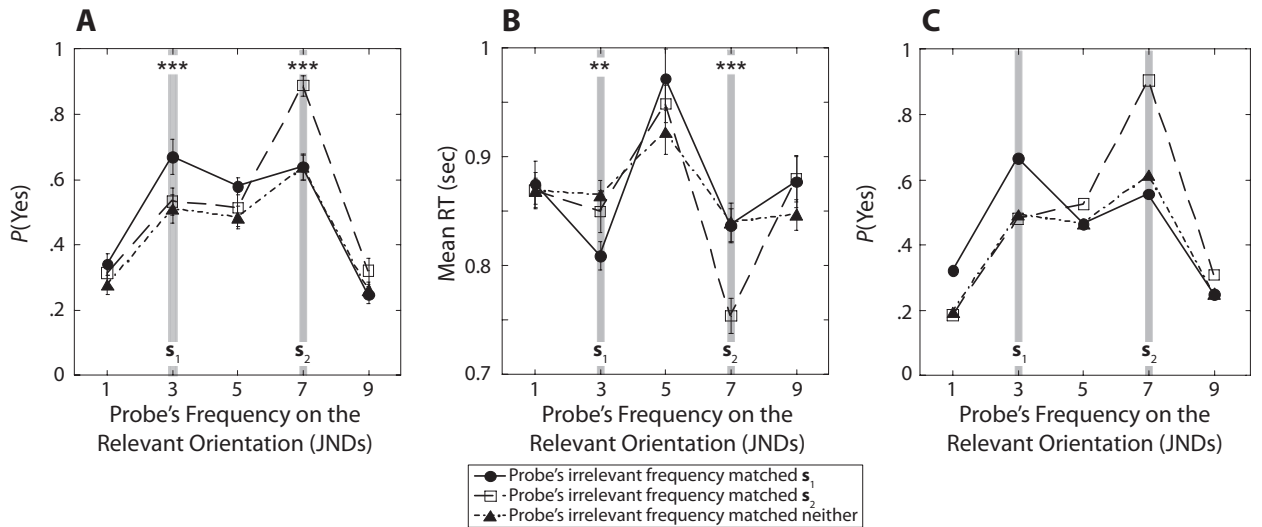


Figure 5. (A) Mnemometric functions showing proportion of *yes* responses as a function of the probe's spatial frequency. (B) Mnemometric functions showing mean reaction time as a function of the probe's spatial frequency. Results shown in both panels are averaged across trials on which s_1 had lower spatial frequency and trials on which s_2 had lower spatial frequency, converged into the case in which s_1 's spatial frequency on the relevant orientation was three JNDs and that of s_2 was seven JNDs, as indicated by the two gray vertical bars. The x-axis is p 's normalized spatial frequency on relevant orientation. The three curves represent conditions in which p matched s_1 , s_2 , or neither on the irrelevant orientation. In panels A and B, error bars represent ± 1 within-subjects standard errors of the means. **Differences significant at $p < .01$. ***Differences significant at $p < .001$. (C) Predicted mnemometric functions based on Model 3. See the text for details. Note the similarity of panels A and C.

[$F(2,18) = 2.867, p = .083$; $F(2,18) = 2.524, p = .108$; and $F(2,18) = 1.704, p = .21$, at JND values labeled 1, 5, and 9 on the x-axis, respectively]. It appears, then, that the irrelevant orientation affected targets but not lures. We will return to this distinction between target and lure trials later, when we consider this distinction in the light of our model-based analysis.

Model-Based Analysis of Irrelevant Information's Effect

In order to evaluate alternative hypotheses about the integration of relevant and irrelevant information, we examined the data through the lens of a global-matching model of recognition memory. In particular, we fit $P(\text{Yes})$ responses from Experiment 2 with several variants of the noisy exemplar model (NEMo), an exemplar-based global-matching model introduced by Kahana and Sekuler (2002). Since our analyses up to this point showed that only target trials were affected by the irrelevant orientation, the model-based analysis only took account of those trials.

The novelty of NEMo lies in its inclusion of η_{s_i, s_j} , a variable representing the similarity of s_i and s_j , the representations of the study items. Specifically, $\eta(s_i, s_j) = e^{-\tau d(s_i, s_j)}$, where d is the weighted distance between the two stimulus vectors and τ determines the exponential generalization gradient. Given a list of study items, $s_1 \dots s_L$, and a probe item p , NEMo will respond *yes* if

$$\underbrace{\sum_{i=1}^L \alpha_i \eta(p, s_i + \epsilon_i)}_{\text{Summed } p \text{ item similarity}} + \underbrace{\frac{2}{L(L-1)} \beta \sum_{i=1}^{L-1} \sum_{j=i+1}^L \eta(s_i + \epsilon_i, s_j + \epsilon_j)}_{\text{Mean interitem similarity}}$$

exceeds the value of the decision criterion, C . The expression's left-hand portion represents the sum of p study item similarities, over all study items; the expression's right-hand portion represents the mean similarity of one study item to another. The importance of both forms of similarity, p item similarity and interitem similarity, in determining recognition responses has been confirmed with sinusoidal gratings (Kahana & Sekuler, 2002; Kahana et al., 2007), chromatic patches (Nosofsky & Kantner, 2006), realistic synthetic human faces (Yotsumoto, Kahana, Wilson, & Sekuler, 2007), and spectro-temporally complex sounds (Visscher, Kaplan, Kahana, & Sekuler, 2006).

When variants of NEMo were fit to the experimental results, the relevant and irrelevant information was allowed to be integrated according to either a Euclidean metric (Minkowski 2-norm distance) or a city block metric (Minkowski 1-norm distance). This choice of alternative metrics was guided by research that linked perceptual integrality and perceptual separability to Euclidean and city block metrics, respectively (Garner, 1974; Lockhead, 1966). In evaluating these metrics, we were mindful of the fact that however separable the dimensions might be on a perceptual level, such separability is not necessarily preserved through subsequent levels of processing (Maddox & Ashby, 1996). Also, when fitting NEMo to the data, we evaluated alternative hypotheses about the route by which irrelevant information influenced recognition. Specifically, different model variants allowed the irrelevant information to enter into the summed p study item similarity, into interitem similarity, into both, or into neither.

