Whatever you do, don't look at the…

Evaluating guidance by an exclusionary attentional template

Valerie M. Beck[1], Steven J. Luck[2], Andrew Hollingworth[1]

[1]University of Iowa

[2]University of California, Davis

Short title: Evaluating the Exclusionary Attentional Template

Word Count: 12,212

Please address correspondence to:

Valerie M. Beck
Department of Psychological and Brain Sciences
University of Delaware
108 Wolf Hall
Newark, DE 19716

e-mail: vbeck@psych.udel.edu
Phone: 302-455-8465

Abstract

People can use a target template consisting of one or more features to guide attention and gaze to matching objects in a search array. But can we also use feature information to guide attention *away* from known irrelevant items? Some studies found a benefit from foreknowledge of a distractor feature, while others found a cost. Importantly, previous work has largely relied on end-of-trial manual responses; it is unclear how feature-guided avoidance might unfold as candidate objects are inspected. In the current experiments, participants were cued with a distractor feature to avoid, then performed a visual search task while eye movements were recorded. Participants initially fixated a to-be-avoided object more frequently than predicted by chance, but they also demonstrated avoidance of cue-matching objects later in the trial. When provided more time between cue stimulus and search array, participants continued to be initially captured by a cued-color item. Furthermore, avoidance of cue-matching objects later in the trial was not contingent on initial capture by a cue-matching object. These results suggest that the conflicting findings in previous negative-cue experiments may be explained by a mixture of two independent processes: initial attentional capture by memory-matching items and later avoidance of known irrelevant items.

Keywords: visual attention, visual search, attentional control, exclusionary template, feature-guided avoidance

Public significance statement: Attention can efficiently be guided toward relevant objects. For example, say you are searching for your friend’s phone. Your friend’s phone used to have a red case so you find your eyes drawn to red objects in the room. However, your friend recently got a new phone case that is not red. Can you use this “not red” information to help you search and avoid looking at red objects? Our work demonstrates that attention is initially drawn toward irrelevant objects (red objects in this example), but these irrelevant objects can also be avoided. Furthermore, later avoidance is not dependent on having attended these irrelevant objects earlier.

Introduction

Most theories of attention propose that a template specifying the features of task-relevant items allows for goal-directed control of selection (e.g., Bundesen, 1990; Duncan & Humphreys, 1989; Wolfe, 1994). Indeed, when participants receive knowledge about a relevant feature prior to search (e.g., a cue specifying that the target will be red), they can largely limit attention to matching items in the search array (Beck, Hollingworth, & Luck, 2012; Green & Anderson, 1956; Vickery, King, & Jiang, 2005; Wolfe, Horowitz, Kenner, Hyle, & Vasan, 2004). Maintaining a representation of the relevant template features requires memory, and because search targets frequently change during real-world behavior, most researchers have proposed that the substrate of the template representation is the visual working memory (VWM) system (Carlisle, Arita, Pardo, & Woodman, 2011; Gunseli, Meeter, & Olivers, 2014; Woodman & Arita, 2011).

To implement template-based guidance, many theories propose that the content of VWM modulates the competition among objects for selection, with a higher attentional weight assigned to features maintained in VWM (Bundesen, 1990; Desimone & Duncan, 1995; Navalpakkam & Itti, 2005; Wolfe, 1994). This raises a key architectural question: Although it is well established that the content of VWM can be used to facilitate the selection of matching items, can the interface between VWM and attentional control be configured so that attention is biased *away* from objects matching VWM content? Do the attentional weights necessarily have to be positive? Or, is it possible to assign a negative attentional weight for a feature value, relative to other features values, so as to implement feature-guided avoidance?

The evidence thus far has been mixed, and this remains one of the central outstanding questions in the field of goal-directed vision. Most studies have demonstrated that when the content of VWM is known to be associated only with distractors, attention is nevertheless captured by matching items (Folk, Remington, & Johnston, 1992; Hollingworth & Luck, 2009; Hollingworth, Matsukura, & Luck, 2013; Olivers, 2009; Olivers, Meijer, & Theeuwes, 2006; Experiment 4 in Soto, Heinke, Humphreys, & Blanco, 2005; Soto, Hodsoll, Rotshtein, & Humphreys, 2008; Soto & Humphreys, 2007), suggesting that participants cannot configure a feature-based, negative template. However, other studies have found some evidence of successful avoidance (Arita, Carlisle, & Woodman, 2012; Woodman & Luck, 2007), as indicated by lower overall search times when memory-matching distractors were present in the search array. In addition, there is growing evidence that people can avoid capture by physically salient objects if these objects contain a predictable feature value (Vatterott & Vecera, 2012), and in some cases this item is suppressed below the level of the other objects in the stimulus array (Gaspelin, Leonard, & Luck, 2015, 2017; Moher, Lakshmanan, Egeth, & Ewen, 2014).

Woodman and Luck (2007) were the first to propose that a VWM representation could be used not only to guide attention toward matching items (a “template for selection”) but also to guide attention away from matching items (a “template for rejection”). Participants were asked to hold a colored square in memory, perform a shaped-defined search task, and then respond to a memory probe. In the critical experiment, the search array contained two items drawn in one color, four items drawn in a second color, and six items drawn in a third color. One of the colors used in the search array always matched the color held in memory, but the number of memory-matching distractors could vary (2, 4, or 6). The search target was a Landolt-C with a gap in the top or bottom and never matched the color in memory. Participants were faster to respond to the target item when there were a greater number of memory-matching distractors (6) than when there were fewer (2 or 4), suggesting they were able to configure a VWM-based “template for rejection” and exclude these memory-matching items from search.

The Woodman and Luck (2007) results are the strongest evidence to date in favor of the capability to configure a negative VWM template. Several other studies have reported converging evidence, but these have been limited in important ways. First, Arita et al. (2012) used a circular search array, with items on the left side presented in one color (e.g., red) and items on the right side presented in another color (e.g., blue). A color cue that preceded the search array could indicate the target item color (positive cue), a distractor color (negative cue), or a color not present in the search array (neutral cue). Arita et al. (2012) found faster response times in the negative-cue condition than in the neutral cue condition, suggesting that participants were able to avoid attending to cue-matching items. However, Beck & Hollingworth (2015) argued that what appeared to be feature-based avoidance could be explained instead by the rapid conversion of the negative feature cue into a simple spatial template (attend left or attend right). When the differently colored items were spatially intermixed, making this conversion strategy more difficult to implement, the response time benefit in the negative-cue condition was eliminated (see also Becker, Hemsteger, & Peltier, 2016).

Moher and Egeth (2012) also proposed that a negative feature cue can support avoidance, but with a caveat. They claimed that avoidance was dependent on directing attention initially to an item or items matching the to-be-avoided color, terming this a “search and destroy” process. Participants were provided a negative cue indicating a distractor color prior to the appearance of a search array. In an initial experiment, response times (RTs) were slower in the negative-cue condition than in a neutral condition, suggesting that participants were attending to the cue-matching distractor even though they knew it was irrelevant. To examine the time-course of selection across the search, Moher and Egeth (2012) used a dot-probe technique and an SOA manipulation to probe the spatial locus of attention early versus late during the search process. They found significantly faster dot-probe RTs at the cue-matching distractor location early during search (117 ms SOA), suggesting that attention was captured initially. They found a non-significant trend toward slower dot-probe RTs at the cue-matching distractor location later during search (167 ms SOA), concluding that initial capture was followed by later avoidance. However, even if the later avoidance effect were robust, the composition of the arrays precluded strong inferences about avoidance. The arrays contained only one cue-matching item. Thus, if attention was initially captured by that item, later "avoidance" may have been the simple consequence of having already attended to it; there need not have been any explicit mechanism of avoidance, just the deployment of attention to the remaining items in the array after initial capture by the cue-matching item.

As an additional test of later avoidance, Moher and Egeth (2012) preceded the search array with a set of placeholders that were the same colors as the search array items. The placeholders were visible for 100, 800, or 1500 ms before the search array appeared. There was a negative-cue cost (relative to a neutral condition) at the 100 ms duration and a negative-cue benefit at the 800 and 1500 ms durations, suggesting that during the placeholder array, participants initially attended to the cue-matching placeholder, but when the placeholders were present for a longer duration, participants had time after initial orienting to de-prioritize those locations. However, it is not clear whether the advantage in search RT at the longer placeholder durations was due to feature-based avoidance per se or due to the ability to mark particular array locations as to-be-avoided (or the complementary set as to-be-attended). That is, participants may have converted the feature information into a spatial template before search commenced, similar to the strategy apparently used in the method of Arita et al (2012; see also Han & Kim, 2009). Note that Moher and Egeth (2012) did not specify a mechanism by which later avoidance was implemented, so this possibility is not necessarily inconsistent with their claims.

Finally, a recent study by Kugler and colleagues (Kugler, ’t Hart, Kohlbecher, Einhäuser, & Schneider, 2015) used an eye tracking method similar to that in the present study and found reliable avoidance of items that matched a negative-cue color. Specifically, half of the items in the array either matched a positive cue or a negative cue. Guidance of gaze toward relevant items was observed for both cue conditions, with more efficient guidance by positive than by negative cues. However, the implications of this result for understanding guidance by VWM are not entirely clear. In Kugler et al., the cued color remained constant across blocks of 10 trials. Thus, guidance by a negative cue in their study may have reflected either a VWM template or a LTM template, as there is a rapid transfer of control from VWM to LTM over the course of several trials with the same search cue (Carlisle et al., 2011).

In sum, the current literature leaves open several key issues. First, it has yet to be determined if participants can use a VWM representation of a negative feature cue to generate any type of avoidance of cue-matching objects, except in the limited circumstance that the feature cue can be converted efficiently into a spatial template (Arita et al., 2012; Beck & Hollingworth, 2015; Han & Kim, 2009; Moher & Egeth, 2012). The results of Woodman and Luck (2007) suggest that such avoidance might be possible, but similar paradigms have produced conflicting results (Olivers, 2009; Olivers et al., 2006; Experiment 4 in Soto et al., 2005; Soto & Humphreys, 2007). Additionally, little is understood about how VWM-guided selection evolves over the course of a trial, either for a positive template or for a negative template. The Moher and Egeth (2012) results suggest that negative templates produce a pattern of initial capture and later avoidance. However, as discussed above, their method yielded ambiguous results with respect to later avoidance. Finally, if such a capture/avoidance pattern were observed, it would need to be determined whether there is a functional relationship between early capture and later avoidance.

In the following three experiments, we examined these issues by recording eye movements while participants performed a visual search task. Prior to the appearance of the search array, they saw either a positive cue (the color of the target), a negative cue (the color of some of the distractors), or a neutral cue (which provided no information about target or distractor colors). Eye tracking during search provided a real-time window on the evolution of selection throughout the trial and made it possible to capture the object-by-object pattern of selection during search. In this manner, we examined the time-course of attentional guidance by positive and exclusionary templates, which we quantified in terms of the probability that a given fixated object matched the cued attribute. In other words, for each moment in time following the onset of the search array, we could assess the probability that attention was directed toward an object of the positively or negatively cued color. Moreover, each search array included multiple objects of the cued color; this allowed us to examine the possible effect of early capture of attention by an object of the (positively or negatively) cued color on the probability of selecting other cue-matching objects later in the trial.

Experiment 1

Experiment 1 was designed to determine whether participants are able to use information about a non-target color to exclude matching items from visual search and how selectivity develops across the trial. In the basic search task used in all three experiments (illustrated in Figure 1), participants viewed an array of circles drawn in different colors. Distractors had a gap on the left or right. The target had a gap on the top or bottom, and participants reported gap location. The gaps were very small, making it difficult to use gap location to guide attention.

Before the onset of the search array, a colored cue was displayed. In the *cue-target* condition, the cue indicated the color of the target item, as in traditional guided search tasks. Participants could use cue information to select cue-matching items during search. In the *cue-avoid* condition, the cue indicated one of the colors in which the target would *not* be drawn. Participants could potentially use the cue information to avoid selecting cue-matching items. Finally, in the *cue-all* condition, the cue was a composite of all possible colors and thus conveyed no information. These different conditions were tested in different trial blocks, but the colors of the target and distractors varied unpredictably from trial to trial.

