VISUAL COGNITION, 2003, *10* (3), 257–298

Memory for rejected distractors in visual search?

Todd S. Horowitz and Jeremy M. Wolfe

Brigham & Women's Hospital and Harvard Medical School, Boston

Theories of visual search have generally assumed that rejected distractors are marked so as to avoid further processing of these items (*memory-driven search*). To test this assumption, Horowitz and Wolfe (1998) developed the randomized search paradigm, in which standard *static* search is compared to *dynamic* search where items are randomly replotted at new locations throughout a trial. Memory-driven search predicts that search slopes should double in the dynamic condition. After reviewing earlier findings that slopes were similar in the two conditions, we present two new experiments. Experiment 1 replicated and extended our previous findings using a larger range of set sizes, a slower rate of change, and adding a fixed location dynamic condition. Experiment 2 employed stimuli that required overt fixation. Neither experiment showed evidence for memory-driven search. We conclude that visual search is best understood as a series of successive judgements of the momentary probability of target presence.

Visual search is a common everyday behaviour (Where are my keys? Are my friends in the crowd of people greeting my aeroplane?) with important practical implications for fields such as medicine (Gale & Walker, 1993), driving (Fairclough & Maternaghan, 1993), and sport (Helsen & Pauwels, 1993). Psychologists have long been interested in visual search behaviour (Ruckmick, 1926). Experimental investigations go back as far as Kingsley (1932, 1934) and continue to proliferate (for recent reviews, see Chun & Wolfe, in press; Pashler, 1997; Wolfe, 1998a).

Although the modern study of attention began with auditory tasks (Broadbent, 1957; Cherry, 1953; Treisman, 1964), since the late 1970s (Schneider &

Please address all correspondence to: T.S. Horowitz, Center for Ophthalmic Research, Brigham & Women's Hospital, 221 Longwood Avenue, Boston, MA 02115, USA.
Email: toddh@search.bwh.harvard.edu

The work in this paper was aided by discussions with many people who have given us ideas, criticisms, and data following the publication of our initial work on this topic. We are particularly indebted to Alex Backer, Kyle Cave, Glyn Humphreys, Beena Khurana, Ray Klein, Diane Kornbrot, Arni Kristjánsson, Herman Müller, David Shore, Anne Treisman, Trish Van Zandt, and an anonymous reviewer. We would also like to thank Megan Hyle and George Alvarez for data collection and helpful comments.

http://www.tandf.co.uk/journals/pp/13506285.html DOI:10.1080/13506280143000005

Shiffrin, 1977; Treisman & Gelade, 1980), attention research has concentrated primarily on the visual domain, and the ever-growing visual search literature has often been at the centre of debates about the nature of attention (Duncan & Humphreys, 1992; Eckstein, 1998; Pashler & Badgio, 1985; Treisman, 1992; Yantis & Johnston, 1990).

A fundamental tenet of visual search models has been that information accumulates over time during a search. The most common (and intuitively understandable) instantiation of this postulate is the assumption of *inhibitory tagging* in serial models of search. In many models, the focus of attention is deployed from place to place or from object to object during the course of a search. Examples of such models include Treisman's ''Feature Integration Theory'' (Treisman, 1993; Treisman & Gelade, 1980), Wolfe's ''Guided Search'' (Cave & Wolfe, 1990; Wolfe, 1994, 1996; Wolfe, Cave, & Franzel, 1989), as well as a number of others (Cohen & Ruppin, 1999; Grossberg, Mingolla, & Ross, 1994; Itti & Koch, 2000; Koch & Ullman, 1985). Other models assume that information is accumulated in parallel from multiple items (Humphreys & Muller, 1993; Kinchla, 1974; McElree & Carrasco, in press; Palmer & McLean, 1995). Some models blur the distinction between serial and parallel accumulation of information (for a discussion see Moore & Wolfe, 2000). Across this diversity of theoretical positions, it has been almost universally assumed that a rejected distractor is somehow marked off in order to prevent attention from being redirected to that distractor at a later time. Sometimes the assumption is explicit. Sometimes it is implicit. It is as ubiquitous as it is reasonable. Nevertheless, in this paper, we present evidence that suggests that the assumption is incorrect.

Mathematically, for the class of serial models of attentional deployment, the question can be framed as asking whether visual search proceeds by *sampling without replacement* (the standard, inhibitory tagging account) or by *sampling with replacement*. We will refer to a system that samples without replacement as *amnesic*. However, it is important to remember that we are concerned with a very narrow sense of terms like ''amnesia'' and ''memory''. It is beyond question that subjects can explicitly remember the identities and locations of targets (Gibson, Li, Skow, Salvagni, & Cooke, 2000; Ward & McClelland, 1989). Moreover, it is clear that implicit memory for the contents and consequences of previous searches can influence current search (Chun & Jiang, 1998; Hillstrom, 2000; Maljkovic & Nakayama, 1994, 1996; for a thorough review of the role of memory in search, see Shore & Klein, 2000). We are concerned with memory for rejected distractors during the course of a single search. Do search mechanisms make use of knowledge of the history of an ongoing search in order to improve the efficiency of that search? In the context of serial models, this question can be rephrased: Does the system behave as if it uses inhibitory tagging of rejected distractors? (Note: ''Serial'' models, for this purpose, include models like Guided Search that propose parallel guidance of serially deployed attention.) We define a *memory-driven* model of search as one which assumes inhibitory tagging, and an *amnesic* model of search as one that does not.

Although we will frame our arguments in the context of serial models of search, our results also have implications for parallel limited- (Ward & McClelland, 1989) and unlimited-capacity (Eckstein, 1998; Palmer & McLean, 1995) models of search. For parallel models, the question is whether information about targets and distractors accumulates over time. The data presented in this paper are problematic for both parallel and serial models. We will return to these issues in the discussion.

Memory-driven and amnesic models represent two ends of a continuum. We can characterize this continuum, again within the context of serial deployments of attention, by asking how many rejected distractors can be marked. If we call this number C, then the amnesic model assumes that $C = 0$, whereas for the memory-driven model, $C = \infty$. Phrased in these terms, the memory-driven model seems prima facie absurd, a ''straw man''; who would believe an infinite memory? Why, then, is the memory-driven model so prevalent, and why has it received such a vigorous defence (e.g., Kristjánsson, 2000; Peterson, Kramer, Wang, Irwin, & McCarley, 2001; Shore & Klein, 2000)? First, in practice, C does not have to be infinite in order for the memory-driven model to appropriately characterize the data. As long as the set size, n, is not much larger than C, such a ''limited capacity memory'' (Horowitz & Wolfe, 2001) would behave as if memory were infinite. Since n rarely exceeds 36 in visual search experiments, a capacious, but not infinite memory would serve. Formally, the experiments described here will test the hypothesis that $C = \infty$, which we argue has been a basic assumption of visual search theory since the 1970s, against the null hypothesis, that $C = 0$. However, we acknowledge that we can really only distinguish between ''C is large'' and ''C is small'', with ''large'' and ''small'' defined relative to the range of set sizes used in visual search experiments. In Horowitz and Wolfe (2001), for instance, we were able to fit a particular limited capacity memory model to our data, and estimated that subjects were able to keep track of 0–3 prior distractors. For the purposes of most visual search experiments, a system with a memory of 1, 2, or 3 items is virtually memoryless, just as a memory for 20 items would be virtually infinite.