Table 1
Key Operations That Define Alternative Models

		Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
Integration		None	Euclidean	Euclidean	Euclidean	City block	City block	City block
Relevant	p item similarity	*	*	*	*	*	*	*
	Interitem similarity	*	*	*	*	*	*	*
Irrelevant	p item similarity		*	*	*	*	*	*
	Interitem similarity		*			*		
Number of free parameters		6	10	10	9	10	10	9

Note—Task-irrelevant information either was omitted from the model altogether (Model 1) or was integrated with task-relevant information, either in a Euclidean metric (Models 2–4) or in a city block metric (Models 5–7). “**p** item similarity” and “interitem similarity” could be defined on the relevant orientation alone, or on both of the relevant and irrelevant orientations. Asterisks indicate that the operation was included in the model. Note that the number of operations included in a model is not the same as the number of free parameters included in that model, as is explained in the text.

Table 1 lists the operations and the number of free parameters included in each model that was evaluated. The complete set of 10 parameters estimated, together with the range of parameter values explored for each, is $\epsilon_{1,R}$, $\epsilon_{2,R}$ (s_1 and s_2 's noise in memory representation on relevant orientation [0,2]), $\epsilon_{1,IR}$, $\epsilon_{2,IR}$ (s_1 and s_2 's noise in memory representation on irrelevant orientation [0,3]), α_1 (degree of forgetting for s_1 relative to s_2 [0,1]), τ (exponential gradient for similarity function [0,3]), β (coefficient for interitem similarity [-2,2]), σ_1 , σ_2 (weights for s_1 and s_2 's irrelevant orientation [0,1]), and C (decision criterion [0,0.5]).

We can clarify the parallels between the operations represented in a model and the free parameters included in that model by considering those parallels in the context of Models 2, 3, and 4. When irrelevant information was allowed to enter into a model, the four free parameters associated with that operation ($\epsilon_{1,IR}$, $\epsilon_{2,IR}$, σ_1 , and σ_2) had to be included in the model. This was true whether the irrelevant orientation contributed to summed **p** study item similarity, to interitem similarity, or to both. As a result, Models 2 and 3 each included the same number of free parameters, despite the fact that the irrelevant information did not contribute to interitem similarity in Model 3. Similarly, whenever interitem distance was included among a model's operations, whether that distance was defined on the relevant orientation alone or on both the relevant and irrelevant orientations, the free parameter β was included in order to represent its contribution to the similarity computation. As a result, because the interitem similarity operation was omitted from Model 4, that model had 9 rather than 10 free parameters because the omission of intersimilarity mandated the omission of the single corresponding free-parameter, β .

Recognition results from target trials were aggregated across all subjects; then, for each trial, differences in spatial frequency (in JND units) for each pair of s_1 , s_2 , and **p**

on relevant and irrelevant orientations were calculated. Trials sharing the same difference in spatial frequency were averaged. Finally, models were fit to averaged subjects' recognition responses, $P(\text{Yes})$, using a multiple-restart stochastic hillclimbing algorithm that optimized $\ln(\text{likelihood})$, hereafter designated LLi.

In order to characterize the success of each variant model, observed values of $P(\text{Yes})$ were regressed against the predicted values from each model variant, and r^2 , the proportion of variance accounted for by the model, was found. The number of free parameters differed among the model variants. As a result, r^2 values were not directly comparable. So, for each model, we calculated the Bayesian information criterion (BIC), which takes account of differences in the number of free parameters in various models (Schwarz, 1978). Under the assumption that model errors are normally distributed,

$$BIC = n \cdot \ln\left(\frac{RSS}{n}\right) + k \cdot \ln(n),$$

where k is the number of free parameters to be estimated, n is the number of observations, and RSS is the residual sum of squares from the estimated model. A better-fitting model produces lower BIC values, which implies fewer explanatory variables, a better fit, or both. Root mean square deviation (RMSD) and LLi were also calculated and were evaluated together with BIC and r^2 in an effort to select among the models.

Table 2 presents each model's goodness-of-fit statistics. Of the seven models, the lowest BIC and RMSD values, and the largest LLi and r^2 values (BIC = -71.7, RMSD = .109, r^2 = .611, LLi = -173.4) were produced by Model 3. In this best-fit model, relevant and irrelevant orientations were integrated in a Euclidean metric, with the irrelevant orientation contributing to the value of summed **p** item similarity, but not to the value of interitem similarity. In

Table 2
Goodness of Fit for Various Models Tested

Index	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
LLi	-325.74	-252.91	-173.4	-292.5	-278.63	-275.21	-291.65
BIC	-69.423	-54.096	-71.700	-55.392	-59.905	-61.004	-63.456
r^2	.214	.497	.611	.559	.390	.396	.389
RMSD	.156	.155	.109	.163	.138	.135	.139

Note—Model 3 was identified by $\ln(\text{likelihood})$ (LLi), the Bayesian information criterion (BIC), and the root mean square deviation (RMSD) indices as the best-fit model.

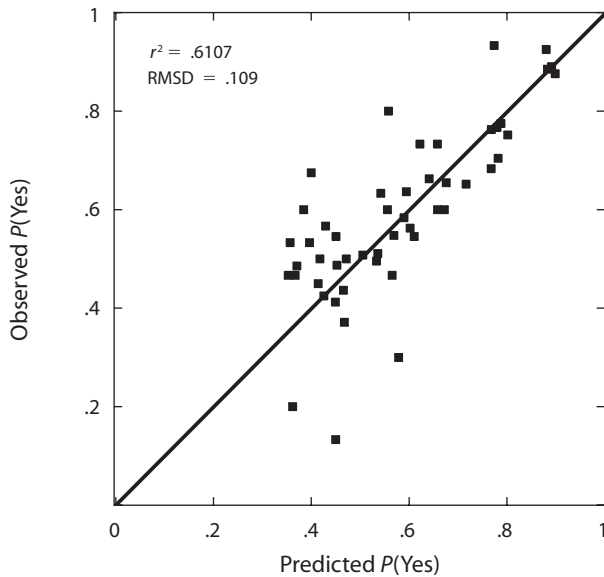


Figure 6. Linear regression of observed proportions of yes responses on target trials against predicted proportions of yes responses by the best-fit model. The model accounts for 61% of the variance.