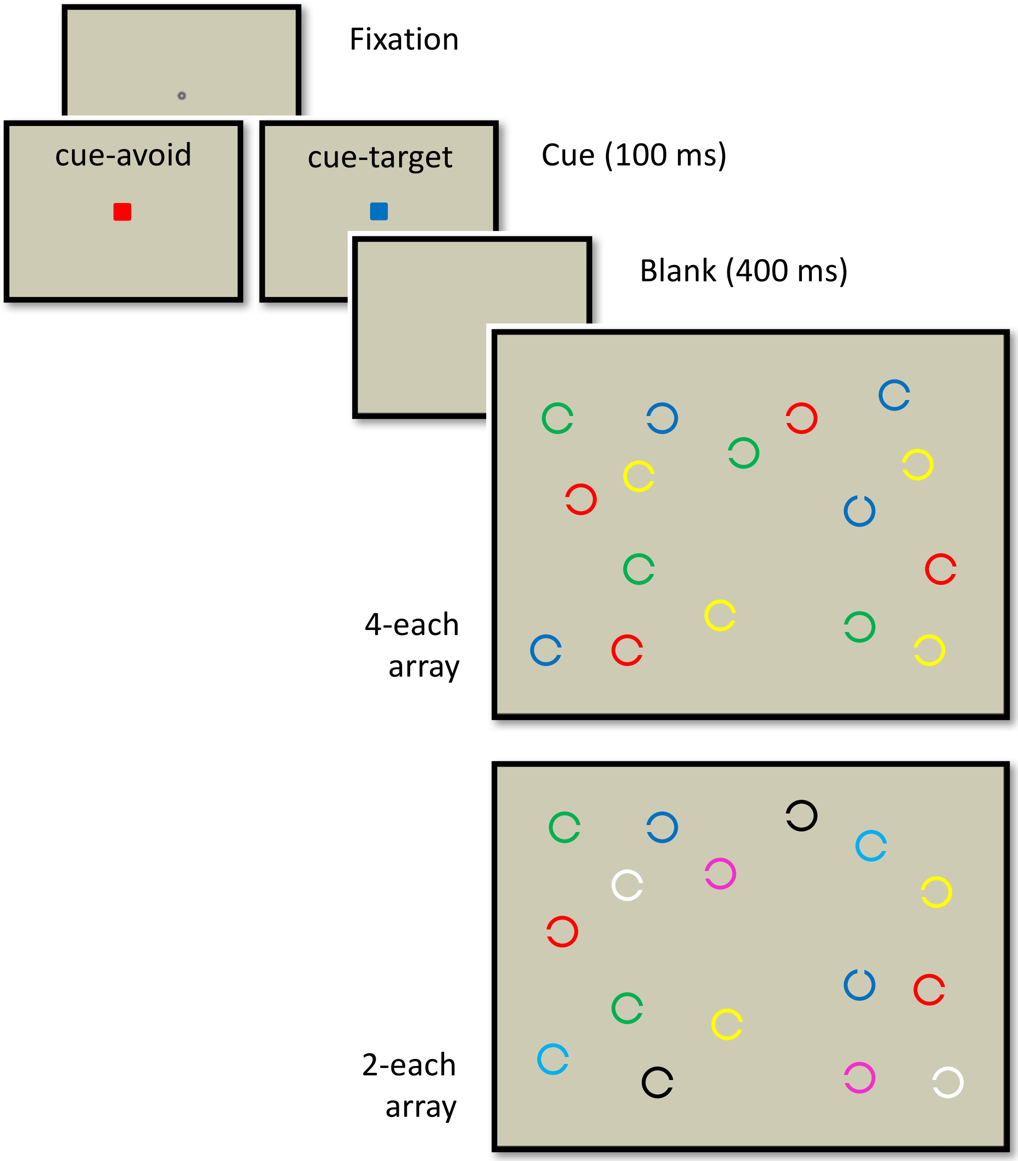

Figure 1: Example trial sequence and search arrays for Experiments 1-3. Participants were instructed to locate the Landolt-C with a top or bottom gap and report the gap location. The cue stimulus could indicate either the color of the target item (cue-target), the color to avoid (cue-avoid), or that the target item could be any color (cue-all; not shown). The search array could contain either four each of four different colors (4-each) or two each of eight different colors (2-each). Cue condition was blocked and the type of search array was intermixed.

Method

Participants. To ensure sufficient sample size, we examined a range of similar studies that have used object fixation as a direct measure of attentional guidance and selection in within-subjects designs (Beck & Hollingworth, 2017; Beck, Hollingworth, & Luck, 2012; Gaspelin, Leonard, & Luck, 2017; Le Pelley, Pearson, Griffiths, & Beesley, 2015; Soto, Humphreys, & Heinke, 2006). For the key analyses in these studies, mean effect size was η_p^2 = .70. Power analysis (using G*power) indicated that six participants would be necessary to achieve 80% power and eight participants to achieve 95% power. Conservatively, our experiments used a sample size of 12. For Experiment 1, the 12 participants (8 female; 18-30 years old) were recruited from the University of California, Davis. They were compensated for their time. All participants reported normal color vision and normal or corrected-to-normal acuity. All procedures were approved by the University of California, Davis's Institutional Review Board.

Stimuli and Procedure. Stimuli were presented on a CRT monitor with a light gray background at a distance of 70 cm. Each search array contained 16 Landolt-C objects (see Figure 1). These objects were 0.67° in diameter, had a line width of 0.10°, and had a gap measuring 0.07°. Objects were placed in random locations on the screen with the following constraints: a minimum distance of 2° from the center of the screen, a minimum distance of 2.07° between objects (center-to-center), and a minimum distance of 2.51° from the edge of the screen. The total visible area of the screen subtended 26.74° x 20.05°, but objects could only appear within an area that subtended 21.72° x 15.03°. New locations were generated for each trial, and the target object was randomly assigned to one of the locations.

The 16 objects in a given array consisted of either four objects of each of four different colors (*4-each*) or two objects of each of eight different colors (*2-each*). Presenting arrays in this manner allowed us to examine search efficiency by varying the number of objects that could be the target (*cue-target*: 4 or 2 items; *cue-avoid*: 12 or 14 items; *cue-all*: always 16 items) without changing the total number of objects on the screen. The eight possible colors were chosen to be highly discriminable: red, yellow, green, blue, white, black, magenta, and cyan. For the *4-each* condition, the four colors on each trial were selected randomly from the set of eight. In a given array, the assignment of each color to each object was determined randomly. Thus, the target color was selected randomly on each trial. There was one target object (top or bottom gap, randomly selected) and 15 distractor objects (left or right gap, randomly selected). Participants reported the target item's gap location by pressing one of two buttons on a game pad. The search array remained visible until the manual response, which terminated the trial. Importantly, the gaps were so small that gap position discrimination required object fixation, and the task therefore required participants to translate covert attentional control into overt shifts of gaze.

At the beginning of each trial, a cue square (0.67° x 0.67°) was presented for 100 ms. After a 400-ms blank ISI, the search array was presented. The cue was either the color in which the target would be drawn (*cue-target*), a color that would be present in the array but would not be the target color (*cue-avoid*), or a checkerboard composed of all possible colors (4x4 grid containing 2 squares of each of the 8 possible colors) indicating that the target item could be any color (*cue-all*; not shown in Figure 1). For the *cue-avoid* condition, the cue square color was selected randomly from the set of distractor colors to appear in the subsequent array. Cue condition was blocked, and block order was randomized across participants. The task began with eight trials of each condition (*cue-target*, *cue-all*, *cue-avoid*), followed by three blocks of 32 trials for each of the three conditions. At the beginning of each condition block, participants received instructions about what the cue signified (Cue = Target Item Color, Cue = NOT Target Item Color, Cue = Target Item is Any Color). The first two trials in each block were considered warm-up trials and were excluded from analysis.

Eye movements were recorded using an Eyelink 1000 eye tracker with a sampling rate of 2000 Hz. Saccades were defined by a combined velocity (>30°/s) and acceleration (>9500°/s^2)

threshold. Gaze position was calibrated using a typical 9-point calibration/validation routine at the beginning of each block and any time the participant failed to meet gaze-contingent fixation criteria at the beginning of a trial. Specifically, each trial began with a gaze-contingent fixation routine that required participants to maintain fixation continuously within a central region (1.67° x 1.67°) for 300 ms, which served to provide a check on tracking accuracy as well as ensure that the participant would see the cue square that appeared at the same location as soon as the gaze-contingent fixation criteria were met.

Data Analysis. For the eye movement analyses, interest areas were defined around each object and the central fixation region. The central fixation interest area was a circle 1.67° in diameter at the center of the screen. Object interest areas were circles centered on each object subtending 2°, which allowed for natural variation of gaze accuracy while also defining non-overlapping regions. An object was considered to be fixated when a fixation occurred within the defined interest area for that object.

As anticipated, manual response accuracy was uniformly high (M = 99% correct) across all conditions (see Table 1 for accuracy by condition and array type); trials with incorrect responses were excluded from all further analyses. All trials with response times that were less than 150 ms or greater than 10,000 ms were excluded from analysis (6.78% of trials). Furthermore, trials with response times that were beyond 2.5 standard deviations from the mean of each condition for each subject were also excluded from all analyses (additional 1.94% of trials).

Experiment	Array Type	Cue Condition		
		Positive (Cue-Target)	Neutral (Cue-All)	Negative (Cue-Avoid)
Experiment 1	4-each	99.0% 1148	98.5% 2696	99.6% 2763
	2-each	97.5% 945	99.6% 2598	99.1% 2683
Experiment 2	4-each	98.9% 1172	99.0% 2595	98.6% 2768
	2-each	98.7% 915	99.6% 2548	99.0% 2977

Table 1:Manual response accuracy and manual response times (ms) by condition and array type for Experiments 1 and 2.

Results and Discussion

We first report end-of-trial measures of search time. Then, we report analyses of object-by-object selection probability, providing the key evidence concerning evolution of selection across the trial. Finally, we report analyses concerning the time taken to initiate search.

Overall Search Time. Because end-of-trial measures cannot illuminate the evolution of selection across the trial, and because such measures have produced contradictory results in previous studies, we did not develop explicit predictions for these analyses beyond the expectation that participants would use the positive cue to improve search efficiency (Beck et al., 2012). Moreover, end-of-trial measures were expected to be more variable than measures reflecting discrete oculomotor selection. Thus, if consistent population effects exist, the present experiments may not have had sufficient power to detect them. Nevertheless, we report inferential statistics to connect our results with the existing literature on this topic.

Overall search time was calculated as the elapsed time until the first fixation in the target region, or time to target fixation (TTF). Manual response time (RT) produced the same pattern of means and statistical significance in all experiments. The results are summarized in Figure 2. We began by comparing the cue-target (positive template) condition against the cue-all (neutral) condition. A condition X array type ANOVA revealed a main effect of condition [$F(1, 11) = 354.93$, $p < .001$, $\eta_p^2 = .97$], a main effect of array type [$F(1, 11) = 22.86$, $p = .001$, $\eta_p^2 = .68$], but no significant interaction. Collapsing across array type, we found faster TTF in the *cue-target* condition ($M = 623$ ms) than in the *cue-all* condition ($M = 2048$ ms), indicating efficient attentional guidance by a positive template (Beck et al., 2012). Similarly, collapsing across cue condition, we found faster TTF for *2-each* ($M = 1241$ ms) than *4-each* arrays ($M = 1430$ ms), possibly driven by the fewer number of relevant items in *2-each* (2) compared to *4-each* (4) arrays in the *cue-target* condition, though this pattern was also found in the *cue-all* condition.[1]

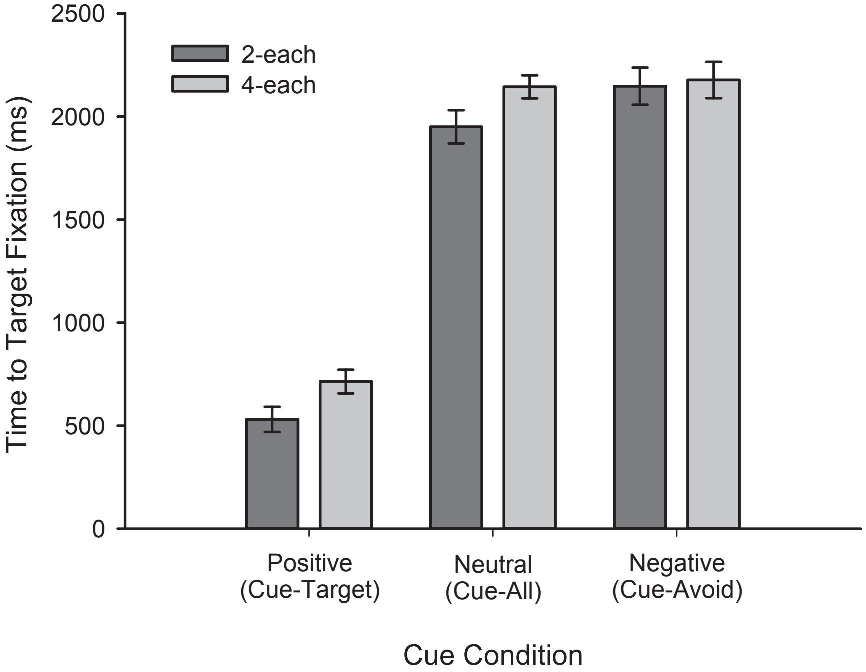

Figure 2: Elapsed time to first fixation on the target item plotted as a function of cue condition (cue-target: Positive; cue-all: Neutral; cue-avoid: Negative) and array type (4-each, 2-each) for Experiment 1 (color cue stimulus). Error bars indicate within-subjects 95% confidence intervals (Morey, 2008).

We conducted a parallel set of analyses to compare the *cue-avoid* (negative template)

---

[1] The reason for this difference between *2-each* and *4-each* arrays in the *cue-all* condition is unclear and likely spurious since we did not observe a similar pattern in Experiment 2.

condition with the *cue-all* (neutral) condition. There were no significant main effects or interaction (all *p*s >.17), suggesting that participants may not have been able to use the cue information in the *cue-avoid* condition to improve search efficiency. Thus, the measures that reflected the total time required for visual search do not provide any clear evidence that participants were able to implement an exclusionary template.[2]

Object-by-Object Analysis of Selectivity. To examine the evolution of selection across the course of a trial, the eye movement data were binned by ordinal object fixated during search (i.e., first object fixated, second object fixated, etc.). That is, the functional unit for the analysis was each object fixated, which may have included multiple individual fixations. For each combination of bin and participant, we compared the observed probability of fixating a cue-matching object against chance by calculating an odds ratio: observed probability of fixating a cue-matching object over the probability of fixating a cue-matching object by chance. Chance probability was calculated for each object in each trial, considering the preceding events on that trial, and then averaged across trials in a bin. To illustrate the calculation of chance probability, consider a trial in the *4-each* array condition in which the first object fixated matched the cue and the second object fixated did not. Given these preceding fixations, the probability of fixating a cue-matching object by chance as the third object would be 3/14 (approximately 21%), as there

---

[2] Measures of mean fixation duration and saccade amplitude were consistent with the use of a positive cue to guide attention. In the *cue-target* condition, fixation durations were reliably shorter (M = 148 ms) and saccades reliably longer (M = 4.7 degrees amplitude) than in the *cue-all* condition (M = 183 ms latency; M = 4.1 degrees amplitude) [$t(11) = 15.38$, $p < .001$, $\eta_p^2 = .96$ for fixation duration; $t(11) = 5.92$, $p < .001$, $\eta_p^2 = .76$ for saccade amplitude]. However, in the *cue-avoid* condition, neither fixation duration (M = 187 ms) nor saccade amplitude (M = 4.1) differed significantly from the *cue-all* condition [$t(11) = 1.62$, $p = .13$, $\eta_p^2 = .19$ for fixation duration; $t(11) = 0.49$, $p = .64$, $\eta_p^2 = .02$ for saccade amplitude], again providing no clear evidence that participants were able to implement an exclusionary template.

were 3 cue-matching objects and 14 total objects remaining that had not yet been fixated.[3] By calculating the ratio of observed probability to chance probability, we were able to control for the history of the types of objects fixated within any particular trial and obtain a direct measure of guidance toward cue-matching objects in the *cue-target* condition and possible avoidance of cue-matching objects in the *cue-avoid* condition.