There are several lines of evidence which point to amnesic visual search: Randomized search experiments (Gibson et al., 2000; Horowitz & Wolfe, 1998; Kristjánsson, 2000), multiple target search experiments (Horowitz & Wolfe, 2001), attentional reaction time search experiments (Alvarez, Horowitz, & Wolfe, 2000), and eye-movement studies (Arani, Karwan, & Drury, 1984; Chan & Courtney, 1998; Gilchrist & Harvey, 2000). The focus of this paper is on the randomized search experiments. These are experiments in which the items in the search display are randomly repositioned during the course of a search. Horowitz and Wolfe (1998) presented an initial, brief account of experiments using this method. The introduction to this paper recapitulates that previous work. The core finding was that randomly repositioning items every 100 ms during a search did not alter the efficiency of search—a finding difficult to explain if that efficiency is based on some sort of inhibitory tagging of rejected distractors in normal search.

There are some problems with the repeated search methodology and there have been several challenges posed by other experiments using the same technique. The central purpose of this paper is to present data from two new experiments. The first is an extension of the original repeated search paradigm and is designed to remedy these problems. To anticipate the results, we continue to find that the efficiency of randomized search is comparable to the efficiency of standard, static search. The second experiment uses the randomized search method to investigate the possibility that visual search keeps track of overt deployments of the eyes even if it does not keep track of covert deployments of attention. Again, our data reject the hypothesis that the visual system makes use of the history of a visual search, even when each item must be overtly fixated.

However, there remain difficulties with these data. Randomized search is consistently slower and less accurate than comparable standard, static search. The difficulties will be discussed as will several lines of converging evidence for amnesic search.

THE RANDOMIZED SEARCH PARADIGM

In the randomized search paradigm (Horowitz & Wolfe, 1998), we compare two search conditions, static and dynamic. In the dynamic condition, we replot the search stimuli periodically during a trial; typically, a new frame is presented every 100 ms. The content of the display remains the same. The same number of distractors are shown on every frame, and the target, if present, is present on every frame. However, the locations of the stimuli are changed randomly from frame to frame. This thwarts any possible marking of rejected distractors, and prevents the parallel accumulation of information about the identity of any stimulus by localized detectors.

The dynamic search condition is always compared to a static search condition, which is as similar to the dynamic search condition as possible (identical stimuli, task requirements, etc.), except that only one frame is shown, so that stimuli do not change locations during the trial and tagging of rejected distractors is possible. Nearly all published search experiments (even those involving motion) are ''static'' searches by this definition, so the static condition serves as a baseline.

Reaction time (RT) is the primary measure of interest in these experiments, and the predictions of the standard, memory-driven model for mean RT under these conditions are straightforward. If items are sampled without replacement, search finishing times (FTs) will be distributed as a negative hypergeometrical distribution (Johnson & Kotz, 1977), and the expected mean value of this distribution $E[S]$, assuming t targets and d distractors, is given by:

$$E[S] = \frac{(t + d + 1)}{t + 1} \tag{1}$$

When items are sampled with replacement, however, a geometric distribution is produced, and the expected mean is given by:

$$E[S] = 1 + \frac{d}{t} \tag{2}$$

If inhibitory tagging is used in the static condition and prevented in the dynamic case, then the ratio of dynamic finishing times to static finishing times will be given by:

$$\mathrm{FT}_{\mathrm{Dynamic}}/\mathrm{FT}_{\mathrm{Static}} = 1+(d/t)/((t + d + 1)/(t + 1)(3)) \tag{3}$$

If we assume $t = 1$ and set size $n = d + 1$, we get:

$$\mathrm{FT}_{\mathrm{Dynamic}}/\mathrm{FT}_{\mathrm{Static}} = 2n/(n + 1) \tag{4}$$

So finishing times in the dynamic condition should be nearly double those in the static condition. However, the FT (when the target is found) is not the same as the observed RT, which also includes non-search-related components such as initial stimulus processing time, decision time, and motor planning and execution. Since these factors should be constant with set size, it is necessary to vary the set size and compute the RT × Set size slopes for the two conditions, which should follow the same ratio: Dynamic slopes (where inhibitory tagging is prevented) should be twice static slopes (where inhibitory tagging is possible).

If search is amnesic, however, then sampling with replacement is used in *both* dynamic and static conditions. In this case, the set size dependent component of RT should be constant across the two conditions, and the two slopes should be the same.

The two hypothetical search processes differ not only on the central tendency of the RTs they produce, but also the overall shape of the distributions themselves. If search is proceeding without replacement through the array, then the target will be found in no more than n steps, with each step equally likely. That is, it is just as likely that the target will be the first item attended ($p = 1/n$) as the last ($p = 1/n$). Therefore, the probability density function will be rectangular (Kontsevich, 2001). For an amnesic search, however, the maximum number of steps is infinite (though in practice of course subjects will give up at some point, curtailing the distribution), and the probability of the target being found on a given step decreases exponentially. Figure 1 shows the idealized forms of these

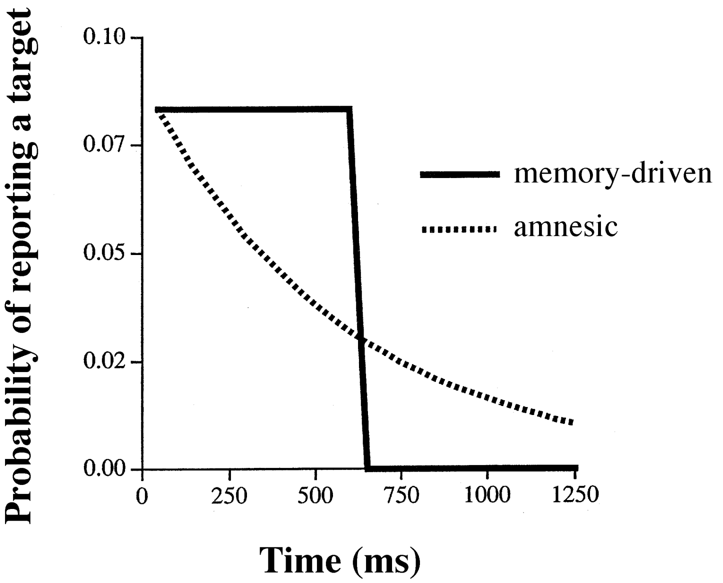

Figure 1. Idealized probability density functions for memory-driven (solid line) and amnesic search (dotted line). Functions generated assuming a sampling rate of 20 Hz and a set size of 12 items.

distributions, and hypothetical resulting RT distributions (assuming a search rate of 50 ms/item and a set size of 12 items for both).