Figure 6, the observed values of $P(\text{Yes})$ are plotted against the values of $P(\text{Yes})$ predicted by the best-fit model. In our model selection process, we compared other models against Model 3, focusing on the two competing models with the highest r^2 , Models 2 and 4. Note, first, that Models 2 and 3 differed in their treatment of the irrelevant orientation’s contribution to the computation of interitem similarity. Specifically, Model 2 excluded the irrelevant information from interitem similarity, but Model 3 included it. Note, however, that the two models had the same number of free parameters (10), which allows a straight-up comparison of goodness of fit in which Model 3 excelled. Next note that Model 4 was nested within Model 3. We calculated $-2\ln\lambda$ for this pair of two models. This statistic is asymptotically distributed as χ^2 , with degree of freedom equal to the difference in the models’ numbers of free parameters. Comparing Models 3 and 4 showed that Model 3, the more complex model, fit the data significantly better than did its competitor ($p < .001$). Note that this advantage was preserved even when the more complex model was penalized (by BIC) for its additional free parameter.

Equation 1 summarizes the best-fit model, including the model’s portrayal of the way that the irrelevant orientation exerted its influence. In this equation, $d_{1p,R}$, $d_{2p,R}$, $d_{12,R}$ are distances in spatial frequency space, expressed in JND units, between the relevant spatial frequencies of s_1 and p , s_2 and p , and s_1 and s_2 , respectively. The corresponding distances between those same pairs’ spatial frequencies on the irrelevant orientation are given by $d_{1p,IR}$, $d_{2p,IR}$, for s_1 and p , and s_2 and p . Note that on any single trial, the *effective* pairwise distance between two stimuli is determined not only by the mean spatial frequency difference (in JNDs) between the stimuli, but also by the noise associated with each stimulus’s representation. Since each trial’s noise val-

ues are random samples from a distribution, the effective distance between stimuli will vary from trial to trial. It is important to note that the irrelevant orientation for study items, s_1 and s_2 , is not represented among Model 3’s operations for interitem similarity. That is, the best-fit model specifically excludes that operation in computing interitem similarity. Values of σ_1 and σ_2 are the weights imputed to s_1 and s_2 ’s irrelevant orientation; these values are expressed as proportions of 1.0, the value that the model assigns to each study item’s weight on the relevant orientation. The value of α_1 represents the degree of forgetting of s_1 , the study item presented first on a trial. Expressed as a fraction of 1.0, the value imputed to s_2 , α_1 would capture any recency effect seen in the data—that is, if accuracy of recognition were systematically higher for s_2 than for s_1 . Finally, β represents the contribution of interitem similarity:

$$\underbrace{\alpha_1 e^{-\tau\sqrt{d_{1p,R}^2 + \sigma_1 d_{1p,IR}^2}} + e^{-\tau\sqrt{d_{2p,R}^2 + \sigma_2 d_{2p,IR}^2}}}_{\text{Summed } p \text{ item similarity}} + \underbrace{\beta e^{-\tau d_{12,R}}}_{\text{Interitem similarity}} \quad (1)$$

In the best-fit model, Model 3, the parameter representing retention by serial position, α_1 , had a value of .54 for s_1 , relative to the value, 1.0, that had been imputed to s_2 , the study item that was more accurately recognized. In addition, in Model 3, the weight, σ_1 , attributed to the irrelevant information associated with s_1 , was .853, and the corresponding weight for s_2 was .938. Again, both of these values are referenced to 1.0, the values imputed to the relevant information for each study item.

The structure of Model 3 can explain another interesting aspect of our results—namely, why irrelevant information affected $P(\text{Yes})$ on target trials but had little or no effect on lure trials (see Figure 5A). As Equation 1 shows, in Model 3, irrelevant and relevant information is integrated in Euclidean fashion and then enters into the model’s summed-similarity term. The model endorses a p with a *yes* response if the value of the expression in Equation 1 exceeds the model’s decision criterion. Consider how this scenario would apply—first, on target trials, and then, on lure trials. By definition, on a target trial, the value of p item similarity on the relevant orientation will tend to be high for one of the study items, although not necessarily so high as to exceed the criterion all on its own. Starting from that already-high base value, the criterion could be exceeded if there were even a relatively modest boost from the p item similarity on the irrelevant orientation. As a result, the irrelevant orientation should have an impact on target trials. Now consider a lure trial. By definition, on a lure trial, the spatial frequency on p ’s relevant orientation must differ from the corresponding spatial frequency of each of the study items. So, the value of the p item similarity on that relevant orientation will tend to be low, which means that in order for the decision criterion to be exceeded, the p item similarity contributed by the irrelevant orientation would have to compensate by being high. So when the value of p item similarity on the relevant orientation is low, as it would tend to be on lure trials, the accompanying similarity value from the irrelevant orientation would have to be quite high in order for either Euclidean sum and, ultimately, the value

of summed similarity to exceed the criterion. As a result, any possible contribution that irrelevant information might make would tend to be inadequate for boosting the value of summed similarity over the criterion. This means that the irrelevant information would tend to be ineffectual on lure trials. Thus, the structure of Model 3 does indeed predict an effect much like the one we observed: The contribution, if any, of irrelevant information would be small on lure trials, or more precisely, smaller than it would be on target trials.