The odds ratio data were then log-transformed so that this measure would be on a linear scale and chance performance would be represented by a value of zero. To avoid undefined values when the probability of fixating a cue-matching object was zero in a bin, 1/32 (one half of the smallest unit of performance increment) was added to each observed probability and to each chance probability prior to log transformation, similar to methods used in the signal detection theory literature (Hautus, 1995). In the final log odds ratio measure (i.e., log of observed probability divided bv chance probability), values greater than zero indicate that a cue-matching object was fixated *more* frequently than predicted by chance, and values less than zero indicate a cue-matching object was fixated *less* frequently than chance.

The goal of this analysis was to determine how the log odds ratio evolved between the first object fixated, the second object fixated, the third object fixated, etc. (the *ordinal object fixated*). However, the maximum number of fixated objects varied across trials and participants, because trials terminated at different points depending on how many items were searched before the target was found. For example, in the *cue-target* condition, participants generally found the target after fixating only 2-3 objects, whereas in the *cue-avoid* condition, participants often fixated 8 or more objects to find the target, though this varied considerably from trial to trial and

---

[3] Note that this method of calculating chance assumes that the probability of object refixation is very low. This was indeed the case. The probability that a given fixation was directed to any of the objects that had already been fixated by that point in the trial was only 4.82%.

across participants. To set consistent criteria across conditions, we therefore limited our analyses to the set of ordinal fixations that contained a reasonable number of trials. Specifically, a given ordinal object fixation was included in the analysis only if at least 11 of the 12 participants had at least 5 trials with that number of fixated objects (and that met the inclusion criteria described previously). Cell means for each condition were entered into one-way ANOVAs with ordinal object fixated as a factor with a minimum of 2 and a maximum of 8 levels, depending on the condition. Because these binned data were not independent and sphericity was likely to be violated, all *p* values reported for the statistical tests on this analysis reflect the Huynh-Feldt correction for heterogeneity of covariance.

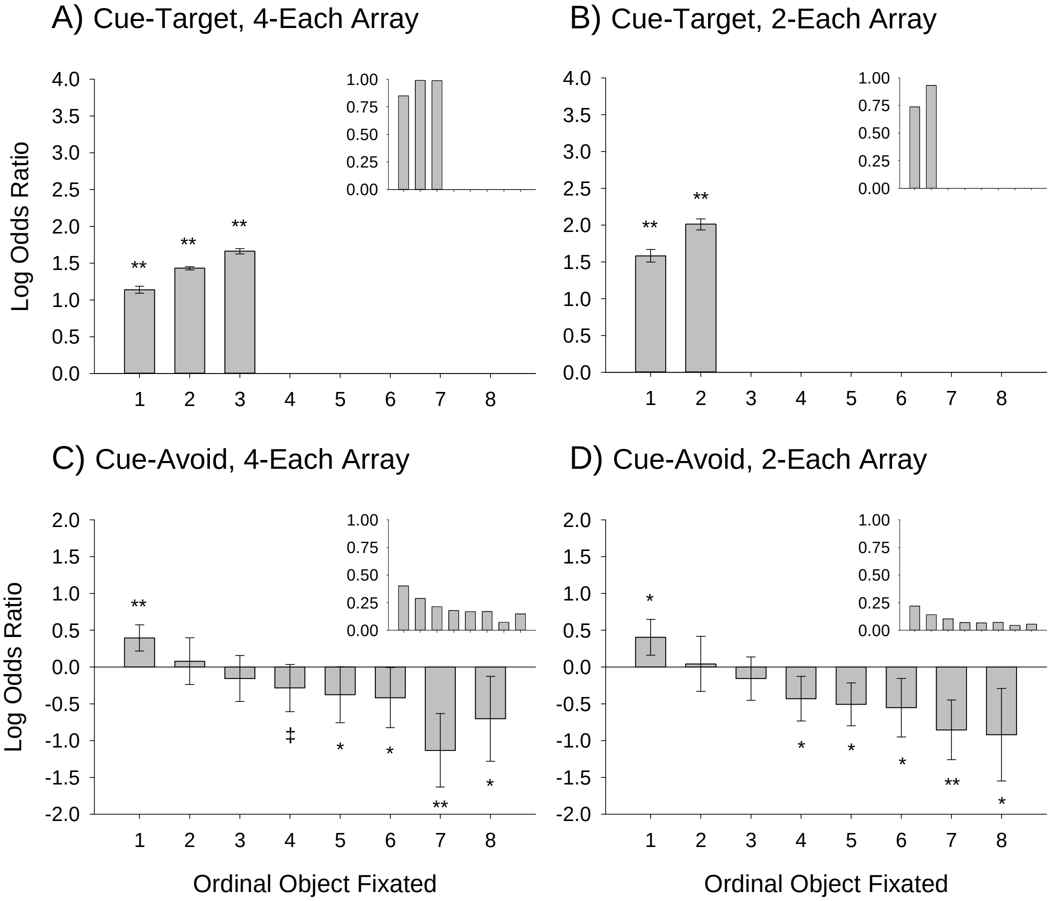

*Figure 3: Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the specific objects fixated thus far, plotted as a function of ordinal object fixated in a trial (first object fixated, second object fixated, etc.). Positive values indicate greater than chance probability of fixating a cue-matching object whereas negative values indicate less than chance probability. Inset plots show the raw observed probability of fixating a cue-matching object (i.e., disregarding which objects were previously fixated) as a function of ordinal object fixated. Data plotted are from the 4-each (A) and 2-each (B) arrays for the cue-target condition and from the 4-each (C) and 2-each (D) arrays for the cue-avoid condition in Experiment 1 (color cue stimulus). Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: ‡ indicates marginal significance ($p < .08$), * indicates $p \leq .05$, and ** indicates $p \leq .001$.*

In the *cue-target* condition, there was a significant main effect of ordinal object fixated for both the *4-each* arrays [$F(1.914, 21.058) = 655.95, p < .001, \eta_p^2 = .98$; Figure 3A] and the *2-each* arrays [$F(1, 11) = 263.01, p < .001, \eta_p^2 = .96$; Figure 3B]. The probability of fixating a cue-matching object increased over the course of a trial, although participants successfully implemented the template even from the first object fixated. Note that a relatively small number of objects was fixated in a trial when the target color was cued (on average, the target was the 2.34th object fixated in the *4-each* and 1.74th object fixated in the 2-each condition). This suggests that participants limited their search to the cued-color items and thus found the target after having fixated only a few objects. Consistent with this pattern, follow-up one-sample *t*-tests indicated that the log odds ratio was significantly greater than zero (i.e., that participants fixated cued-colored objects significantly more often than predicted by chance) for each ordinal object fixated in both the *4-each* arrays (Figure 3A) and the *2-each* arrays (Figure 3B). The probability of fixating a cue-matching object increased over the first few saccades on the array. This reflects the general finding that very early saccades on a scene or array (particularly the first) are not as strongly guided as later saccades, perhaps due to the influence of global scene properties on selection of the first saccade target location (Zelinsky, 2008).

In the *cue-avoid* condition, there was a significant main effect of ordinal object fixated for both *4-each* arrays [$F(5.742, 63.165) = 9.65$, $p < .001$, $\eta_p^2 = .47$; Figure 3C] and *2-each* arrays [$F(7, 70) = 7.28$, $p < .001$, $\eta_p^2 = .42$; Figure 3D]. Since the central question was whether the probability of fixating cue-matching items in the *cue-avoid* condition decreased systematically over the course of a trial, the key analysis concerned the linear trend across the 8 levels of ordinal object fixated. There was a reliable linear trend for both the *4-each* arrays [$F(1, 11) = 44.2$, $p < .001$, $\eta_p^2 = .80$ (90% CI: .52, .87)] and *2-each* arrays [$F(1, 10) = 27.3$, $p < .001$, $\eta_p^2 = .73$ (90% CI: .37, .83)], indicating a robust reduction in the probability of fixating cue-matching items as search progressed (Figures 3C and 3D). Moreover, gaze tended to be biased *toward* cue-matching items at the beginning of the trial and *away from* cue-matching items toward the end of the trial. One-sample *t*-tests revealed that the first object fixated was more likely to be a cued-color object than predicted by chance, that the second and third objects fixated were not significantly biased toward or away from the cued color, and that the remaining objects fixated (up to the eighth object) were significantly less likely than chance to be the cued color; this pattern was found for both the *4-each* (Figure 3C) and *2-each* (Figure 3D) arrays. Thus, there were two notable effects that emerged on *cue-avoid* trials: 1) initial capture of attention by cue-matching objects, and 2) subsequent avoidance of cue-matching objects.

Was early capture functionally related to later avoidance? In their "search and destroy" characterization of feature-guided avoidance, Moher and Egeth (2012) suggested that initial capture by a known irrelevant feature enables later avoidance of similar items. To test this, we sought to examine whether the reduced probability of fixating a cue-matching object late in a trial was contingent on early capture, dividing the trials by whether capture did or did not occur at the beginning of the trial. Although the ability to conduct this analysis was limited in Experiment 1 by small numbers of trials in each of the cells after division, we discuss the method here and provide preliminary results. A more comprehensive test is reported in Experiment 3. Early capture trials were defined as trials on which the to-be-avoided color was fixated as either the first or second object fixated. This analysis was necessarily limited to trials for which three or more objects were fixated (*4-each*: 91% trials retained), and to ensure that there were at least several cue-matching objects left to fixate, the analysis was limited to the *4-each* condition. As illustrated in Figure 4, participants demonstrated avoidance of cue-matching objects both on trials for which they showed initial oculomotor capture (Figure 4A) and on trials for which they did not show initial oculomotor capture (Figure 4B). One-sample *t*-tests revealed reliable avoidance of cue-matching objects by the fifth and subsequent objects on capture trials and reliable avoidance of cue-matching objects by the third and fourth objects fixated on trials without initial capture.[4] These results suggest that fixation of a cue-matching object early in the

---

[4] When the first or second saccade (or both) were not directed toward the cued to-be-avoided color, the target was naturally found after fewer fixations, so valid data were present only up through four objects fixated on these trials (Figure 4B).

trial is not necessary to demonstrate avoidance of cue-matching objects later in the trial, although thorough analysis is precluded by the many bins in the "no capture" trials (Figure 4B) that could not be analyzed due to limitations in the number of trials available.

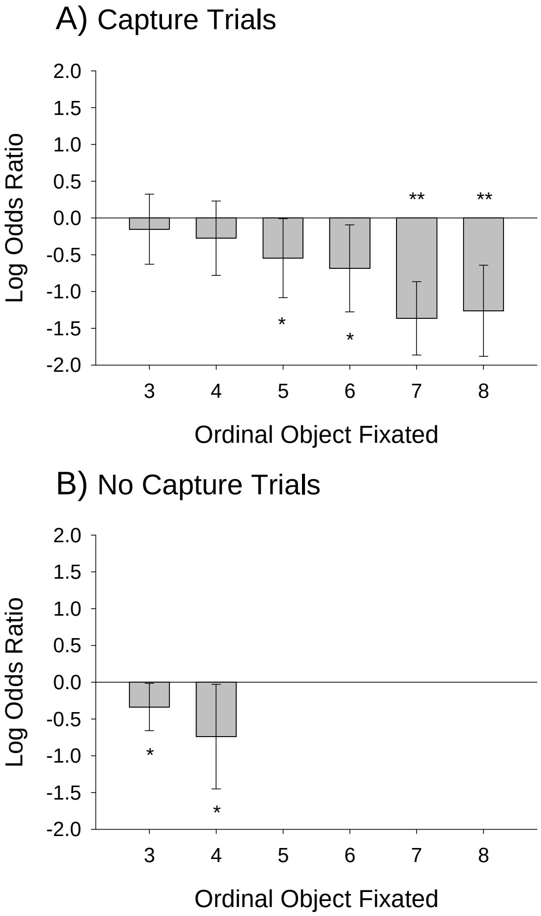

*Figure 4: Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Data plotted are from the 4-each array in the cue-avoid condition split into trials with initial capture (A) and without initial capture (B) from Experiment 1 (color cue stimulus). Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates p ≤ .05, and ** indicates p ≤ .001.*

Search Initiation Time. In the *cue-avoid* condition, participants showed a pattern of early capture and later avoidance. The potential benefit of later avoidance may have been offset by the cost of early capture, yielding overall search times that were similar to those found in the *cue-all* condition (Figure 2). However, cue condition also may have influenced overall search times by generating differences in the time taken to *initiate* search. For example, if the use of a negative cue requires a translation process (from a feature representation to a spatial template), then we would expect guided search to be initiated more slowly in that condition that in the *cue-all* or *cue-target* conditions. Search initiation time was operationalized as the elapsed time from the onset of the array to the first fixation on one of the array objects, collapsing across the *2-each* and *4-each* conditions. Where appropriate, we conditionalized this analysis on whether the first fixated object matched or did not match the cue.