Horowitz and Wolfe (1998) tested observers in dynamic and static search in three experiments. In Experiments 1 and 2, the target was the letter T, and distractors were Ls. The stimuli were black on a grey background, each measuring $1.28^\circ \times 1.28^\circ$. Strokes were 0.26° in width. Masks were 2×2 black grids constructed from the same strokes as the letters, such that any (orthogonal) rotation of either a T or an L would be completely obstructed by the mask. Stimuli were presented on a blank grey field. Targets and distractors could appear in any of the four cardinal orientations. In Experiment 1, stimuli were presented for 83.33 ms on each frame, followed by a 27.78 ms masking frame consisting of a mask grid substituted at the location of each letter. In Experiment 2, there were no masks; each stimulus frame was presented for 106.70 ms and followed immediately by the next frame.

An important methodological problem in the randomized search paradigm is ensuring that subjects use similar search procedures in both static and dynamic conditions. For instance, if the target can be replotted at any random location on a given frame, then it will be as efficient for the observer to fixate in one spot and wait for the target to appear as to go and search for it. Our experiments included manipulations designed to thwart this ‘‘sit-and-wait'' strategy. In Experiment 1, we generated four independent search frames and repeated them through 5.25 cycles (21 frames). Stimuli were placed in an unmarked 8 × 8 grid of possible locations; the centre of each cell was 2.1° from its neighbour, with the entire field measuring 15.52° × 15.52°. Since the target could only appear in four locations (out of a possible sixty-four), the sit-and-wait strategy would fail on 93.75% of trials. In Experiment 2, we generated 20 independent frames (presented for one cycle), but fixed the eccentricity of the target. Potential stimulus locations were arranged on four concentric circles at radii of 2°, 4°, 6°, and 8°. There were 4 locations in the inner circle, 8 locations in the next, 12 locations in the next, and 16 locations in the outermost circle. All locations on a given circle were equally spaced around the circumference. On a given trial, the target was presented at locations on only one of the four circles. Subjects were not informed about this manipulation. In this situation, a sit-and-wait on one location strategy should succeed on 23.7% of trials, and fail on 76.3% of trials. This calculation relies on the assumption that each location is equally likely to be chosen for monitoring by the subject, which is probably not true; there is a bias for central locations over eccentric locations (Carrasco, Evert, Chang, & Katz, 1995; Wolfe, O'Neill, & Bennett, 1998). However, this would only make the sit-and-wait strategy even less plausible, since central locations were less likely to hold a target than eccentric locations in this experiment.

Figure 2 plots the RT × Set size functions for the masked and unmasked experiments on separate panels, along with the predictions of the standard (memory-driven) model. RT × Set size slopes are given in Table 1.

We can draw three conclusions from these data:

(1) The predictions of the standard, memory-driven model are disconfirmed. The target-present slopes for the dynamic conditions do not remotely approach twice the slope of the static conditions. In fact, dynamic and static target-present slopes are statistically indistinguishable, as predicted by an amnesic account of visual search.

TABLE 1
RT × Set size slopes (ms/item) from Experiments 1 and 2 of Horowitz and Wolfe

Masks	*Target*	*Static*	*Dynamic*
Yes (Exp. 1)	Present	18.76	18.13
	Absent	50.42	23.74
No (Exp. 2)	Present	20.89	11.51
	Absent	42.00	12.18

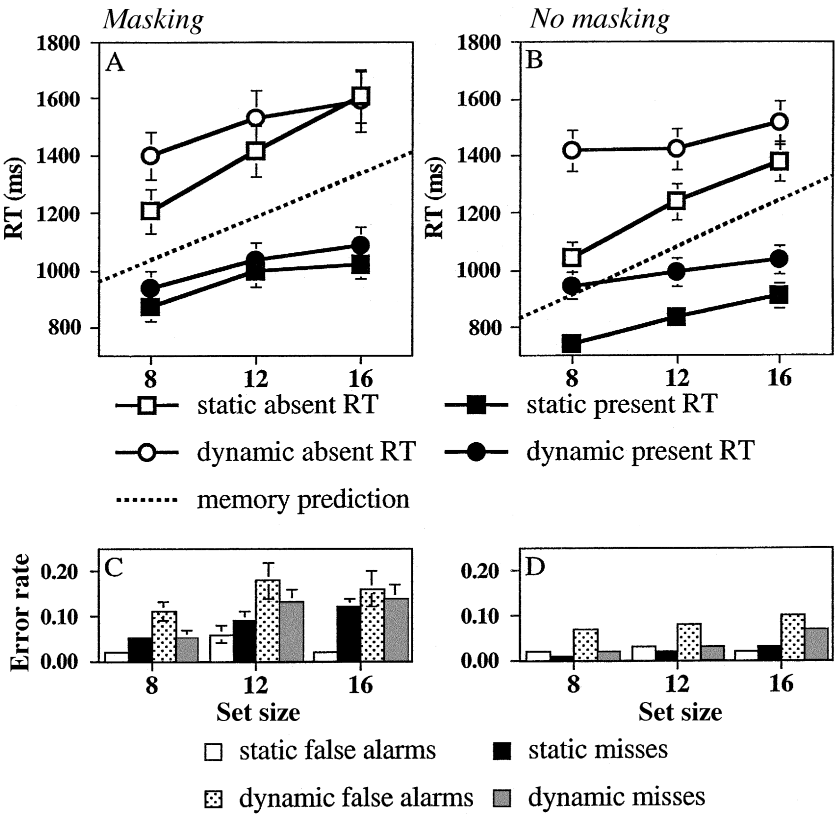

Figure 2. RT and error data from the first two experiments of Horowitz and Wolfe. Panels A and C show data from Experiment 1, and panels B and D show data from Experiment 2. RT data are shown in panels A and B, errors in panels C and D. In this and all subsequent figures, error bars indicate the standard error of the mean. Dotted lines in panels A and B depict a function with twice the slope of the static target-present function.