In order to illustrate quantitatively Model 3's ability to predict a smaller effect on lure trials than on target trials, we ran Model 3 on all the trials on which the human subjects had been tested. This produced a set of predicted values for $P(\text{Yes})$, which are plotted as mnemometric functions in Figure 5C. As can be easily seen, the pattern of the predicted data resembles the pattern of $P(\text{Yes})$ obtained from subjects (see Figure 5A). In particular, for a target probe (corresponding to 3 and 7 on the x -axis), when the spatial frequencies on both orientations match those of a study item, either s_1 or s_2 , the predicted $P(\text{Yes})$ is substantially increased. In contrast, for a lure probe (corresponding to 1, 5, and 9 on the x -axis), when the spatial frequency on irrelevant orientation matches that of s_1 or s_2 , the predicted $P(\text{Yes})$ is only very slightly affected. So, simulations with the model do reproduce the empirical result that the probe's irrelevant orientation had a considerably smaller impact on $P(\text{Yes})$ on lure trials than it had on target trials.

Analyses of Reaction Times

To supplement the model-based analysis of recognition responses, we examined the speed of subjects' judgments—that is, their RTs. Specifically, we carried out two distinct analyses of RTs. The first analysis, which worked with the trimmed-mean RTs from various conditions, was meant to parallel the mnemometric analysis of recognition accuracy shown in Figure 5A; a companion analysis examined the distributional properties of the RTs, interpreting those properties within the framework provided by an ex-Gaussian template (Ratcliff & Rouder, 1998; Rotello & Zeng, 2008) for RT distributions. Of particular interest was whether these two RT analyses would agree with the model-based analysis of $P(\text{Yes})$ as to the irrelevant information's effect on visual memory.

RT-based mnemometric functions. For comparability with the mnemometric functions derived from values of $P(\text{Yes})$, our first RT analysis included RTs from all trials on which subjects responded *yes*—that is, both hits (on target trials) and false alarms (on lure trials). The combination of three match possibilities on the relevant orientation (\mathbf{p} matched s_1 , s_2 , or neither) and three match possibilities on the irrelevant orientation yields a total of nine \mathbf{p} subtypes. To avoid spurious effects on mean RTs that would be produced by extremely quick or extremely slow responses, which could reflect guessing or lapses in attention, a subject's RT distribution for each of the nine \mathbf{p} subtypes was trimmed by excluding the shortest and longest 10% of the RTs. Figure 5B shows the mnemometric functions based on trimmed-mean RTs. Each curve represents a different \mathbf{p} subtype. The \mathbf{p} 's relevant frequency is expressed along the horizontal axis, just as it was in Figure 5A.

Generally, the RT curves in Figure 5B are inverted versions of the corresponding curves in Figure 5A: Where $P(\text{Yes})$ is high, mean RT is short, and vice versa. The one obvious exception to this general rule occurs when the \mathbf{p} 's relevant spatial frequency lies midway between those of s_1 and s_2 (this value is plotted against 5 on the x -axis). The correspondences between the RT results and those from $P(\text{Yes})$ are striking. Thus, RTs to a supertarget \mathbf{p} were shorter than RTs to a subtarget \mathbf{p} [$F(2,18) = 7.091$, $p < .005$, and $F(2,18) = 25.324$, $p < .001$, for \mathbf{p} s that matched s_1 and s_2 , respectively]. In addition, responses to a supertarget \mathbf{p} that matched s_2 were faster than responses to a supertarget \mathbf{p} that matched s_1 ($p < .01$). However, the speeds of the responses to superlure and sublure \mathbf{p} s were not significantly different from one another [$F(2,18) = 0.093$, $p = .912$; $F(2,18) = 1.823$, $p = .19$; and $F(2,18) = 1.121$, $p = .348$, on 1, 5, and 9 JND values on the x -axis, respectively]. So, subtypes of targets that produced significant differences in $P(\text{Yes})$ produced significant differences in mean RTs, and subtypes of lures that failed to produce differences in $P(\text{Yes})$ also tended to fail to produce a significant difference in mean RTs. For the sake of completeness, note that a slightly different picture emerges from examining RTs for correct *no* responses rather than the aggregated *yes* responses shown in Figure 5B. For correct *no* responses, a superlure \mathbf{p} tended to be rejected more quickly than a sublure \mathbf{p} ($p < .022$).

Distributional analysis of RTs. Our second evaluation of the way that the irrelevant orientation influenced recognition judgments examined the distributions of RTs associated with various subtypes of trials. As was explained earlier, the mnemometric functions shown in Figure 5B are based on the means of trimmed distributions of RTs. It is well known that the mean of any distribution is distinctly nonrobust; that is, the mean is highly sensitive to the characteristics of a distribution's tails (e.g., Wilcox, 2003). Although this lack of robustness is a challenge for statistical applications, that same lack of robustness can be exploited for theoretical advantage. Specifically, the mean RT is sensitive to several distinct influences, which, in turn, may correspond to distinct psychological processes. Thus, the means of RT distributions might differ because the distributions' locations have been displaced along the time axis, because the distributions' shapes differed, or because both locations and shapes changed. Having identified significant, systematic variation in mean RTs across conditions, we examined the RTs' distributional properties. In this analysis, we attempted to identify the source of any observed variation in mean RT, adjudicating between (1) changes in distribution location and (2) changes in distribution shape (Rotello & Zeng, 2008). This distinction is important because theorists have offered different functional attributions for changes in location and shape.¹

We fit RT distributions from different conditions to an ex-Gaussian template. As has long been recognized, the ex-Gaussian distribution, which is the convolution of a Gaussian with an exponential distribution, yields a good fit to empirical RT distributions from a wide range of tasks and conditions (Andrews & Heathcote, 2001; Spieler, Balota, & Faust, 2000). For such distributions, μ and σ are the mean

and standard deviation of the Gaussian distribution, and τ is the rate at which the exponential component declines. Variation in the relative values of these parameters generates a family of functions that range from decidedly normal to decidedly exponential. Most important for our purposes, the three parameters have been linked to distinct, theoretically meaningful psychological processes that exert characteristic influences on the empirical RT distribution. Generally, the value of τ is taken to be an index of decision-related processes, and μ is taken to reflect an early input process or processes, such as stimulus encoding, or late output motor response (Ratcliff & Rouder, 1998; Rotello & Zeng, 2008). Consequently, the trio of parameter values that characterize the ex-Gaussian distributions from different experimental conditions makes it possible to identify cognitive processes that differentiate those conditions. Note that under some conditions, a clear failure of ex-Gaussian parameters to differentiate conditions can also carry important theoretical weight (Rotello & Zeng, 2008).