For the *cue-all* condition, mean search initiation time was 335 ms. For the *cue-avoid* condition, this measure was 322 ms when gaze was first directed to a matching item (i.e., capture) and 390 ms when gaze was first directed to a non-matching item (i.e., to one of the possible target items). This indicates relatively rapid capture by a cue-matching item and relatively slow initiation of guided search to non-matching items. The 390-ms initiation latency for guided search in the *cue-avoid* condition was reliably longer than that for the *cue-all* condition (335 ms), $t(11) = 2.86$, $p = .015$, $\eta_p^2 = .43$. Thus, the benefit of avoiding cue-matching items later in the trial was not only offset by the cost of early oculomotor capture, it was also offset by a cost on search initiation time for trials when gaze was *not* captured initially. Finally, in the *cue-target* condition, mean initiation time was 345 ms when gaze was first directed to a cue-matching item (i.e., to one of the possible target items) and 219 ms when directed to a non-matching item. Note that these latter trials constituted only a small proportion of *cue-target* trials and were plausibly generated by rapid, global processing of the array (Zelinsky, 2008). Initiation of guided search for positive cues (345 ms) was marginally faster than initiation of guided searched for negative cues (390 ms), $t(11) = 1.91$, $p = .083$, $\eta_p^2 = .25$. Finally, there was no difference in guided search initiation time between the cue-target (345 ms) and cue-all conditions, $t(11) = 0.50$, $p = .629$, $\eta_p^2 = .022$. Although we cannot determine with confidence the precise source of the delay in search initiation for the *cue-avoid* condition, it could plausibly reflect the time necessary to convert a feature representation to a spatial template, or it could reflect covert capture of attention and suppression of a saccade to a negative-cue-matching object on some proportion of trials.

Summary. When the cue indicated the target color, participants used this information to efficiently restrict search to relevant items. This effect was observed on overall search times and on the probability that each fixated object matched the cued color. When the cue indicated a color to be avoided, measures of overall search time indicated no advantage relative to a neutral cue, suggesting that participants may not have been able to successfully implement a negative template or did not attempt to implement a negative template. However, inspection of the evolution of selection across the trial showed a systematic effect of negative cue use that was obscured in the overall measures of search time: early in the trial, attention was captured by the to-be-avoided color; later in the trial, participants successfully avoided that color. Preliminary evidence suggested that later avoidance was not necessarily contingent on early oculomotor capture. Finally, the use of a negative cue introduced a delay in the time taken to initiate search.

Experiment 2

One possible explanation for the early capture effect in the *cue-avoid* condition of Experiment 1 is that it resulted from low-level priming (Maljkovic & Nakayama, 1994) generated by sensory processing of the cue color patch. To test this possibility, we replicated Experiment 1, replacing the colored cue stimulus with the color name printed in dark grey. If the initial capture effect in the *cue-avoid* condition was driven primarily by perceptual priming, it should be eliminated with this modification. If, however, capture was caused by activation of the cued color value as a template in VWM, the capture effect should be preserved, although its magnitude might be reduced (e.g., Soto & Humphreys, 2007).

Method

Participants. The key linear trends in Experiment 1 produced effect sizes of η_p^2 = .80 and .73, consistent with our expectations from previous studies using object fixation to probe attentional guidance and selection during search. Thus, we continued to use a sample size of 12 participants in Experiments 2 and 3, which was more than sufficient to ensure appropriate levels of power for effects of this magnitude. Twelve participants (9 female; 18-30 years old) from the University of Iowa completed Experiment 2 and were compensated for their time. All participants reported normal color vision and normal or corrected-to-normal visual acuity. All procedures were approved by the University of Iowa's Institutional Review Board.

Stimuli and Procedure. Stimuli and procedure were the same as for Experiment 1 except that the colored cue square was replaced with the cue color name ("red", "yellow", "green", "blue", "white", "black", "purple", and "aqua") printed in dark grey at fixation. In the *cue-all* condition, the checkerboard cue square was replaced with the word "any" printed in dark grey.

Data Analysis. As in Experiment 1, manual response accuracy was uniformly high (*M* = 99%) across all conditions (see Table 1 for accuracy by condition and array type), and trials with incorrect responses were excluded from all further analyses. All trials with response times that were less than 150 ms or greater than 10,000 ms were excluded from analysis (8.24% of trials). Furthermore, trials with response times that were beyond 2.5 standard deviations from the mean of each condition for each subject were also excluded from all analyses (additional 1.57% of trials).

Results and Discussion

Overall Search Time. The results are summarized in Figure 5. For the positive template, a condition X array type ANOVA run on mean TTF revealed a main effect of condition [$F(1, 11) = 735.03$, $p < .001$, $\eta_p^2 = .99$], a main effect of array type [$F(1, 11) = 7.15$, $p = .022$, $\eta_p^2 = .39$], and a significant interaction [$F(1, 11) = 6.39$, $p = .028$, $\eta_p^2 = .37$; see Figure 5]. Collapsing across array type, we found faster TTF in the *cue-target* condition ($M = 612$ ms) than in the *cue-all* condition ($M = 1996$ ms), indicating attentional guidance by a positive template. Similarly, collapsing across cue condition, we found faster TTF for *2-each* ($M = 1241$ ms) than *4-each* arrays ($M = 1367$ ms), primarily driven by the fewer number of relevant items in the *2-each* (2) compared to *4-each* (4) arrays in the *cue-target* condition, which also explains the significant interaction.

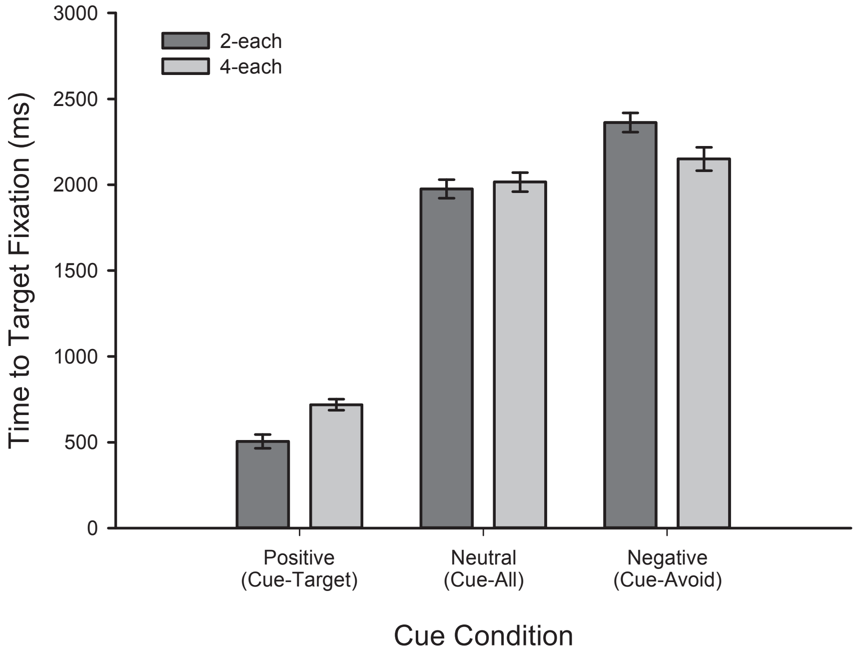

Figure 5: Elapsed time to first fixation on the target item plotted as a function of cue condition (cue-target: Positive, cue-all: Neutral, cue-avoid: Negative) and array type (4-each, 2-each) for Experiment 2 (word cue stimulus). Error bars indicate within-subjects 95% confidence intervals (Morey, 2008).

For the negative template, a condition X array type ANOVA run on mean TTF revealed a significant main effect of condition [$F(1, 11) = 19.39$, $p = .001$, $\eta_p^2 = .64$], but no main effect of array type or significant interaction (all ps > .1). Collapsing across array type, we found slower TTF in the *cue-avoid* condition ($M = 2257$ ms) than in the *cue-all* condition ($M = 1996$ ms), which is the opposite pattern from what we would expect if participants were able to benefit from the negative cue information.[5]

Object-by-object Analysis of Selectivity. The evolution of selection across the trial corresponded closely to the pattern observed in Experiment 1 (Figure 6). In the *cue-target* condition, there was a significant main effect of ordinal object fixated on the probability of fixating a matching object (log-odds ratio) for both the *4-each* [$F(2, 22) = 451.61$, $p < .001$, $\eta_p^2 = .98$; Figure 6A] and *2-each* [$F(1, 11) = 154.86$, $p < .001$, $\eta_p^2 = .93$; Figure 6B] arrays. The small number of objects fixated in a trial (on average the target was the 2.45th object fixated in the *4-each* and the 1.78th object fixated in the *2-each* condition) indicated that participants limited selection to relevant, cued-color items. Follow-up one-sample t-tests revealed that participants fixated cue-matching objects significantly more often than predicted by chance at each of the first three objects in the *4-each* array (Figure 6A) and at each of the first two objects in the *2-each* array (Figure 6B).

---

[5] As in Experiment 1, measures of mean fixation duration and saccade amplitude were consistent with the patterns observed in end-of-trial measures of search efficiency. In the *cue-target* condition, fixation durations were reliably shorter ($M = 136$ ms) and saccades were reliably longer ($M = 4.6$ degrees amplitude) than in the *cue-all* condition ($M = 175$ ms; $M = 3.9$ degrees amplitude) [$t(11) = 14.58$, $p < .001$, $\eta_p^2 = .95$ for fixation duration; $t(11) = 8.14$, $p < .001$, $\eta_p^2 = .86$ for saccade amplitude]. However, in the *cue-avoid* condition, neither fixation duration ($M = 176$ ms) nor saccade amplitude ($M = 3.8$ degrees amplitude) differed significantly from the *cue-all* condition, [$t(11) = 0.65$, $p = .53$, $\eta_p^2 = .04$ for fixation duration; $t(11) = 0.76$, $p = .46$, $\eta_p^2 = .05$ for saccade amplitude].

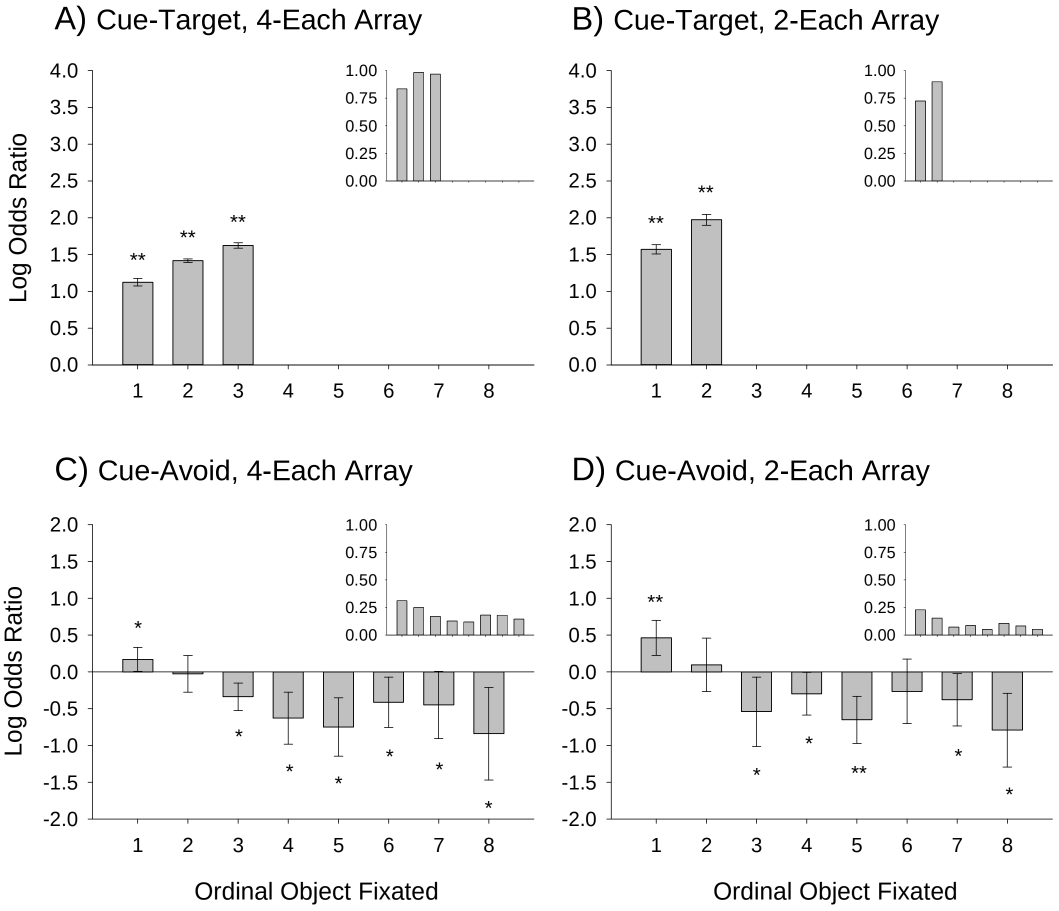

*Figure 6: Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Inset plots show the raw observed probability of fixating a cue-matching object as a function of ordinal object fixated. Data plotted are from the 4-each (A) and 2-each (B) arrays for the cue-target condition and from the 4-each (C) and 2-each (D) arrays for the cue-avoid condition in Experiment 2 (word cue stimulus). Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates p ≤ .05, and ** indicates p ≤ .001.*

In the *cue-avoid* condition, we again observed a significant main effect of ordinal object fixated for both *4-each* [$F(4.113, 41.132) = 4.59$, $p = .003$, $\eta_p^2 = .32$; Figure 6C] and *2-each* [$F(7, 77) = 6.67$, $p < .001$, $\eta_p^2 = .38$; Figure 6D] arrays. Critically, there was a reliable linear trend for both the *4-each* arrays [$F(1, 10) = 13.5$, $p = .004$, $\eta_p^2 = .57$ (90% CI: .16, .73)]; and *2-each* arrays [$F(1, 11) = 13.8$, $p = .003$, $\eta_p^2 = .56$ (90% CI: .16, .72)], indicating systematic reduction in the probability of fixating cue-matching items across the course of search. One-sample *t*-tests revealed that the first object fixated was more likely to be a cued-color object than predicted by chance, that the second object fixated was not significantly biased toward or away from the cued color, and that the remaining objects fixated (up to the eighth object) were significantly less likely than chance to be the cued color (except where indicated in Figure 6). This pattern held for both the *4-each* (Figure 6C) and *2-each* (Figure 6D) arrays.