(2) Local masking does not change the slope relationship between static and dynamic search. Dynamic and static target-present slopes are statistically identical when stimuli are masked on each frame, and statistically identical when the stimuli are unmasked.

(3) Static and dynamic search do not produce identical results: the target-present slopes are the same in both conditions, regardless of masking, but there are differences in target-absent slopes and in the target-present intercepts, as well as in error rates.

Of these differences between static and dynamic search results, the target-absent data are the least worrisome. Although it is clear that target-present trials can be terminated when the target is found, termination rules for target-absent trials are more ambiguous and criterion-dependent (Chun & Wolfe, 1996; Zenger & Fahle, 1997). Target-absent slopes may differ between the dynamic and static conditions for reasons that have nothing to do with the question of memory. For instance, if observers search for a set-size-dependent length of time before concluding that no target is present, and the dynamic manipulation makes it more difficult to ascertain set size, then we might expect the flattening of slopes observed in Figure 2.

Turning to the target-present results, randomized search experiments reliably produce longer RTs in the dynamic than in the static condition. This difference in intercept of the RT $\times$ Set size function suggests that dynamic and static search are treated differently at some, non-search-related processing stage (Pashler & Badgio, 1985; Roberts & Sternberg, 1993). The manipulation did not influence the efficiency of the search, as reflected in the slope of the RT $\times$ Set size functions. The RT increase could have several causes. It is possible that it may take subjects longer to commit to a response after they find the target, perhaps because the disappearance of the target within 100 ms degrades its representation. The differences between the masked and unmasked versions of the experiment support this intuition. Masking increased the static intercept by 160 ms, $t(8) = 2.62$, $p < .05$, but had no effect on the dynamic intercepts, $t(8) < 1$. This suggests that the dynamic condition is, in effect, already masked by the rapid onsets and offsets of search items.

The error rates are the most troublesome aspect of the results shown in Figure 2. Error rates are fairly high in the dynamic conditions (6–13% overall, compared to 2–6% in the static conditions). This complicates the interpretation of RT data. Observers may trade off speed for error, achieving fast RTs in a difficult condition at the cost of more errors so that observed RT is equivalent to that in an easier condition. Could the errors explain the failure to find the predicted 2:1 ratio between dynamic and static target-present slopes? Horowitz and Wolfe (1998) argued that the answer was ‘‘no’’.[1] However, the disparities in error rates, intercept, and target-absent slopes weaken the argument. Therefore, we included a third experiment. We reasoned that the high error rates and the aberrant target-absent slopes had a similar cause, the difficulty of deciding when to stop searching when one has not found a target. Accordingly, we switched to a 2AFC version of the task. The stimuli consisted of 24 letters of the alphabet (excluding ‘‘I’’ and ‘‘J’’). The letters were drawn in 48-point Arial Bold font in black capitals. At a viewing distance of 57 cm, the letters all subtended approximately 1.3° in height and varied from 1.0° to 1.7° in width. The masks were 2 $\times$ 2 black grids constructed from strokes 0.26° in width and 1.7° long.

Stimuli were presented on a blank grey field. Potential stimulus locations were arranged in four concentric circles at radii of 2°, 4°, 6°, and 8°. There were

[1] See search.bwh.harvard.edu

4 locations in the inner circle, 8 locations in the next, 12 locations in the next, and 16 locations in the outermost circle. All locations on a given circle were equally spaced around the circumference.

The targets were the letters ''E'' and ''N''. One target was present on every trial, and observers were asked to report which target was present. Distractor letters were drawn at random from the remaining letters in the set. All the letters in a given frame were unique. In the dynamic condition, each frame was presented for 106.7 ms and then immediately replaced the following frame, with no ISI or mask. As before, dynamic targets were restricted to a single eccentricity on a given trial. In the static condition, only the first frame was presented, and it remained on the screen for 20*106.7 or 2134 ms. In either condition, after 2134 ms had passed, the stimuli were replaced with a masking frame which consisted of masks at each of the possible stimulus locations. If the subject had not responded by this time, the masking frame remained on until the response. Otherwise, the mask was presented for 13.9 ms and was replaced by a feedback message indicating accuracy and RT for that trial.

RTs, slopes, and errors from this experiment are shown in Figure 3. We can see that errors were substantially reduced in this experiment. Although error rates are still higher for the dynamic condition (5.4% compared to 2.7%), the rates are low enough to give us confidence in interpreting the RTs. It is highly improbable that the relatively few errors in dynamic condition are masking a ''true'' slope which is twice the static slope. In fact, the dynamic slope in this experiment is somewhat (though not significantly) shallower than the static slope, suggesting that observers were engaging in a small speed–error tradeoff, masking a ''true'' slope roughly equal to the static slope.

There is still one important difference between the static and dynamic conditions: A robust 182 ms intercept difference. Again, this may be due to some masking effect of the sudden onsets and offsets in the dynamic condition.

Since we reported these results, several investigators have employed the randomized search paradigm to test the hypothesis that visual search is amnesic. Gibson et al. (2000) have replicated our Experiment 2 in the context of an experiment that demonstrates that observers do have a memory for targets discovered during search. In addition to a straight replication of our results, they also tested a condition in which there were one or two identical targets presented simultaneously. Observers were asked to report how many targets were presented. In order to perform this task, subjects must be able to distinguish between tokens of the same target type, distinguished only by their location. This would only be possible if subjects were able to apprehend both targets during a single frame, and attach some sort of location index (Pylyshyn, 1989) to at least one target. Since subjects were able to perform the task easily in the static condition, we can conclude that, given sufficient time, the visual system can spatially index at least one target. As accuracy was near chance in the dynamic condition, 100 ms must be insufficient time to locate a target, index it,