The distributions needed for the ex-Gaussian analysis comprised every RT that was associated with a *yes* response, including both responses on target trials and responses on lure trials. Since extreme RTs pose no obstacle to an ex-Gaussian analysis, we fit ex-Gaussian functions to each subject's raw, untrimmed RT distributions. The combination of nine **p** subtypes and 10 subjects produced 90 RT distributions to be fit. To evaluate goodness of fit, χ^2 values were computed for each distribution's fit to an ex-Gaussian. The mean χ^2 value for all 90 fits was 9.26 ($SD = 6.61$), mean $p = .355$ ($SD = .307$). This overall success in fitting RT distributions to an ex-Gaussian template encouraged us to carry the analysis one step further, comparing the triplets of ex-Gaussian parameter values for different **p** subtypes. We especially wanted to identify any of the three ex-Gaussian parameters that differentiated **ps** that were

supertargets from **ps** that were not; previous analyses, one with values of $P(\text{Yes})$ and the other with values of trimmed mean RTs, revealed substantial differences between these **p** subtypes. In addition, we looked to condition-dependent variation in the ex-Gaussian parameters to support or contradict another finding with the $P(\text{Yes})$ values—namely, that irrelevant information seemed to have little or no effect upon $P(\text{Yes})$ on lure trials.

Figure 7 shows the mean ex-Gaussian parameter values from the best-fit ex-Gaussian models. Values of μ , σ , and τ are shown in the left-hand, middle, and right-hand panels, respectively. Within each panel, values are separated according to which study item matched the **p**'s relevant orientation— s_1 , s_2 , or neither (a lure trial). We applied a contrast vector $-1, 0.5, 0.5$ to compare the parameter value of supertargets against the values for both types of subtargets. For the μ parameter, this contrast was significant for both s_1 and s_2 [$F(1,9) = 9.809, p < .012$, and $F(1,9) = 135.379, p < .001$, respectively]. The same contrast was significant also for the σ parameter, but only for s_2 [$F(1,9) = 6.577, p < .03$], not for s_1 [$F(1,9) = 2.443, p = .152$]. For the τ parameter, the contrast proved to be nonsignificant, for either s_1 or s_2 (both $ps > .05$). With RT distributions from lure trials, none of the three parameters distinguished among different conditions of match to the spatial frequency on the irrelevant orientation. The absence of such effects is consistent with the lack of effect seen earlier, with $P(\text{Yes})$ on various types of lure trials.

Adopting Ratcliff and Rouder's (1998) functional interpretation of the ex-Gaussian parameters, the behavior of the μ parameter suggests that the irrelevant orientation exerted its primary influence early in processing or late in outputting motor response, rather than at a late stage that entails decision making. The lone statistically significant result with the σ parameter confirms that the irrelevant

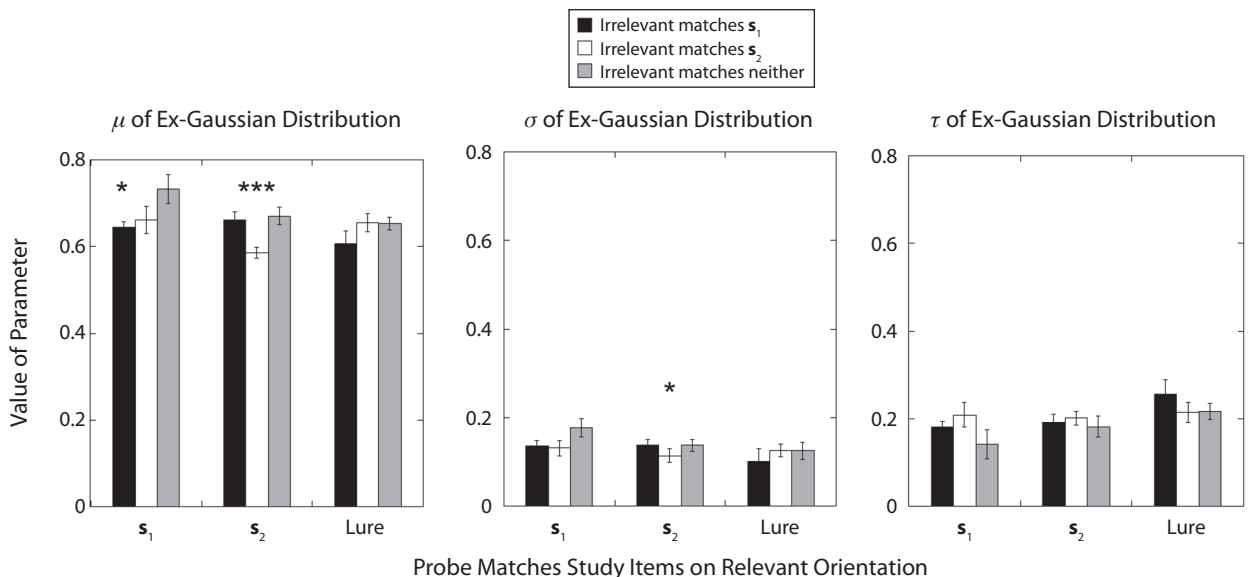


Figure 7. The mean values of the ex-Gaussian parameters for the nine **p** types. The left panel gives values of μ , which is the mean of the Gaussian distributional component; the middle panel gives values of σ , the standard deviation of the Gaussian component; and the right panel gives values of τ , the gradient of the exponential component. ***Differences significant at the level of $p < .001$. * $p < .05$. Error bars are ± 1 within-subjects standard errors of the means.

orientation's clear effect at an early stage of processing is reinforced by some later process—for example, a process that is related to motor output.