To probe whether later avoidance was contingent on early capture, we again divided the trials by whether capture did or did not occur at the beginning of the trial and limited the analysis to trials for which three or more objects were fixated (*4-each*: 94% retained), as described in Experiment 1. Participants fixated cue-matching objects significantly less often than predicted by chance both when early capture occurred (Figure 7A) and when it did not (Figure 7B). One-sample *t*-tests revealed reliable avoidance of cue-matching objects by the fourth and fifth objects on capture trials and by the third and subsequent objects on trials without capture. As in Experiment 1, these results suggest that fixation of a cue-matching object early in the trial is not necessary to produce avoidance of cue-matching objects later in the trial, although the analysis must again be considered preliminary given the small number of observations available.

A) Capture Trials

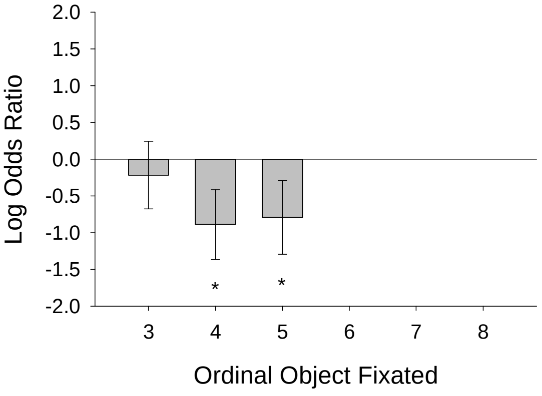

B) No Capture Trials

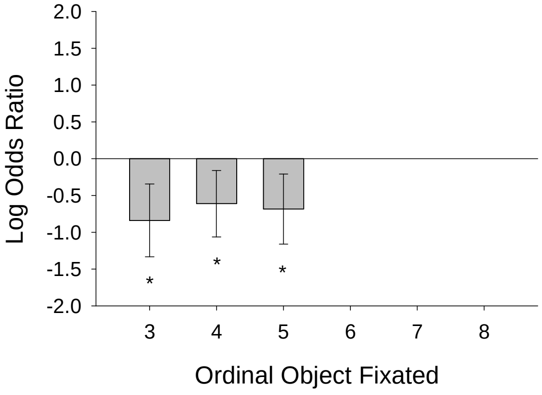

*Figure 7: Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Data plotted are from the 4-each array in the cue-avoid condition split into trials with initial capture (A) and without initial capture (B) from Experiment 2 (word cue stimulus). Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates p ≤ .05, and ** indicates p ≤ .001.*

Search Initiation Time. For the *cue-all* condition, mean search initiation time was 305 ms. For the *cue-avoid* condition, this measure was 332 ms when gaze was first directed to a matching item (i.e., capture) and 359 ms when gaze was first directed to a non-matching item (i.e., to one of the possible target items). As in Experiment 1, the 359-ms initiation latency for guided search in the *cue-avoid* condition was reliably longer than that for the *cue-all* condition (305 ms), $t(11) = 3.76$, $p = .003$, $\eta_p^2 = .56$. In the *cue-target* condition, mean initiation time was 339 ms when gaze was first directed to a cue-matching item (i.e., to one of the possible target items) and 224 ms when directed to a non-matching item. Initiation of guided search for positive cues (339 ms) was not significantly faster than initiation of guided searched for negative cues (359 ms), $t(11)$ = 0.88, p = .399, η_p^2 = .07, although the numerical difference was in the same direction as in Experiment 1. Finally, there was a marginal difference in guided search initiation time between the cue-target (339 ms) and cue-all conditions, $t(11)$ = 2.00, p = .071, η_p^2 = .267, potentially indicating that additional time was also necessary to implement a positive template.

Summary. Experiment 2 replicated most of the principal results observed in Experiment 1. End-of-trial measures of search efficiency again obscured a more complicated pattern of selection across the trial. Unlike Experiment 1, there was an overall *cost* associated with the negative-cue condition relative to the neutral condition. Yet, object-by-object selection indicated the same pattern as in Experiment 1—early capture and later avoidance—highlighting the need to assess selection across the course of the trial. Moreover, the initiation of search was again delayed in the *cue-avoid* condition relative to the *cue-all* condition. Finally, the early capture effect in the *cue-avoid* condition was observed using a text label rather than a color square, demonstrating that the effect was unlikely to be caused by low-level priming.

Experiment 3

A possible explanation for the delayed implementation of avoidance in Experiments 1 and 2 is that the delay between the cue and the search array was simply too short for participants to configure an exclusionary template. Previous work has demonstrated that, after as little as 200 ms, participants are able to efficiently use cue information to guide search toward matching items (Vickery et al., 2005; Wolfe et al., 2004). It is possible, however, that more time is required to configure an exclusionary template. In Experiments 1 and 2, the cue-stimulus delay was 500 ms, and participants consistently demonstrated avoidance of cue-matching objects by the third (Experiment 2) or fourth (Experiment 1) object fixated. Experiment 3 systematically extended the cue-stimulus delay past the point at which we observed avoidance in the previous experiments (Experiment 1: 1929 ms; Experiment 2: 1528 ms) to test whether avoidance could be observed at the beginning of the trial. Specifically, the cue-stimulus delay was increased from the original 500 ms to a maximum of 2000 ms. Additionally, because participants could have occasionally fixated both of the cued-color items in the *2-each* array and then not have any unvisited cued-color objects left to avoid, we only used *4-each* arrays.

Method

Participants. Twelve new participants (5 female; 18-30 years old) were recruited from the University of Iowa and were compensated for their time. All participants reported normal color vision and normal or corrected-to-normal visual acuity. All procedures were approved by the University of Iowa's Institutional Review Board.

Stimuli and Procedure. Participants in Experiments 1 and 2 demonstrated avoidance of cue-matching items by the third or fourth object fixated in a trial (Exp 1: approximately 1900 ms after cue onset; Exp 2: approximately 1500 ms after cue onset). Therefore, the delay between the cue and the search array was increased to a maximum of 2000 ms to allow sufficient time for participants to establish an exclusionary search template. The stimulus onset asynchrony (SOA) between the cue stimulus and the search array was 500 (same SOA used in Experiments 1 and 2), 1000, 1500, or 2000 ms. The SOA interval was randomly intermixed within each cue condition. Lastly, we eliminated the *cue-all* condition to focus on the *cue-target* and *cue-avoid* conditions. Again, cue condition was blocked and condition order was counterbalanced across participants. The session began with a 12-trial practice block (6 trials each for *cue-target* and *cue-avoid*). Then there were eight blocks of 24 trials for each of the two cue conditions. The first two trials in each block were considered warm-up trials and were excluded from all analyses. This yielded 44 trials per SOA, per condition.

Data Analysis. Manual response accuracy was uniformly high (*M* = 99% correct) across all conditions (see Table 2 for accuracy by condition and SOA), and trials with incorrect responses were excluded from all further analyses. All trials with response times that were less than 150 ms or greater than 10,000 ms were excluded from analysis (8.36% of trials). Furthermore, trials with response times that were beyond 2.5 standard deviations from the mean of each condition for each subject were also excluded from all analyses (additional 1.96% of trials).

SOA	Cue Condition	
	Positive (Cue-Target)	Negative (Cue-Avoid)
500	99.6% 1298	98.7% 2866
1000	98.4% 1237	98.7% 3002
1500	99.1% 1248	98.4% 2875
2000	99.0% 1242	98.8% 2976

Table 2:Manual response accuracy and manual response times (ms) by condition and SOA for Experiment 3.

Results and Discussion

Overall Search Time. A condition (*cue-target*, *cue-avoid*) X SOA (500, 1000, 1500, 2000) ANOVA run on mean TTF revealed a main effect of condition [$F(1, 11) = 153.88$, $p < 0.001$, $\eta_p^2 = .93$], but no main effect of SOA and no significant interaction (all *p*s > .30; see Figure 8). Unsurprisingly, participants were able to locate the target item more quickly in the *cue-target* ($M = 777$ ms) than in the *cue-avoid* ($M = 2265$ ms) condition, reflecting attentional guidance by a positive template (Beck et al., 2012). However, there was no effect of SOA, even in the *cue-avoid* condition, suggesting that participants did not benefit from the additional time to prepare an exclusionary template, at least as reflected in overall search time.

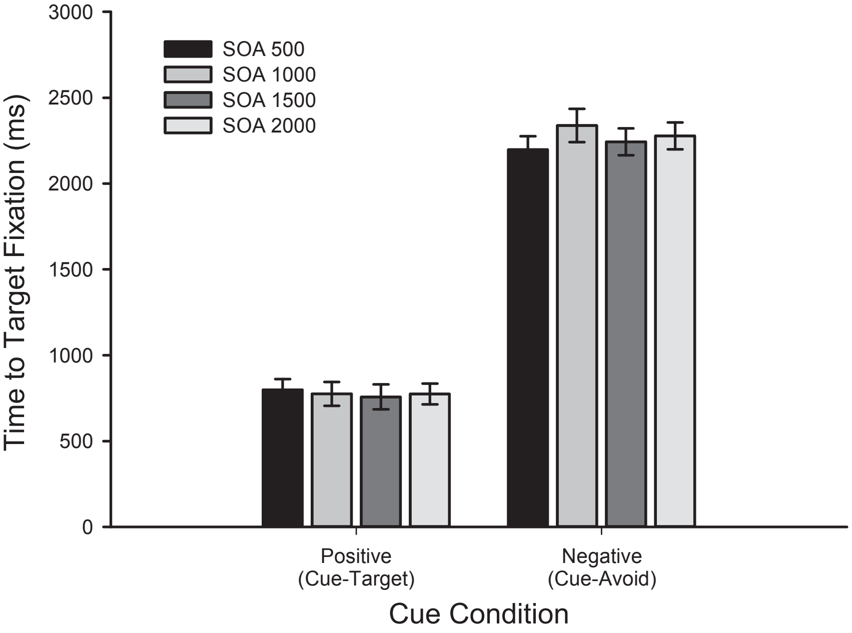

Figure 8: Elapsed time to first fixation on the target item plotted as a function of cue condition (cue-target: Positive, cue-avoid: Negative) and SOA (500, 1000, 1500, 2000) for Experiment 3. Error bars indicate within-subjects 95% confidence intervals (Morey, 2008).

Object-by-object Analysis of Selectivity. As in Experiments 1 and 2, the eye movement data were binned by ordinal object fixated during search and log-transformed odds ratios were calculated to measure the probability of fixating a cue-matching object for each bin. In the *cue-target* condition, there was a significant main effect of ordinal object fixated for SOA 500 [$F(1.842, 18.418) = 355.00$, $p < .001$, $\eta_p^2 = .97$; Figure 9A], SOA 1000 [$F(2, 22) = 161.31$, $p < .001$, $\eta_p^2 = .94$; Figure 9B], SOA 1500 [$F(1.662, 18.283) = 197.35$, $p < .001$, $\eta_p^2 = .95$; Figure 9C], and SOA 2000 [$F(1.820, 20.015) = 511.366$, $p < .001$, $\eta_p^2 = .98$; Figure 9D]. Follow-up one-sample *t*-tests examining whether each bin differed from zero revealed that participants fixated cue-matching objects significantly more often than predicted by chance in each cell that was included in the analysis (Figure 9A-D). Again, these data indicate participants were able to quickly restrict selection to relevant, cued-color items.