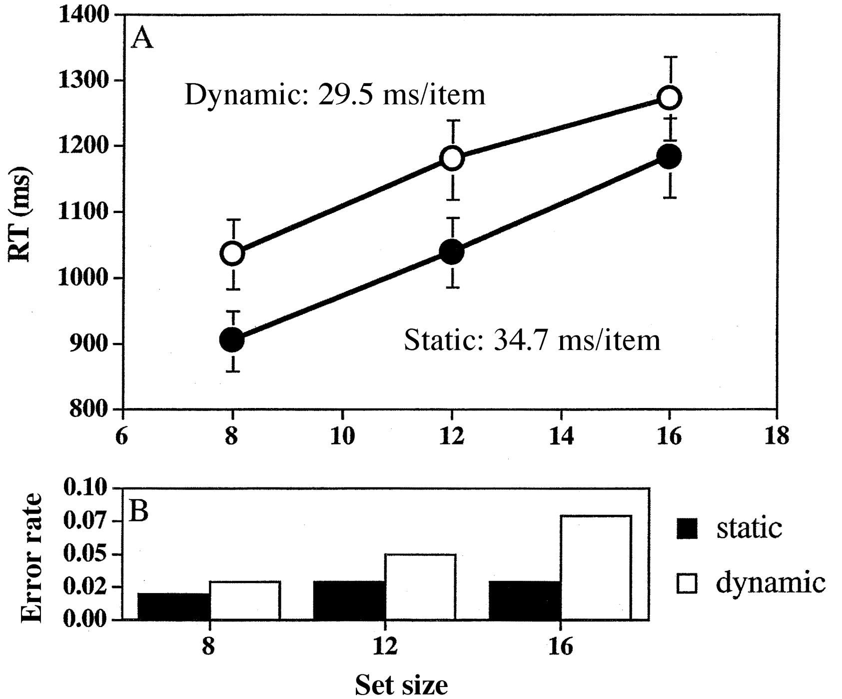

Figure 3. Data from Experiment 3 of Horowitz and Wolfe. RTs are shown in panel A, along with the slopes of the RT × set size functions. Errors are depicted in panel B.

and locate a second target. Furthermore, these data are consistent with the idea that subjects can only take a small number of samples from the display on each frame, rather than developing global information about stimulus identity over time. If the visual system could accumulate information in a parallel, global fashion, then a dynamic display with one target should have been distinguishable from one with two targets.

However, it is important to keep in mind that the Gibson et al. data do not demonstrate any memory for rejected distractors. In the standard, 1 versus 0 targets task, information about the locations of stimuli is irrelevant to the task, and therefore randomizing the locations of stimuli does not affect performance. In the 1 versus 2 targets task employed by Gibson et al., the location of targets becomes relevant, because it is the only feature which distinguishes the two targets. In the static case, target location information is available. In the dynamic condition, target location is not available, unless both targets are apprehended on the same frame, which is unlikely at 100 ms/frame. The lack of location information in the dynamic search for 1 versus 2 targets means that both amnesic and memory-driven models will predict disrupted performance in this case.

The fact that subjects can perform this task in the static case does prove that subjects can remember the spatial location of at least one target (see also Schneider, 1999). Remember that the form of memory that we are concerned with here is the ability to avoid attending to a previously attended item. Can the visual system use spatial memory to avoid attending to previously identified targets? The burden of the experiments presented here, as well as Gibson et al.'s replication of our original paradigm (not to mention Horowitz & Wolfe, 2001), is that the visual system does not avoid resampling previously rejected distractors. Although it is possible that targets are different, it is not necessary to assume that the system avoids re-attending to previously identified targets. All we have to assume is that subjects can encode the location of the first target they encounter, so that the next time a target is found, they can determine whether or not it was a previously identified target.

Brady and Lawrie (1999) have also replicated our findings using faces instead of letters as stimuli.

METHODOLOGICAL PROBLEMS WITH THE RANDOMIZED SEARCH PROCEDURE

The finding that dynamic search yields search efficiencies comparable to static search appears to be robust (for exceptions, see later). However, there remain some problems with the procedure we originally used. The difference in intercept between the two conditions (also observed in the data of Brady & Lawrie, 1999; Gibson et al., 2000) remains puzzling. A more serious methodological problem arises from the fixed duration of stimuli. On each trial, stimuli were presented for 20 frames, or roughly 2100 ms. A truly amnesic search could, in theory, produce RTs of any arbitrary length since, by chance, an observer might not deploy attention to the target until attention had been incorrectly deployed 10, 100, or 1000 times. A completely memory-driven search would not produce these long RTs because, if the items are sampled without replacement, the maximum number of attentional deployments could never exceed the number of items in the display.

However, limiting stimulus presentation has the effect of terminating search after 2 s. This could eliminate more long RTs in the dynamic case than in the static case (assuming the static case to be memory-driven). This, in turn, would lower the mean RTs, artificially flattening the dynamic slopes more than the static slopes (Backer & Preal, 1999).

In addition to this theoretical problem, there have also been some experiments, published and unpublished, challenging our conclusions. First, Kristjánnson (2000) conducted two experiments using the randomized search procedure. In the first experiment, instead of selecting completely random new locations on each frame of the dynamic condition, he randomly shuffled stimuli amongst the same set locations on each frame of a trial. This form of dynamic search stimuli yielded steeper slopes than the static control case (Treisman, personal communication, has also reported similar data). Next, he replicated our original Experiment 2 (the no-masking experiment), but extended the range of set sizes out to 60 items. For smaller set sizes (4–16), Kristjánnson's data replicate ours: Dynamic and static conditions show similar slopes. However, above 20 items, the curves start to diverge, with the dynamic case showing steeper slopes then the static control. Furthermore, the static slope remains roughly constant over the full range of set sizes, so the difference is due to an increase in the dynamic slope, rather than to some flattening of the static slope.

Finally, static and dynamic conditions are simply *never* going to be identical because the dynamic conditions are continuously changing. The multiple frames of the dynamic conditions may mask each other. Moreover, the dynamic condition might encourage subjects to wait for targets to appear at a location or to reappear after being glimpsed at a location.

With these limitations in mind, we designed an experiment which would

(1) Use a larger range of set sizes, to replicate Kristjánnson (2000).
(2) Compare the fully randomized dynamic condition with the fixed-location randomization, replicating Kristjánnson (2000) and Treisman (personal communication).
(3) Allow stimuli to be visible until response, addressing Backer and Preal's (1999) concerns.
(4) Present stimuli at a slower frame rate, in order to reduce the difference between dynamic and static conditions.

EXPERIMENT 1: IMPROVED RANDOMIZED SEARCH

For this experiment, we adapted the task used in Experiment 3 of Horowitz and Wolfe (1998): 2AFC search for ''E'' or ''N''. There were four critical differences between this experiment and the experiments reported in Horowitz and Wolfe (1998):

(1) Display set size was varied from 9 to 60 items, allowing us to verify Kristjánnson's (2000) findings.

(2) There were two dynamic conditions. The *free dynamic* condition was identical to the previously reported dynamic conditions, in which the stimulus locations for each frame were drawn from all possible locations in the display. The *fixed dynamic* condition was modelled after the displays used by Kristjánnson (2000) and Treisman (personal communication), where the set of locations was constant from frame to frame within a trial but which location held which stimulus was varied randomly.

(3) The display was present until response, which should reduce any potential truncation of the RT distribution in the dynamic condition.