GENERAL DISCUSSION

Experiment 1 showed that visual discrimination of one Gabor's spatial frequency is affected by the presence of a spatially overlapping, orthogonally oriented component. This result held despite the fact that the orthogonal Gabor was explicitly task irrelevant, and despite the fact that the subjects had been explicitly instructed to ignore it. This result is generally consistent with previous research in which comparable stimuli were used in a related but somewhat different task (Olzak & Thomas, 1991). Experiment 2 showed that short-term recognition memory for the spatial frequencies associated with a single orientation in two different, successively presented Gabors was also affected by the presence of a task-irrelevant orthogonal orientation. Moreover, our model-based analysis suggested that the task-irrelevant orientation had a potent influence on recognition. Specifically, the model (NEMo) assigned the task-irrelevant orientation's contribution to memory a weight only slightly smaller than the weight assigned to the task-relevant orientation. Also, the influence of the task-irrelevant orientation was selective, with its largest impact coming when the probe item's spatial frequency of the task-relevant orientation was highly similar to the relevant spatial frequencies of the study items. The impact of the irrelevant orientation was selective in another sense as well: The best-fit global-matching model portrayed the irrelevant orientation as contributing to summed similarity, but not to interitem similarity. In addition, model variants that incorporated a Euclidean metric outperformed those that incorporated a city block metric, suggesting that for purposes of recognition, relevant and irrelevant orientations cannot be considered entirely separable dimensions (Garner & Felfoldy, 1970).

The irrelevant orientation's strong but selective effect on recognition memory was not entirely expected. After all, recognition memory was assessed using stimuli that had been scaled to compensate for differences in the visual discriminability of spatial frequency with 1-D and 2-D stimuli. The scaling was intended to discount effects that visual limitations might exert on recognition memory. Moreover, several aspects of Experiment 2 could have strengthened the subjects' ability to filter out the task-irrelevant orientation. First, for each subject, the same orientation (either horizontal or vertical) was task relevant throughout the experiment, which might have allowed the subjects to adopt and maintain a consistent strategy. Second, the stability of performance across all three sessions in which memory was tested suggests that the influence of the task-irrelevant orientation did not come from simple lack of practice or from confusion about the task.

So, an explanation is clearly needed for the irrelevant orientation's selective effect on recognition memory. One possibility is that the stimulus scaling was inadequate or imperfect. On this hypothesis, the interference seen in Experiment 2 would reflect some residual, uncompensated aspect of visual discrimination, rather than some process

associated with memory itself. This low-level, vision-based explanation is difficult to rule out with complete certainty. The intended compensation for visual discrimination was based on individual subjects' Weber fractions, which were collected in Experiment 1. Like any psychophysical measures, those Weber fractions came from a probabilistic process. In our case, the Weber fractions represented the mean difference ratio between two spatial frequencies that allowed the higher frequency stimulus to be correctly identified 79% of the time. The Weber fraction's probabilistic character means that, on some trials, compensation based on a Weber fraction might be greater than the minimum necessary but, on other trials, might be less. Given that the Weber fraction corresponded to the 79% point on the psychometric function, trials with a subminimum compensation would be far fewer than trials with a superminimum compensation. However plausible this hypothesis might be, without supplementary assumptions, it cannot account for the frequency dependence of the interference, nor can it account for the pattern of selectivity revealed by the model-based analysis of the results.

An alternative explanation for the task-irrelevant orientation's impact on recognition focuses on differences between task requirements for visual discrimination and for recognition memory. Because our recognition memory task required that two items be held in memory, the task entailed a higher cognitive load than the subjects faced in making a simple visual discrimination. Specifically, having to maintain two different items in memory in order to make a judgment, as was the case in Experiment 2, likely demands more attention and effort than having to maintain just one item, as was the case in Experiment 1. Lavie and colleagues have shown that task demands on cognitive load modulate the influence of irrelevant distractor stimuli on visual perception and short-term memory (Bahrami, Carmel, Walsh, Rees, & Lavie, 2008; Lavie, Hirst, de Fockert, & Viding, 2004). Specifically, they have demonstrated that increased cognitive load amplifies the effect of a distractor stimulus over what it would have otherwise. Analogous effects of cognitive load have been observed with functional neuroimaging (de Fockert, Rees, Frith, & Lavie, 2001). These demonstrations of cognitive load effects support the idea that the higher cognitive load in Experiment 2's memory task could have reduced the subjects' ability to filter out task-irrelevant information that would have been successfully rejected under the lower cognitive load of Experiment 1's visual discrimination task. Note, however, that key aspects of our results defy explanation either by some load-dependent attenuation of filtering or by incomplete compensation for stimulus discriminability. Either one of these could explain some general failure in the intentional exclusion of the task-irrelevant orientation from memory. However, neither of these hypotheses could explain the selectivity of that failure. Specifically, neither hypothesis can explain why task-irrelevant information should affect target trials but not lure trials, a selectivity that was seen with recognition accuracy and with RTs. As was explained earlier in the Model-Based Analysis of Irrelevant Information's Effect section, this selectivity is a natural consequence of the global-matching process embodied in the best-fit model of recognition memory.

The analysis of mean RTs showed that in some conditions, the spatial frequency on the task-irrelevant orientation significantly speeded recognition responses. Interestingly, this effect paralleled the same independent variable's effect on $P(\text{Yes})$. To confirm the stage or stages of processing at which irrelevant information exerted its effect, we fit empirical response latencies to an ex-Gaussian template. By identifying the distributional parameters that distinguished conditions in which irrelevant information was or was not effective, we sought to distinguish among effects operating at different stages of processing. As was explained earlier, among the three parameters that define an ex-Gaussian distribution, μ reflects mean latency, σ reflects variability in latency, and τ reflects latency extremes. A number of researchers have claimed that an ex-Gaussian dissection captures task-dependent changes in the distinct functional determinants of RT (e.g., Andrews & Heathcote, 2001; Hockley, 1984; Ratcliff & Rouder, 1998; Spieler et al., 2000). Hohle (1965) proposed that the three ex-Gaussian parameters correspond to processes associated with distinct stages of the overall recognition process. In particular, it has been claimed that processing at the input end (*stimulus encoding*) and output end (*motor response*) are reflected in the Gaussian parameters (μ and σ), whereas central, decision-related processes are captured by the exponential parameter (τ). In our study, the speedup of recognition judgments when both relevant and irrelevant orientations of the probe matched one study item is associated with a significant decrease in the value of the Gaussian variable μ for both s_1 and s_2 and in the value of σ , but only for s_2 . Thus, our working hypothesis is that the irrelevant orientation affects processing's early stages, and perhaps motor output as well. Furthermore, the absence of the irrelevant orientation's influence on the τ parameter suggests that decision processes may not be significantly affected by the stimulus's irrelevant information. This finding comports with the structure of the model that provided the best fit to the results based on $P(\text{Yes})$. In that model also, the irrelevant orientation was not allowed to affect the decision process, which, in that model, is represented by the interitem similarity operation. Before making firm conclusions on the basis of the ex-Gaussian template, we must acknowledge that some have raised serious questions about the validity of attempting to link particular ex-Gaussian parameters to distinct psychological processes (Matzke & Wagenmakers, 2009).