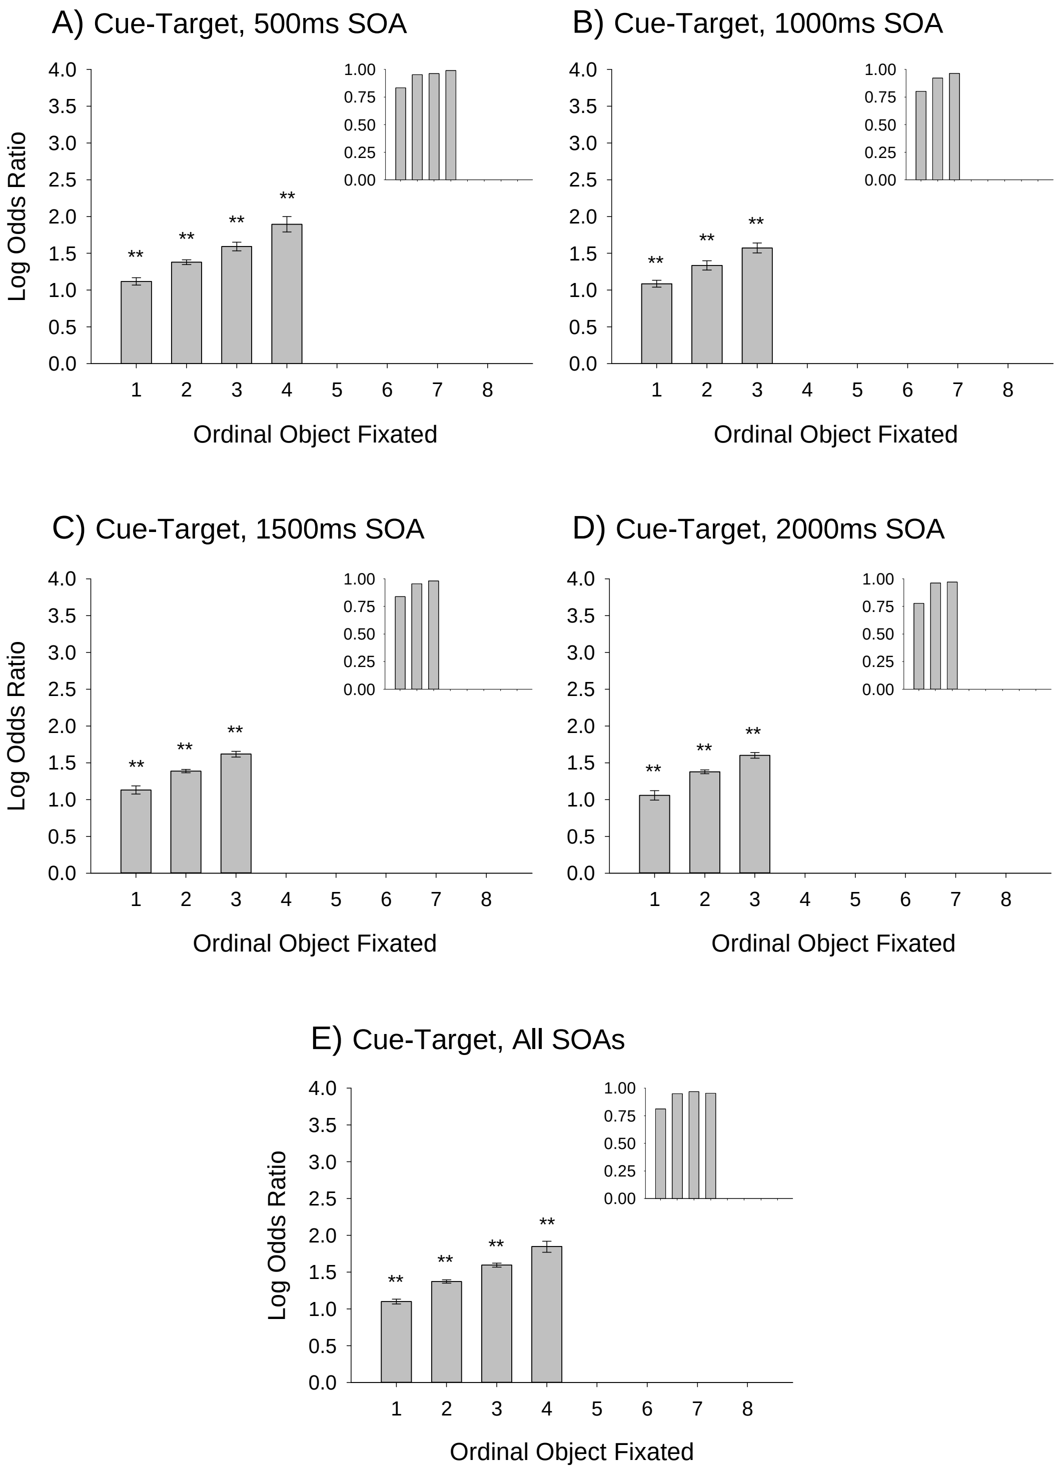

Figure 9: Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Inset plots show the raw observed probability of fixating a cue-matching object as a function of ordinal object fixated. Data plotted are from the SOA 500 (A), SOA 1000 (B), SOA 1500 (C), SOA 2000 (D), and all SOAs (E) for the cue-target *condition in Experiment 3. Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates* p *≤ .05, and ** indicates* p *≤ .001.*

In the *cue-avoid* condition, there was a significant main effect of ordinal object fixated for SOA 500 [$F(5.833, 58.330) = 6.152, p < .001, \eta_p^2 = .38$; Figure 10A], SOA 1000 [$F(5.106, 56.168) = 8.48, p < .001, \eta_p^2 = .44$; Figure 10B], SOA 1500 [$F(6.644, 73.085) = 6.58, p < .001, \eta_p^2 = .37$; Figure 10C], and SOA 2000 [$F(5.631, 61.943) = 6.70, p < .001, \eta_p^2 = .38$; Figure 10D]. Critically, there was a reliable linear trend in all four SOA conditions: SOA 500 [$F(1, 10) = 15.8, p = .003, \eta_p^2 = .61$ (90% CI: .20, .76)], SOA 1000 [$F(1, 11) = 74.8, p < .001, \eta_p^2 = .87$ (90% CI: .67, .92)], SOA 1500 [$F(1, 11) = 20.7, p < .001, \eta_p^2 = .65$ (90% CI: .27, .78)], and SOA 2000 [$F(1, 11) = 16.0, p = .002, \eta_p^2 = .59$ (90% CI: .20, .74)]. Thus, there was systematic reduction in the probability of fixating cue-matching items across the course of search, and this was observed at all cue-search SOAs.

Follow-up one-sample *t*-tests revealed that participants fixated cue-matching objects more frequently than predicted by chance for the first object (SOA 1000 was marginal, $p = .065$), and reliably less often than chance by the third object for all SOAs (by the second object for SOA 1000; see Figure 10A-D). Even at the longer SOAs—at a post-cue time by which participants were able to avoid cue-matching objects in Experiments 1 and 2—we again observed early capture during search. These results suggest that failure to find evidence of avoidance of cue-matching objects early in the trial in Experiments 1 and 2 was not because participants needed more time between appearance of the cue stimulus and the search array to configure an exclusionary template.

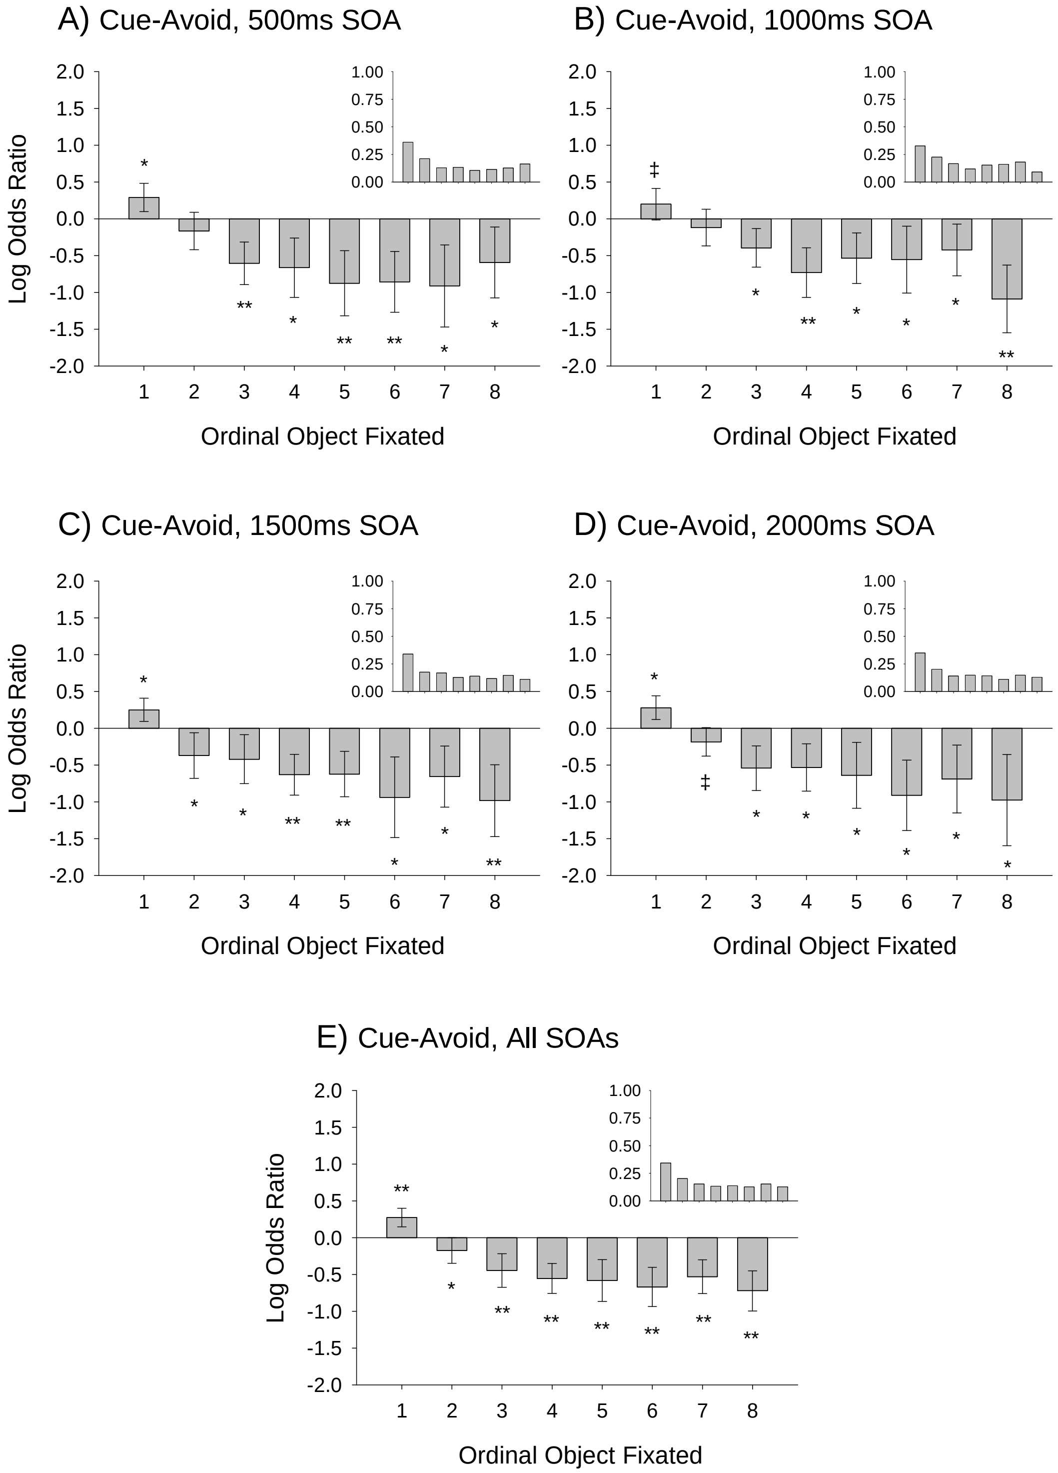

Figure 10: Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Inset plots show the raw observed probability of fixating a cue-matching object as a function of ordinal object fixated. Data plotted are from the SOA 500 (A), SOA 1000 (B), SOA 1500 (C), SOA 2000 (D), and all SOAs (E) for the cue-avoid *condition in Experiment 3. Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: ‡ indicates marginal significance (*p *≤ .065, * indicates* p *≤ .05, and ** indicates* p *≤ .001.*

To examine whether later avoidance was contingent on early capture, we split the ordinal object fixated data by the same capture criteria used previously (collapsed across SOA). As in Experiments 1 and 2, this analysis was limited to trials for which three or more objects were fixated (92% retained across all SOAs). Participants fixated cue-matching objects significantly less often than predicted by chance both when early capture occurred (Figure 11A) and when it did not (Figure 11B). One-sample *t*-tests comparing each bin against zero revealed reliable avoidance of cue-matching objects by the third object fixated for trials with and without capture, and this avoidance of cue-matching objects remained reliable through the eighth object fixated. As in Experiments 1 and 2, these results suggest that fixation of a cue-matching object early in the trial is not necessary to produce avoidance of cue-matching objects later in the trial. Furthermore, this avoidance was robust for all objects beyond the third object fixated in a trial.

Omnibus estimate of effect size. The key effect (reliable reduction in the probability of fixating a to-be-avoided color as a function of ordinal object fixated) was observed in all eight implementations in this study (the two array-type conditions in each of Experiments 1 and 2 and the four SOA conditions of Experiment 3). Effect sizes for the linear trend ranged from $\eta_p^2 = .56$ to .87. Thus, we have confidence that the decrease in the probability of fixating the to-be-avoided color over the course of the trial was a large and replicable effect. To obtain a more precise estimate of effect size, we combined the *cue-avoid* data from Experiments 1-3 in an omnibus analysis, collapsing across the within-experiment implementations and treating experiment as a between-subjects factor. There was a reliable linear trend as a function of ordinal object fixated, $F(1, 33) = 104.4$, $p < .001$, $\eta_p^2 = .76$ (90% CI: .62, .82). This effect size estimate and confidence range can be used to guide future work seeking to replicate and extend the present findings.

A) Capture Trials, All SOAs

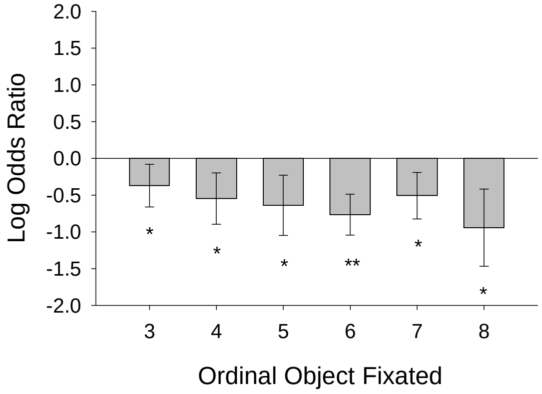

B) No Capture Trials, All SOAs

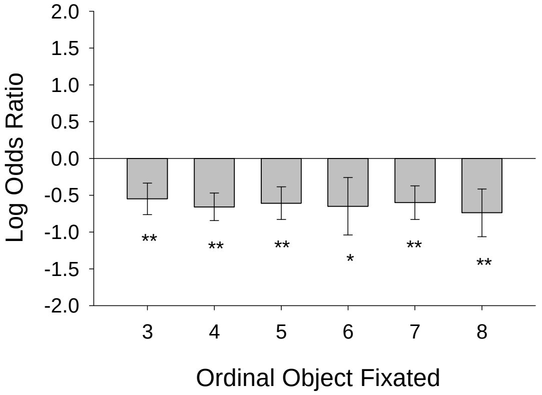

Figure 11: Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Data plotted are collapsed across SOA in the cue-avoid *condition split into trials with initial capture (A) and without initial capture (B) from Experiment 3. Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates* p *≤ .05, and ** indicates* p *≤ .001.*

Search Initiation Time. For the search initiation time analysis, we collapsed across the SOA conditions. In the *cue-avoid* condition, mean search initiation time was 359 ms when gaze was first directed to a matching item (i.e., capture) and 429 ms when gaze was first directed to a non-matching item (i.e., to one of the possible target items). In the *cue-target* condition, mean initiation time was 379 ms when gaze was first directed to a cue-matching item (i.e., to one of the possible target items) and 251 ms when directed to a non-matching item. Initiation of guided search for positive cues (379 ms) was significantly faster than initiation of guided searched for negative cues (429 ms), $t(11) = 2.99$, $p = .012$, $\eta_p^2 = .448$, again indicating a delay associated with search initiation in the cue-avoid conditon.