(4) Finally, the SOA from one frame to the next was extended to 500 ms in order to reduce the difference between dynamic and static conditions by reducing the number of transients.

Method

Subjects. Sixteen subjects from our volunteer subject pool served as subjects in return for compensation of $7/hour. All subjects had normal or corrected to normal visual acuity and passed the Ishihara colour screen. All gave informed consent prior to their participation.

Equipment and stimuli. Stimuli were presented on a 21-inch Mitsubishi monitor running at a refresh rate of 75 Hz and controlled by a PowerMacintosh 4400 computer. The experiment was programmed in Matlab (Mathworks, Inc.) using the Psychophysics Toolbox (Brainard, 1997). Responses were collected using an Apple ADB keyboard.

The targets were the letters ''E'' and ''N'', and the distractors were drawn randomly from the remaining letters of the alphabet (except for ''I'', ''J'', and ''W''). Letters were drawn in 48-point Arial font in black on a grey background. At a viewing distance of 57 cm, the letters all subtended approximately 1.3° in height and varied from 1.0° to 1.7° in width.

Procedure. Set size could be 9, 12, 18, 24, 42, 48, or 60 items. There were 60 possible stimulus locations, arranged on five imaginary circles, concentric around the fixation cross at eccentricities of 1.94°, 3.88°, 5.82°, 7.76°, and 9.70° (at a viewing distance of 57 cm). The innermost circle had 4 positions, the second ring had 8, the third 12, the fourth 16, and the fifth and outermost 20.

For each trial in the dynamic condition, 50 search frames were generated. Each frame contained exactly one target. Across trials, the two targets ''E'' and ''N'' were equally likely. Within a trial, all frames contained the same target letter. The same set of distractors was present on each frame (distractors were selected to minimize the duplication of distractor letters on a given trial). On the first frame, a set of locations equal to the set size was selected at random, and the stimuli were randomly distributed among these locations. In the fixed dynamic condition, this same set of locations was used on each succeeding frame, but the distractors were redistributed among the fixed set of locations at random. In the free dynamic condition, a new set of distractor locations was selected on each frame. Target locations were not selected entirely at random. To defeat the ''sit-and-wait'' strategy described earlier, targets could only appear in a fixed set of locations in both the fixed and free dynamic conditions. One-third of the stimulus locations selected on the first frame were randomly designated as potential target locations, and these locations were constant throughout the trial. One of the potential target locations was selected at random on each frame, with the proviso that the target not be plotted at the same location on two successive frames.

In the static condition, only a single frame was generated and presented. As in the dynamic conditions, there was a target on every trial, which could be ''E'' or ''N'' with equal probability.

Each trial began with a fixation cross (''+'' drawn in the same font and size as the letter stimuli) presented alone for 500 ms. At that point the fixation cross remained on the screen and the first search frame was presented. In the dynamic conditions, successive frames were presented every 500 ms until a response was made or 5 s, whichever came first. In the static condition, one frame was presented until response (or 5 s). After response, observers were given feedback as to their RT and whether or not they made the correct response. Feedback was presented for 500 ms.

The three conditions were presented in separate blocks, and the order of blocks was counterbalanced across subjects. Each block began with 50 practice trials, from which no data were collected, followed by 350 experimental trials. Subjects were allowed to press the space bar on the keyboard to initiate a pause whenever necessary.

Data analysis. As in Horowitz and Wolfe (1998), the test of the memory-driven hypothesis is focused on the slope of the RT $\times$ Set size functions. As we argued earlier, observed RT by itself does not provide a test of sampling regimes. The memory-driven and amnesic hypotheses make different predictions about the set-size-dependent component of RT, which reflects FT. But on what measure of RT should the slopes be computed? We will conduct three separate analyses. As before, we will analyse mean RT. For mean RT, the memory-driven model predicts that slopes in the dynamic condition should be double those in the static condition.

The memory-driven search hypothesis can be tested on the median RTs as well as on the mean. Recall that the prediction of a doubled slope in the dynamic condition if search is memory driven is produced in part by the long tail of the amnesic RT distribution. In theory, dynamic/amnesic RTs can be extremely long. However, the extremely long RTs are also extremely rare. This produces sampling problems in experiments with a reasonable number of trials. The effects of this tail of long RTs are reduced if we look at the median rather than the mean of the RT distribution. If RT $\times$ Set size slopes are based on medians, the theoretical ratio between dynamic:static slopes should be $\ln(.5)/(.5)$, or roughly 1.39 if the dynamic condition is amnesic while the static condition is memory driven.[2]

[2] We are grateful to Trisha Van Zandt for this calculation.

The two models, as we have presented them, do not explicitly predict error rates. In fact, the models assume perfect accuracy. This is of course an unreasonable assumption, on theoretical grounds as well as according to common sense. If search truly were amnesic, on a small but finite number of trials, the target might not be found for minutes, hours, or even days. Therefore, subjects must employ some arbitrary criterion for when to stop searching and guess, whether or not they can be sure that there is always a target to be found. Even if we assume that there are no stimulus errors (distractors incorrectly identified as targets, and vice versa), subjects must commit some mistakes. If error rates were low and did not vary across conditions, we could still ignore them, as is the typical practice in the search literature. Although we expect errors to be reduced in this experiment when compared to our previous study (Horowitz & Wolfe, 1998), it is possible that we will observe significantly greater errors in the dynamic condition. Furthermore, errors may increase more steeply with set size in the dynamic condition. Such results complicate the interpretation of RTs and the RT $\times$ Set size slopes. Instead of analysing the RT and error data separately, a different approach is to combine the two into a single measure. Townsend and Ashby (1978, 1983) have suggested that, in cases where accuracy levels are not near ceiling, dividing RT by accuracy provides a useful measure of the overall capacity of the system, regardless of the various speed/error tradeoffs made by different subjects in different conditions. Since RT/accuracy has the dimensions of milliseconds, we will call this measure ‘‘corrected RT’’.[3] Keep in mind that corrected RT takes into account both latency and accuracy. Increased RT and decreased accuracy alike lead to increases in corrected RT.

We analysed all three RT measures, as well as error rates, using 3×7 within-subjects ANOVAs with factors condition (fixed dynamic, free dynamic, and static) and set size (9, 12, 18, 24, 36, 48, and 60). The RT $\times$ Set size functions were computed for each RT measure individually for each subject and both slopes and intercepts were analysed with one-way within-subjects ANOVA with condition as the factor.

Results

Figure 4 plots RTs for correct trials against set size using means, medians, and corrected medians. Error rates are also shown. Slopes and intercepts are given in Table 2.