We can refine the identification of the stage at which task-irrelevant information influences subjects' responses by adopting the perspective of a global-matching model for recognition memory. Some global-matching models for recognition memory postulate that judgments that the probe is *old* or *new* are based on \mathbf{p} 's summed similarity to all study items. When the value of summed similarity exceeds some criterion value, the probe is judged *old* (e.g., Kahana & Sekuler, 2002; Nosofsky, 1991; Yotsumoto et al., 2007). Kahana and Sekuler's NEMo introduced a novel extension of this basic framework. In NEMo, the value of summed probe-study item similarity on any trial is supplemented by another variable, the value of similarity among the study items. Generally, this variable reduces $P(\text{Yes})$ when study item sets are homogeneous (Kahana & Sekuler, 2002; Kahana et al., 2007). Nosofsky and Kantner (2006) suggested that homogeneity,

as depicted by the interitem similarity, affects the criterion that subjects employ in deciding whether a familiarity signal is sufficiently strong to justify a *yes* response. It is important to recognize that the distinction between the two categories of similarity signals embodied in NEMo is more than a mere formalism embodied in an equation for a model. In fact, recent research confirms the functional distinction between (1) the similarity of the probe and study items and (2) the similarity of study items to one another. Thus, Visscher, Kahana, and Sekuler (2009) demonstrated a clear functional separation in the trial-to-trial carryover of the two putative similarity signals: Probe item similarity on one trial affected $P(\text{Yes})$ on the subsequent trial, but interitem similarity's effect was confined to just the single trial on which the signal was generated. Moreover, using intracranial measurements of brain oscillations captured during performance of a recognition task, van Vugt et al. (2009) demonstrated distinct correlates of these two kinds of similarity signals.

Because each of the two categories of similarity signals—summed probe item similarity and interitem similarity—affects recognition, the observed effect of the irrelevant orientation in Experiment 2 could have arisen from any of several different sources. To select among them, we generated several alternative forms of NEMo and evaluated each one's ability to account for the results of Experiment 2. We wanted to know (1) whether the irrelevant orientation contributed to the computation of summed similarity, interitem similarity, neither, or both, and (2) in what way the two orientations were integrated, in a Euclidean or in a city block metric. Our model selection process identified the preferred model as one in which the irrelevant orientation contributed only to the value of summed similarity and in which the relevant and irrelevant information was integrated in Euclidean fashion (see Equation 1). It seems, then, that the irrelevant orientation's spatial frequency contributes to the value of summed similarity, which reflects the study items' representations in memory. On a target trial, by definition, the spatial frequency of the probe's relevant orientation must replicate that of one of the study items; if the spatial frequency of the probe's irrelevant orientation also replicates that of the same study item, the resulting summed similarity is increased, thereby increasing the likelihood that the probe will be recognized as *old*. However, the spatial frequency on the irrelevant orientation seems to make no contribution to the value of interitem similarity. This result suggests that attentional filtering selectively excludes irrelevant information, but only from participating in the computation of interitem similarity. The dissociation of task-irrelevant information's effect on the two signals that contribute to visual short-term recognition reveals both the deficiency and the effectiveness of attentional filtering. This irrelevant information's selective effect is consistent with the view that interitem similarity does not contribute to summed similarity but, rather, is linked to some central decision process, such as the setting of a decision criterion (Nosofsky & Kantner, 2006).

AUTHOR NOTE

This research was supported by NIH Grants MH068404 and MH61975. Correspondence concerning this article should be addressed to J. Huang,

Volen Center for Complex Systems, Brandeis University, Mailstop 013, Waltham, MA 02454 (e-mail: jiehuang@brandeis.edu).