Summary. If participants simply needed more time between the cue stimulus and the search array to configure an exclusionary template, the initial capture effect observed in Experiments 1 and 2 should have been eliminated or should have diminished as the SOA increased. However, the initial capture effect was observed across SOAs and remained robust at the longest SOAs. Moreover, in the *cue-avoid* condition, we observed robust avoidance of cue-matching objects after initial capture.

General Discussion

In three experiments, we examined the implementation of positive templates (specifying the target color) and exclusionary templates (specifying the color of items that would not be the target) in a search paradigm optimized to observe the evolution of selection across the course of a trial. Targets contained a feature that could be discriminated only by foveation, and the sequence of eye fixations on individual objects during search provided data about selectivity across time. For positive templates, we replicated previous experiments demonstrating that participants efficiently restrict selection to template-matching items (Beck et al., 2012; Green & Anderson, 1956; Vickery et al., 2005; Wolfe et al., 2004), and this selectivity was implemented from the very first fixated object in the array. For exclusionary templates, there were four main findings. First, early in the trial, template-matching objects were fixated *more* often than would be expected by chance: gaze was preferentially oriented to template-matching items despite the demand to ignore them. Second, later in the trial, template-matching objects were fixated *less* often than would be expected by chance, indicating that participants ultimately configured guidance to avoid items cued as irrelevant. Third, later avoidance was not contingent on early oculomotor capture: avoidance was observed robustly both on trials with and without initial capture. Finally, the time taken to initiate search after array onset was longer for exclusionary templates than for positive templates or for a neutral control condition.

The primary question was whether feature-based attentional control settings can be configured to guide attention away from known irrelevant items. Woodman and Luck (2007) originally proposed that a VWM representation could be used flexibly either to guide attention toward or away from matching objects. In their study, participants saw a memory color that would never be the target value. Search RT was systematically reduced as the number of matching items in the display increased. In contrast, a large number of studies have demonstrated capture by memory matching objects, even when those objects are known never to be targets (Folk et al., 1992; Hollingworth & Luck, 2009; Hollingworth et al., 2013; Olivers, 2009; Olivers et al., 2006; Soto et al., 2005, 2008; Soto, Humphreys, & Heinke, 2006; Soto & Humphreys, 2007). These latter studies have found an increase in search RT when a memory-matching distractor is present in the display, suggesting that participants were not able to implement an exclusionary template.

This empirical inconsistency may have arisen, in part, from the fact that selectivity in search was assessed only by end-of-trial measures of search time. In the present experiments, the object-by-object eye movement data revealed both early capture by cue-matching distractors (consistent with studies observing RT costs of cue-matching items) and later avoidance of cue-matching items (consistent with Woodman & Luck, 2007). The studies observing RT costs tended to include only a single cue-matching item (e.g., Olivers et al., 2006; Soto et al., 2005). Thus, there would have been limited opportunity for later avoidance to offset the cost of early capture. However, in the Woodman and Luck study, a benefit of later avoidance may have overshadowed a cost of early capture, since there were a relatively large number of cue-matching objects in the display, increasing the potential benefit of avoidance (see also Kugler et al., 2015). Thus, our analysis of the microstructure of selection during search has the potential to resolve conflicting evidence in the literature: different patterns of end-of-trial search measures may be explained by relative differences in the magnitudes of early capture costs and later avoidance benefits. Moreover, the benefits of later avoidance appeared to be further countered by a delay in the time required to initiate search on trials without overt capture. It is therefore possible that previous studies may have failed to find negative cue benefits on end-of-trial search measures, at least in part, because participants took longer to initiate the search operation in the negative cue condition compared with a neutral condition.

Does evidence of later avoidance during search indicate that participants were ultimately able to configure a feature-based exclusionary template? This is probably the least likely of several possible explanations. The early capture effect clearly indicates that participants did not establish a feature-based exclusionary template before the trial began, despite seeing the relevant color in the cue and despite being given as long as 2000 ms to configure the template. Moreover, it is unlikely that the appearance of the array was necessary to configure a feature-based exclusionary template. The array provided no color information that was not available in the cue stimulus. If participants needed to attend to the relevant color in order to establish an exclusionary template, they should have been able to do so when the cue itself was presented at central fixation. As we have argued previously (Beck & Hollingworth, 2015), it may be extremely difficult, perhaps impossible, to configure an online, feature-based exclusionary template in a trial-by-trial manner. To know what feature to avoid when the array appears, a person must remember the cue value by maintaining it in VWM. This is likely to engage active, sustained representations in sensory cortex (Emrich, Riggall, LaRocque, & Postle, 2013; Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009), which will interact with new sensory processing to increase the salience of matching items (Hamker, 2004; Schneegans, Spencer, Schöner, Hwang, & Hollingworth, 2014) or to otherwise bias sensory competition in favor of those items (Desimone & Duncan, 1995). Attention will be attracted to the very items a person intends to avoid (Cunningham & Egeth, 2016; Han & Kim, 2009; Moher & Egeth, 2012; Tsal & Makovski, 2006). In more general terms, it may not be possible to actively remember the relevant feature value and simultaneously set a negative attentional weight for that value (relative to the weights for other values), because VWM maintenance overlaps, to some significant extent, with mechanisms for establishing positive attentional control.[6]

Note that although it may not be possible to configure a feature-based exclusionary template in VWM, several recent studies suggest that it is possible to configure an exclusionary template based on *long-term* learning across many trials of an experiment in which a particular feature is always associated with distractors (Gaspelin et al., 2015, 2017; Moher et al., 2014; Vatterott & Vecera, 2012). This is likely to reflect the gradual tuning of perceptual sensitivity to particular feature values. This same learning can be used to establish a positive feature template when the target feature remains consistent across trials (Carlisle et al., 2011). Thus, there appears to be a fundamental dissociation between feature-based attentional templates formed over long-term learning and those created online in VWM: the former supports templates both for selection and rejection; the latter supports only templates for selection.

---

[6] We do not claim that maintenance of a feature value in VWM is *equivalent* to an attentional template. The guidance effects observed from incidental maintenance of a feature value in VWM tend to be weaker than the guidance effects observed from maintenance of the target feature in VWM.

A second possible explanation for later avoidance in the present experiments is that participants translated the to-be-avoided cue color into a positive template specifying the values of the relevant array colors. For example, on a *4-each* trial in which the avoid cue was yellow and the array was composed of yellow, red, blue, and green items, participants may have established a positive template for red, blue, and green, allowing them to indirectly avoid attending to yellow. Note that because there were eight possible colors, a positive template for the three relevant colors in this example could be established only after the array appeared. Thus, early capture may have reflected the period of time when a still-active representation of the to-be-avoided color was being transformed into a positive template specifying the remaining colors in the array. Although possible, we think this alternative is unlikely given well-established limitations on search guidance from a positive feature template. There is currently debate over whether positive guidance is limited to a single value on a dimension (Olivers, Peters, Houtkamp, & Roelfsema, 2011; Wolfe, 2007) or can span multiple values on a dimension (Beck et al., 2012; Stroud, Menneer, Cave, & Donnelly, 2011). Although there now exists strong evidence supporting the latter view (Beck & Hollingworth, 2017; Beck et al., 2012; Grubert & Eimer, 2015, 2016; Hollingworth & Beck, 2016; Irons, Folk, & Remington, 2012; Moore & Weissman, 2010), guidance by even two feature values is significantly less efficient than guidance by a single feature value (Beck et al., 2012; Stroud et al., 2011). In the present study, we found avoidance in the *2-each* condition, in which there were eight colors in the array. It is unlikely that participants avoided the cued color in this condition by establishing a positive template specifying the remaining seven colors, especially since such a template representation would typically exceed the capacity of VWM. Given the very similar pattern of avoidance in the *2-each* and *4-each* conditions, it is therefore also unlikely that avoidance in the *4-each* condition was accomplished by conversion to a positive template specifying the three relevant colors.

A third possible explanation for later avoidance, and the one that seems most likely given the broader set of findings in the field, is that upon appearance of the search array, participants translated the feature-cue information into a *spatial* template specifying either the array locations to be avoided or the array locations to be attended (e.g., Kugler et al., 2015). Since the locations could be marked only after array onset, early capture would be caused by the fact that the active representation of the to-be-avoided color was maintained in VWM at the beginning of search, until it could be transformed. Spatial marking has been demonstrated to require initial attention to the to-be-marked locations (Humphreys, Stalmann, & Olivers, 2004), potentially consistent with the delayed implementation of guidance from a negative cue observed in the present study. Spatial translation is also consistent with our earlier finding that the utility of an exclusionary template is strongly dependent on the spatial configuration of array elements (Beck & Hollingworth, 2015). Finally, spatial translation provides a plausible explanation of exclusionary cue benefits in two additional studies that provided a preview of the search locations and colors, allowing participants to mark the locations before search commenced (Han & Kim, 2009; Moher & Egeth, 2012, Experiment 4). In the present experimental context, it is to be expected that a later benefit of spatial marking would not necessarily exceed the costs of early capture, since the colors were spatially intermixed (Beck & Hollingworth, 2015), the number of to-be-avoided colors never constituted more than 25% of the array, and there was an additional cost associated with search initiation for negative cue trials. However, our view predicts that manipulations influencing the ease of spatial recoding should directly influence the probability of later avoidance during the search trial.

To account for a similar pattern of early capture and later avoidance, Moher and Egeth

(2012) proposed a “search and destroy” mechanism as a means of implementing a negative attentional template. “Search and destroy” is an extension of Tsal and Makovski’s (2006) “process all” mechanism that was proposed to account for allocation of attention to any location where an item was expected to appear, regardless of task relevance. Specifically, Moher and Egeth (2012) proposed that attention was initially deployed to a cue-matching distractor, even though it was known to be task irrelevant, facilitating later avoidance of cue-matching items. However, in their critical Experiment 3, Moher and Egeth (2012) only included a single cue-matching distractor in the search array, so it is unclear whether the trend toward an avoidance effect reflected avoidance per se or was caused by the fact that avoidance could not be assessed independently of early attention to the cue-matching object. That is, having attended to the only cue-matching object early in the trial, “avoidance” may have simply reflected the allocation of attention to the remaining items in the array. In the present study, we included multiple cue-matching objects to provide a strong test of possible avoidance. Indeed, there was robust avoidance of cue-matching objects later in the trial.

Moreover, we tested whether there was a relationship between early capture and later avoidance, as claimed by the “search and destroy” account. We found no such relationship: robust avoidance occurred independently of early oculomotor capture. Of course, since we only assessed overt attention, it is possible that there was a relationship between early *covert* capture and later avoidance, but given the close association between covert and overt selection in this type of task, it is reasonable to expect that covert biases would be reflected, at least to some significant extent, in overt biases. The lack of a relationship here, combined with the fact that no functional relationship between early capture and later avoidance was demonstrated in Moher and Egeth (2012), suggests a simpler account based on spatial marking of to-be-searched or to-be-avoided locations. According to this account, the early bias toward cue-matching objects reflects the well-established phenomenon of VWM-based capture (e.g., Olivers et al., 2006; Soto et al., 2005), rather than a strategic process designed to facilitate later avoidance. The to-be-avoided color must be maintained in memory until the array appears and for some time during the process of template translation; during this period it supports capture. Later avoidance reflects the successful translation of the feature information into a spatial template for selection (Beck & Hollingworth, 2015).

Conclusion

In a search task requiring sequential oculomotor selection, a cue specifying a to-be-avoided feature generated a clear pattern of early capture and later avoidance. The relative magnitudes of the capture and avoidance effects has the potential to account for key empirical inconsistencies in the literature, in which end-of-trial search measures have sometimes shown exclusionary template costs and sometimes shown benefits. However, it is unlikely that the ultimate avoidance of negatively cued items was based on an exclusionary *feature-based* template. Instead, the most plausible explanation is that participants converted the feature information into a spatial template upon the appearance of the array. This could be considered as consistent with the general characterization of “search and destroy” developed by Moher and Egeth (2012). However, we found no evidence to suggest that later avoidance was contingent on early capture. Thus, the more conservative account is that capture occurs automatically as a consequence of maintaining the cued feature in VWM, and that the template conversion process is implemented in a manner that is largely independent of early capture.

Author Notes

This study was supported by grants from NIH (R01MH076226, R01MH065034, and R01EY017356) and by an NSF Graduate Research Fellowship to Valerie M. Beck. The authors would like to thank J. Toby Mordkoff for helpful discussions of the research.

Figure Captions

Figure 1. Example trial sequence and search arrays for Experiments 1-3. Participants were instructed to locate the Landolt-C with a top or bottom gap and report the gap location. The cue stimulus could indicate either the color of the target item (*cue-target*), the color to avoid (*cue-avoid*), or that the target item could be any color (*cue-all*; not shown). The search array could contain either four each of four different colors (*4-each*) or two each of eight different colors (*2-each*). Cue condition was blocked and the type of search array was intermixed.