Mean RT (Figure 4A) showed a strong effect of set size, $F(6, 90) = 58.55$, $p < .0001$. However, the effect of condition was not significant, $F(2, 30) = 1.77$, $p = .187$, nor was there any interaction, $F(12, 180) < 1$. To evaluate the possibility that the variance between the two dynamic conditions was blurring the overall difference between dynamic and static conditions, we analysed fixed dynamic

[3] We thank Kyle Cave for this suggestion.

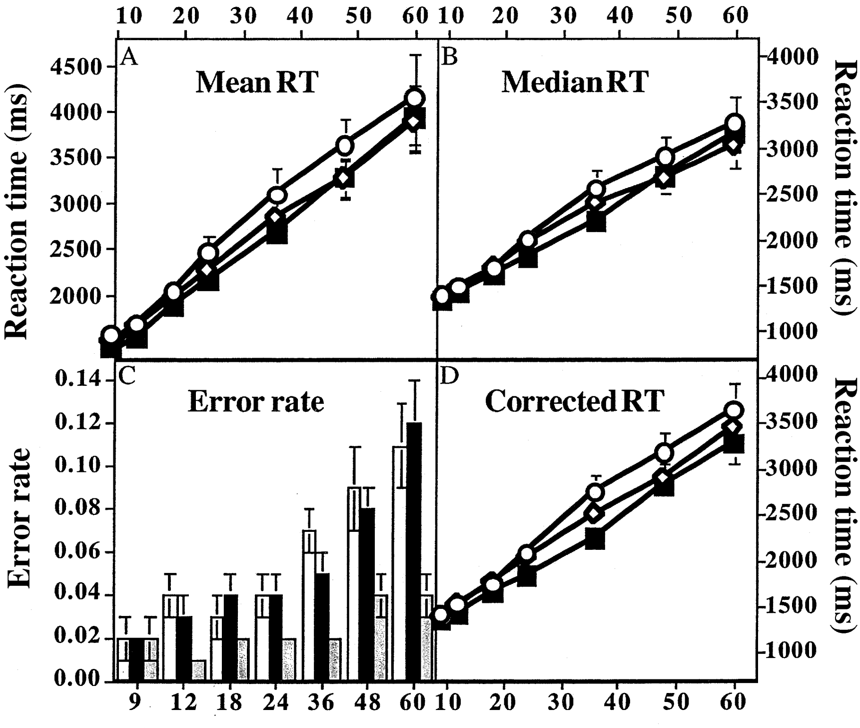

Figure 4. Results from Experiment 1. Mean RTs are shown in panel A, median RTs in panel B, error rates in panel C, and corrected median RTs in panel D. In panels A, B, and D, solid squares denote the static condition, open diamonds the free dynamic condition, and open circles the fixed dynamic condition. In panel C, open bars represent data from the fixed dynamic condition, solid bars the free dynamic condition, and stippled bars the static condition.

TABLE 2
Slopes and intercepts of the RT × Set size functions from Experiment 1

Based on	*Measure*	*Condition*		
		Fixed	*Free*	*Static*
Means	Slope	51.82	46.18	48.53
	Intercept	1119.98	1106.13	967.37
Medians	Slope	38.15	32.66	35.48
	Intercept	1050.97	1107.05	980.37
Corrected medians	Slope	45.10	39.56	37.68
	Intercept	997.36	1045.68	971.99

and static RTs in one analysis and free dynamic versus static RTs in another, separate analysis. The results of these sub-set ANOVAs were identical to the primary ANOVA reported previously, although the trend toward slower RTs for the fixed dynamic condition was marginal, $F(1, 15) = 3.05$, $p = .101$.

The lack of a Condition × Set size interaction suggests that there was little difference in search efficiency among the three conditions. This concurs with the analysis of RT × Set size slopes, which showed no effect of condition, $F(2, 30) < 1$. Again, in case heterogeneity between the two dynamic conditions was blurring the overall effect of condition, the two dynamic conditions were compared separately to the static condition by planned comparisons; however, neither comparison produced an F value above 1. Remember that for means the memory-driven model predicts that the dynamic:static slope ratios should be 2:1. Actual slope ratios were 1.07:1 for fixed dynamic and 0.95:1 for free dynamic.

RT × Set size intercepts, however, produced a marginally significant, $F(2, 30) = 3.20$, $p = .055$, effect of condition. Planned comparisons showed that both fixed dynamic, $F(1, 15) = 5.23$, $p < .05$, and free dynamic, $F(1, 15) = 4.32$, $p < .05$, had higher intercepts than the static condition. The two dynamic intercepts did not differ, $F(1, 15) < 1$.

Analysis of the medians (Figure 4B) revealed only a main effect of set size, $F(6, 90) = 73.24$, $p < .0001$. The Condition × Set size effect was marginal, $F(12, 180) = 1.691$, $p = .072$; however, three separate ANOVAs comparing fixed and static, free and static, and free and fixed median RTs, respectively, failed to show any significant effect of condition or any Condition × Set size interaction. An ANOVA performed on the median slopes did not show any difference among the three conditions, $F(2, 30) = 1.89$, $p > .1$. Observed slope ratios were 1.08:1 and 0.92:1, for the fixed and free dynamic conditions, respectively, compared to the memory-driven prediction of 1.39:1.

Interestingly, the median intercepts also did not differ by condition, $F(2, 30) = 2.27$, $p > .1$, suggesting that the intercept difference on the mean RTs may have been due to the presence of a few large RTs in the dynamic conditions.

Error rates (Figure 4C) increased with set size, $F(6, 90) = 21.98$, $p < .0001$, and varied significantly by condition, $F(2, 30) = 13.60$, $p < .0001$, subject to an interaction, $F(12, 180) = 4.14$, $p < .0001$. The interaction arises because errors are increasing more rapidly with set size for the two dynamic conditions than for the static condition.

Corrected median RTs (median RT/accuracy) are plotted in Figure 4D. The general pattern of results changes little compared with the uncorrected data. The ANOVA, however, indicates significant effects not only of set size, $F(6, 90) = 92.08$, $p < .0001$, but also of condition, $F(2, 30) = 4.39$, $p < .05$, and the Condition × Set size interaction, $F(12, 180) = 2.00$, $p < .05$. To explore this interaction, we turn to an analysis of the effect of condition on corrected RT × Set size slopes. The overall effect of condition is marginal, $F(2, 30) = 2.87$, $p =$

.072. Paired comparisons indicate that the fixed dynamic slope is steeper than the static slope, $F(1, 15) = 5.44$, $p < .05$; ratio 1.20, whereas the free dynamic slope did not significantly differ from the static slope, $F(1, 15) < 1$; ratio 1.05. The difference between free and fixed dynamic slopes was marginal, $F(1, 15) = 3.26$, $p = .091$. Intercepts computed on the corrected medians showed no effect of condition, $F(2, 30) < 1$, again suggesting that the difference among uncorrected mean intercepts is due to the presence of outliers in the dynamic distributions.