REFERENCES

- ANDREWS, S., & HEATHCOTE, A. (2001). Distinguishing common and task-specific processes in word identification: A matter of some moment? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *27*, 514-544.
- BAHRAMI, B., CARMEL, D., WALSH, V., REES, G., & LAVIE, N. (2008). Unconscious orientation processing depends on perceptual load. *Journal of Vision*, *8*(3, Art. 12), 1-10.
- BOYNTON, G. M. (2009). A framework for describing the effects of attention on visual responses. *Vision Research*, *49*, 1129-1143.
- BRAINARD, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433-436.
- CHRISTIE, L. S., & LUCE, R. D. (1956). Decision structure and time relations in simple choice behavior. *Bulletin of Mathematical Biophysics*, *18*, 89-112.
- DE FOCKERT, J. W., REES, G., FRITH, C. D., & LAVIE, N. (2001). The role of working memory in visual selective attention. *Science*, *291*, 1803-1806.
- DESIMONE, R., & DUNCAN, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193-222.
- GARNER, W. R. (1974). *The processing of information and structure*. Hillsdale, NJ: Erlbaum.
- GARNER, W. R., & FELDOLDY, G. L. (1970). Integrality of stimulus dimensions in various types of information processing. *Cognitive Psychology*, *1*, 225-241.
- HASHER, L., STOLTZFUS, E. R., ZACKS, R. T., & RYPM, B. (1991). Age and inhibition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *17*, 163-169.
- HOCKLEY, W. E. (1984). Analysis of response time distributions in the study of cognitive processes. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *10*, 598-615.
- HOHLE, R. H. (1965). Inferred components of reaction times as functions of foreperiod duration. *Journal of Experimental Psychology*, *69*, 382-386.
- 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. S., & SEKULER, R. (2007). Lure similarity affects visual episodic recognition: Detailed tests of a noisy exemplar model. *Memory & Cognition*, *35*, 1222-1232.
- LAVIE, N., HIRST, A., DE FOCKERT, J. W., & VIDING, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General*, *133*, 339-354.
- LIU, T., SLOTNICK, S. D., SERENCES, J. T., & YANTIS, S. (2003). Cortical mechanisms of feature-based attentional control. *Cerebral Cortex*, *13*, 1334-1343.
- LOCKHEAD, G. R. (1966). Effects of dimensional redundancy on visual discrimination. *Journal of Experimental Psychology*, *72*, 95-104.
- MADDOX, W. T., & ASHBY, F. G. (1996). Perceptual separability, decisional separability, and the identification-speeded classification relationship. *Journal of Experimental Psychology: Human Perception & Performance*, *22*, 795-817.
- MATZKE, D., & WAGENMAKERS, E.-J. (2009). Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. *Psychonomic Bulletin & Review*, *16*, 798-817.
- MCGILL, W. J. (1963). Stochastic latency mechanisms. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1, pp. 309-360). New York: Wiley.
- NOSOFSKY, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception & Performance*, *17*, 3-27.
- NOSOFSKY, R. M., & KANTNER, J. (2006). Exemplar similarity, study list homogeneity, and short-term perceptual recognition. *Memory & Cognition*, *34*, 112-124.
- OLZAK, L. A., & THOMAS, J. P. (1991). When orthogonal orientations are not processed independently. *Vision Research*, *31*, 51-57.
- PELLI, D. G., ROBSON, J. G., & WILKINS, A. J. (1988). The design of a new letter chart for measuring contrast sensitivity. *Clinical Vision Sciences*, *2*, 187-199.
- RATCLIFF, R., & ROUDER, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*, 347-356.
- ROTELLO, C. M., & ZENG, M. (2008). Analysis of RT distributions in the remember-know paradigm. *Psychonomic Bulletin & Review*, *15*, 825-832.
- ROWE, G., HASHER, L., & TURCOTTE, J. (2008). Age differences in visuospatial working memory. *Psychology & Aging*, *23*, 79-84.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, *6*, 461-464.
- SEKULER, R., & KAHANA, M. J. (2007). A stimulus-oriented approach to memory. *Current Directions in Psychological Science*, *16*, 305-310.
- SERENCES, J. T., SCHWARZBACH, J., COURTNEY, S. M., GOLAY, X., & YANTIS, S. (2004). Control of object-based attention in human cortex. *Cerebral Cortex*, *14*, 1346-1357.
- SPIELER, D. H., BALOTA, D. A., & FAUST, M. E. (2000). Levels of selective attention revealed through analyses of response time distributions. *Journal of Experimental Psychology: Human Perception & Performance*, *26*, 506-526.
- STERNBERG, S. (1966). High-speed scanning in human memory. *Science*, *153*, 652-654.
- STEVENS, W. D., HASHER, L., CHIEW, K. S., & GRADY, C. L. (2008). A neural mechanism underlying memory failure in older adults. *Journal of Neuroscience*, *28*, 12820-12824.
- VALBERG, A. (2005). *Light vision color*. Hoboken, NJ: Wiley.
- VAN VUGT, M. K., SCHULZE-BONHAGE, A., SEKULER, R., LITT, B., BRANDT, A., BALTUCH, G., & KAHANA, M. J. (2009). Intracranial electroencephalography reveals two distinct similarity effects during item recognition. *Brain Research*, *1299*, 33-44.
- VISSCHER, K. M., KAHANA, M. J., & SEKULER, R. (2009). Trial-to-trial carryover in auditory short-term memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *35*, 46-56.
- VISSCHER, K. M., KAPLAN, E., KAHANA, M. J., & SEKULER, R. (2006). Auditory short-term memory behaves like visual short-term memory. *PLoS Biology*, *5*, e56.
- VOGEL, E. K., WOODMAN, G. F., & LUCK, S. J. (2005). Pushing around the locus of selection: Evidence for the flexible-selection hypothesis. *Journal of Cognitive Neuroscience*, *17*, 1907-1922.
- WATSON, A. B., & PELLI, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, *33*, 113-120.
- WILCOX, R. R. (2003). *Applying contemporary statistical techniques*. San Diego: Academic Press.
- YOTSUMOTO, Y., KAHANA, M. J., WILSON, H. R., & SEKULER, R. (2007). Recognition memory for realistic synthetic faces. *Memory & Cognition*, *35*, 1233-1244.
- YOTSUMOTO, Y., & SEKULER, R. (2006). Out of mind, but not out of sight: Intentional control of visual memory. *Memory & Cognition*, *34*, 776-786.
- ZHOU, F., KAHANA, M. J., & SEKULER, R. (2004). Short-term episodic memory for visual textures: A roving probe gathers some memory. *Psychological Science*, *15*, 112-118.

NOTE

1. A half century ago, researchers noted that many choice tasks generated RT distributions that could be described as a convolution of an exponential and a Gaussian (Christie & Luce, 1956; McGill, 1963). For these exponential and Gaussian components of RTs, researchers proposed functional interpretations, although their proposals did not always agree with one another. Generally, changes in the location of an RT distribution have been attributed to differences in some early, perhaps automatic process associated with stimulus encoding or to late motor response. If the contribution of this process differed between conditions, so too would the location of the resulting RT distributions. In contrast, changes in the RT distribution's shape have generally been attributed to differences in later strategic or decision processes. We must caution, however, that recently, some have questioned the legitimacy of aligning distinct psychological processes with values of parameters from the ex-Gaussian template (Matzke & Wagenmakers, 2009).

(Manuscript received September 26, 2008;
revision accepted for publication May 24, 2009.)