Figure 2. Elapsed time to first fixation on the target item plotted as a function of cue condition (*cue-target*: Positive, *cue-all*: Neutral, *cue-avoid*: Negative) and array type (*4-each*, *2-each*) for Experiment 1 (color cue stimulus). Error bars indicate within-subjects 95% confidence intervals (Morey, 2008).

Figure 3. Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the specific objects fixated thus far, plotted as a function of ordinal object fixated in a trial (first object fixated, second object fixated, etc.). Positive values indicate greater than chance probability of fixating a cue-matching object whereas negative values indicate less than chance probability. Inset plots show the raw observed probability of fixating a cue-matching object (i.e., disregarding which objects were previously fixated) as a function of ordinal object fixated. Data plotted are from the *4-each* (A) and *2-each* (B) arrays for the *cue-target* condition and from the *4-each* (C) and *2-each* (D) arrays for the *cue-avoid* condition in Experiment 1 (color cue stimulus). Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: ‡ indicates marginal significance ($p < .08$), * indicates $p \leq .05$, and ** indicates $p \leq .001$.

Figure 4. Log-transformed odds ratios indicating the probability of fixating a cue-

matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Data plotted are from the *4-each* array in the *cue-avoid* condition split into trials with initial capture (A) and without initial capture (B) from Experiment 1 (color cue stimulus). Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates $p \leq .05$, and ** indicates $p \leq .001$.

Figure 5. Elapsed time to first fixation on the target item plotted as a function of cue condition (*cue-target*: Positive, *cue-all*: Neutral, *cue-avoid*: Negative) and array type (*4-each*, *2-each*) for Experiment 2 (word cue stimulus). Error bars indicate within-subjects 95% confidence intervals (Morey, 2008).

Figure 6. Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Inset plots show the raw observed probability of fixating a cue-matching object as a function of ordinal object fixated. Data plotted are from the *4-each* (A) and *2-each* (B) arrays for the *cue-target* condition and from the *4-each* (C) and *2-each* (D) arrays for the *cue-avoid* condition in Experiment 2 (word cue stimulus). Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates $p \leq .05$, and ** indicates $p \leq .001$.

Figure 7. Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values

indicate less than chance probability. Data plotted are from the *4-each* array in the *cue-avoid* condition split into trials with initial capture (A) and without initial capture (B) from Experiment 2 (word cue stimulus). Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates $p \leq .05$, and ** indicates $p \leq .001$.

Figure 8. Elapsed time to first fixation on the target item plotted as a function of cue condition (*cue-target*: Positive, *cue-avoid*: Negative) and SOA (500, 1000, 1500, 2000) for Experiment 3. Error bars indicate within-subjects 95% confidence intervals (Morey, 2008).

Figure 9. Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Inset plots show the raw observed probability of fixating a cue-matching object as a function of ordinal object fixated. Data plotted are from the SOA 500 (A), SOA 1000 (B), SOA 1500 (C), SOA 2000 (D), and all SOAs (E) for the *cue-target* condition in Experiment 3. Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates $p \leq .05$, and ** indicates $p \leq .001$.

Figure 10. Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Inset plots show the raw observed probability of fixating a cue-matching object as a function of ordinal object fixated. Data plotted are from the SOA 500 (A), SOA 1000 (B), SOA 1500 (C), SOA 2000 (D), and all SOAs (E) for the *cue-avoid* condition

in Experiment 3. Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: ‡ indicates marginal significance (*p* ≤ .065, * indicates *p* ≤ .05, and ** indicates *p* ≤ .001.

Figure 11. Log-transformed odds ratios indicating the probability of fixating a cue-matching object given the types of objects fixated thus far plotted as a function of ordinal object fixated in a trial. Positive values indicate greater than chance probability whereas negative values indicate less than chance probability. Data plotted are collapsed across SOA in the *cue-avoid* condition split into trials with initial capture (A) and without initial capture (B) from Experiment 3. Error bars indicate standard 95% confidence intervals. Values in each bin were compared against zero with significance levels as follows: * indicates *p* ≤ .05, and ** indicates *p* ≤ .001.

References

Arita, J. T., Carlisle, N. B., & Woodman, G. F. (2012). Templates for rejection: Configuring attention to ignore task-irrelevant features. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(3), 580–4. https://doi.org/10.1037/a0027885

Beck, V. M., & Hollingworth, A. (2015). Evidence for negative feature guidance in visual search is explained by spatial recoding. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(5), 1190–1196. https://doi.org/10.1037/xhp0000109

Beck, V. M., & Hollingworth, A. (2017). Competition in saccade target selection reveals attentional guidance by simultaneously active working memory representations. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(2), 225–230. https://doi.org/10.1037/xhp0000306

Beck, V. M., Hollingworth, A., & Luck, S. J. (2012). Simultaneous Control of Attention by Multiple Working Memory Representations. *Psychological Science*, *23*(8), 887–898. https://doi.org/10.1177/0956797612439068

Becker, M. W., Hemsteger, S., & Peltier, C. (2016). No templates for rejection: A failure to configure attention to ignore task-irrelevant features. *Visual Cognition*. https://doi.org/10.1080/13506285.2016.1149532

Bundesen, C. (1990). A theory of visual attention. *Psychol Rev*, *97*(4), 523–547.

Carlisle, N. B., Arita, J. T., Pardo, D., & Woodman, G. F. (2011). Attentional templates in visual working memory. *The Journal of Neuroscience*, *31*(25), 9315–22. https://doi.org/10.1523/JNEUROSCI.1097-11.2011

Cunningham, C. a, & Egeth, H. E. (2016). Taming the White Bear: Initial Costs and Eventual Benefits of Distractor Inhibition. *Psychological Science*. https://doi.org/10.1177/0956797615626564

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222. https://doi.org/10.1146/annurev.ne.18.030195.001205

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychol Rev*, *96*(3), 433–458.

Emrich, S. M., Riggall, A. C., LaRocque, J. J., & Postle, B. R. (2013). Distributed Patterns of Activity in Sensory Cortex Reflect the Precision of Multiple Items Maintained in Visual Short-Term Memory. *Journal of Neuroscience*, *33*(15), 6516–6523. https://doi.org/10.1523/JNEUROSCI.5732-12.2013

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(4), 1030–44.

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2015). Direct Evidence for Active Suppression of Salient-but-Irrelevant Sensory Inputs. *Psychological Science*. https://doi.org/10.1177/0956797615597913

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2017). Suppression of overt attentional capture by salient-but-irrelevant color singletons. *Attention, Perception, & Psychophysics*, *79*(1), 45–62. https://doi.org/10.3758/s13414-016-1209-1

Green, B. F., & Anderson, L. K. (1956). Color coding in a visual search task. *Journal of Experimental Psychology*, *51*(1), 19–24. Retrieved from http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=pubmed&cmd=Retrieve&dopt=AbstractPlus&list_uids=13286435

Grubert, A., & Eimer, M. (2015). Rapid Parallel Attentional Target Selection in Single-Color and Multiple-Color Visual Search. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(1), 86–101. https://doi.org/10.1037/xhp0000019

Grubert, A., & Eimer, M. (2016). All set, indeed! N2pc components reveal simultaneous attentional control settings for multiple target colors. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(8), 1215–1230. https://doi.org/10.1037/xhp0000221

Gunseli, E., Meeter, M., & Olivers, C. N. L. (2014). Is a search template an ordinary working memory? Comparing electrophysiological markers of working memory maintenance for visual search and recognition. *Neuropsychologia*, *60*, 29–38. https://doi.org/10.1016/j.neuropsychologia.2014.05.012

Hamker, F. H. (2004). A dynamic model of how feature cues guide spatial attention. *Vision Research*, *44*(5), 501–521. https://doi.org/10.1016/j.visres.2003.09.033

Han, S. W., & Kim, M.-S. (2009). Do the contents of working memory capture attention? Yes, but cognitive control matters. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(5), 1292–302. https://doi.org/10.1037/a0016452

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632–5. https://doi.org/10.1038/nature07832

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d′. *Behavior Research Methods, Instruments, & Computers*, *27*(1), 46–51. https://doi.org/10.3758/BF03203619

Hollingworth, A., & Beck, V. M. (2016). Memory-Based Attention Capture When Multiple Items Are Maintained in Visual Working Memory. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(7), 911–917. https://doi.org/10.1037/xhp0000230

Hollingworth, A., & Luck, S. J. (2009). The role of visual working memory (VWM) in the control of gaze during visual search. *Attention, Perception, & Psychophysics*, *71*(4), 936–949. https://doi.org/10.3758/APP.71.4.936

Hollingworth, A., Matsukura, M., & Luck, S. J. (2013). Visual working memory modulates rapid eye movements to simple onset targets. *Psychological Science*, *24*(5), 790–6. https://doi.org/10.1177/0956797612459767

Humphreys, G. W., Stalmann, B. J., & Olivers, C. N. L. (2004). An analysis of the time course of attention in preview search. *Perception & Psychophysics*, *66*(5), 713–30. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/15495898

Irons, J. L., Folk, C. L., & Remington, R. W. (2012). All set! Evidence of simultaneous attentional control settings for multiple target colors. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(3), 758–775. https://doi.org/10.1037/a0026578

Kugler, G., 't Hart, B. M., Kohlbecher, S., Einhäuser, W., & Schneider, E. (2015). Gaze in Visual Search Is Guided More Efficiently by Positive Cues than by Negative Cues. *PloS One*, *10*(12), 1–17. https://doi.org/10.1371/journal.pone.0145910

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, *22*(6), 657–672.

Moher, J., & Egeth, H. E. (2012). The ignoring paradox: Cueing distractor features leads first to selection, then to inhibition of to-be-ignored items. *Attention, Perception, & Psychophysics*, *74*(8), 1590–1605. https://doi.org/10.3758/s13414-012-0358-0

Moher, J., Lakshmanan, B. M., Egeth, H. E., & Ewen, J. B. (2014). Inhibition Drives Early Feature-Based Attention. *Psychological Science*, *25*(2), 315–324. https://doi.org/10.1177/0956797613511257

Moore, K. S., & Weissman, D. H. (2010). Involuntary transfer of a top-down attentional set into the focus of attention: Evidence from a contingent attentional capture paradigm. *Attention, Perception, & Psychophysics*, *72*(6), 1495–1509. https://doi.org/10.3758/APP

Morey, R. D. (2008). Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutorial in Quantitative Methods for Psychology*, *4*(2), 61–64.

Navalpakkam, V., & Itti, L. (2005). Modeling the influence of task on attention. *Vision Research*, *45*(2), 205–231. https://doi.org/10.1016/j.visres.2004.07.042

Olivers, C. N. L. (2009). What drives memory-driven attentional capture? The effects of memory type, display type, and search type. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(5), 1275–91. https://doi.org/10.1037/a0013896

Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(5), 1243–1265. https://doi.org/10.1037/0096-1523.32.5.1243

Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: when it guides attention and when it does not. *Trends in Cognitive Sciences*, *15*(7), 327–34. https://doi.org/10.1016/j.tics.2011.05.004

Schneegans, S., Spencer, J. P., Schöner, G., Hwang, S., & Hollingworth, A. (2014). Dynamic interactions between visual working memory and saccade target selection. *Journal of Vision*, *14(11)*(9), 1–23. https://doi.org/10.1167/14.11.9

Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, *20*(2), 207–14. https://doi.org/10.1111/j.1467-9280.2009.02276.x

Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(2), 248–261. https://doi.org/10.1037/0096-1523.31.2.248

Soto, D., Hodsoll, J. P., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends Cogn Sci*, *12*(9), 342–348. https://doi.org/10.1016/j.tics.2008.05.007

Soto, D., & Humphreys, G. W. (2007). Automatic guidance of visual attention from verbal working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(3), 730–7. https://doi.org/10.1037/0096-1523.33.3.730

Soto, D., Humphreys, G. W., & Heinke, D. (2006). Working memory can guide pop-out search. *Vision Research*, *46*(6–7), 1010–8. https://doi.org/10.1016/j.visres.2005.09.008

Stroud, M. J., Menneer, T., Cave, K. R., & Donnelly, N. (2011). Using the dual-target cost to explore the nature of search target representations. *Journal of Experimental Psychology: Human Perception and Performance*. https://doi.org/10.1037/a0025887

Tsal, Y., & Makovski, T. (2006). The attentional white bear phenomenon: the mandatory allocation of attention to expected distractor locations. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(2), 351–363. https://doi.org/10.1037/0096-1523.32.2.351

Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychonomic Bulletin and Review*, *19*(5), 871–8. https://doi.org/10.3758/s13423-012-0280-4

Vickery, T. J., King, L.-W., & Jiang, Y. V. (2005). Setting up the target template in visual search. *Journal of Vision*, *5*(1), 81–92. https://doi.org/10:1167/5.1.8

Wolfe, J. M. (1994). Guided Search 2.0 A revised model of visual search. *Psychonomic Bulletin & Review*, *1*(2), 202–38. https://doi.org/10.3758/BF03200774

Wolfe, J. M. (2007). Guided Search 4.0: Current Progress With a Model of Visual Search. In W. D. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 99–120). Oxford University Press.

Wolfe, J. M., Horowitz, T. S., Kenner, N. M., Hyle, M., & Vasan, N. (2004). How fast can you change your mind? The speed of top-down guidance in visual search. *Vision Res*, *44*(12), 1411–1426. https://doi.org/10.1016/j.visres.2003.11.024

Woodman, G. F., & Arita, J. T. (2011). Direct electrophysiological measurement of attentional templates in visual working memory. *Psychological Science*, *22*(2), 212–5. https://doi.org/10.1177/0956797610395395

Woodman, G. F., & Luck, S. J. (2007). Do the contents of visual working memory automatically influence attentional selection during visual search? *Journal of Experimental Psychology: Human Perception and Performance*, *33*(2), 363–377. https://doi.org/10.1037/0096-1523.33.2.363

Zelinsky, G. J. (2008). A theory of eye movements during target acquisition. *Psychological Review*, *115*(4), 787–835. https://doi.org/10.1037/a0013118