Discussion

The results of this experiment essentially replicate our earlier results with an improved method. Dynamic search does not produce the dramatically steeper slopes predicted by standard memory-driven accounts of search. In fact, there were no reliable differences in the slope data. The methodological changes in the present experiment allow several new issues to be addressed. In these data, the RT $\times$ Set size functions for the dynamic conditions do not diverge from the static function at large set sizes. This conclusion does not depend on whether stimuli are shuffled about in a fixed set of locations from frame to frame or replotted at random new locations on each frame. Finally, our prior results were not an artefact of cutting off the stimulus displays at 2 s.

The similarity between conditions is underscored if we look at the overall distribution of RTs for both correct and error trials, shown pooled across subjects and set sizes in Figure 5. Remember that if search were memory-driven, then the static condition should result in a fairly rectangular distribution, whereas the dynamic conditions should produce exponential distributions. In fact, the distributions for all three conditions are basically identical exponential distributions.

However, intercept differences between dynamic and static search remain. We had hypothesized that extending the frame duration from 100 to 500 ms might eliminate the intercept differences between dynamic and static search. Any or all of several factors might have been ameliorated by the use of the slower rate of change. Slower change could have discouraged subjects from waiting to see if the target reappeared. It could have reduced the masking caused by 10 Hz offsets and onsets. Finally, the slower rate would reduce the chance that a stimulus would disappear while being processed. Apparently, none of these factors is the cause of the intercept difference. The 138 ms intercept difference found in this experiment is comparable to that in our previous studies. Since condition stubbornly refused to interact with set size, the intercept difference must be in some non-search process we have yet to identify.

Differences in error rates between the dynamic and static conditions also remain in this experiment. Since the error rates begin to diverge above 24 items, is it possible that we are observing the same effect of large set sizes that

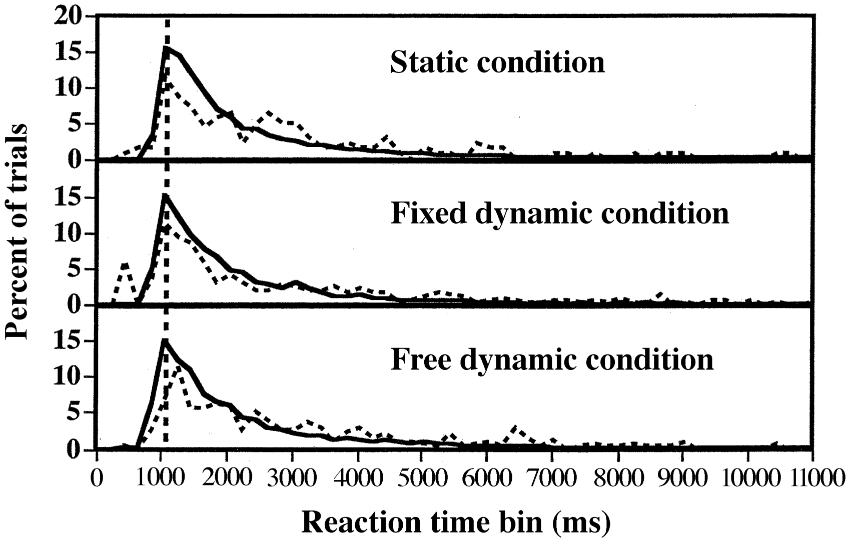

Figure 5. Histogram of RTs from the three conditions of Experiment 1, collapsed across subjects and set sizes. Solid lines indicate correct RTs, and dotted lines the RTs on error trials.

Kristjánsson found, only pushed into the error rates instead? In other words, an obvious possibility is that we are observing a simple speed–error tradeoff here. According to this hypothesis, the equivalence in RT × Set size slopes between the dynamic and static conditions is achieved at the cost of higher error rates for the dynamic conditions, and if we were able to force errors to be equivalent, then the dynamic slopes would increase towards the predictions of the memory-driven model. We think this is unlikely.

First, as noted, the static and dynamic distributions of RTs should be qualitatively different. Figure 5 shows that they are not. Second, when RT and error data are combined into a single measure (Townsend & Ashby, 1978, 1983), the findings were substantially the same. The free dynamic slope did not differ from the static slope. Although fixed dynamic slope was significantly steeper than the static slope, the difference was fairly small, compared to what we would expect, based on the memory-driven model. As noted previously, we should expect a slope ratio of ≈ 1.39 between the dynamic and static conditions if search were memory driven. The actual ratios are 1.05 for free:static and 1.20 for fixed:static, both of which are significantly different from 1.39, $t(15) = 12.21$, 8.07, respectively, $p < .0001$ for both. The differences in error rate point to an unavoidable feature of the randomized search paradigm. The dynamic stimuli are more degraded than the static stimuli. At the very least, the static target will remain in place until the observer responds. The dynamic target, however, is likely to disappear before response, making errors more likely (''I know I found it. Now, was that an N or an E? I forget.'') Even with the unavoidable differences between conditions, the results of the randomized search experiments argue against the hypothesis that standard, static search is memory driven.

EXPERIMENT 2: DOES MEMORY FOR PREVIOUS EYE MOVEMENTS IMPROVE SEARCH EFFICIENCY?

Taken together with data from other lines of research (reviewed in the General Discussion), the data from the randomized search paradigm argues that observers are not marking rejected distractors in visual search. At least, distractors are not marked in any manner that can be used to prevent sampling with replacement. This seems unfortunate. It seems like such a good idea to have a way to prevent re-examination of rejected distractors in visual search. Perhaps such a mechanism does exist, acting on overt eye movements rather than covert attentional deployments.

There are a number of reasons to think that this might be the case. First, the leading candidate for a mechanism to prevent re-examination of rejected distractors is inhibition of return (IOR; Klein, 1988; Klein & MacInnes, 1999; Taylor & Klein, 1998). In typical IOR tasks, attention is directed to a location and then redirected away from that location. Subjects are then asked to respond to a stimulus that is presented either at the previously attended location or elsewhere. The basic finding is that processing is impaired (RTs longer, accuracy reduced) at the previously attended location, as if it were harder to get attention back to that location (Klein, 2000; Posner, 1980; Posner & Cohen, 1984; reviewed in Taylor & Klein, 1998).

Is IOR operating during visual search tasks? The matter has been controversial (Klein, 1988, 1994, p. 466; Wolfe & Pokorny, 1990) but it now appears that IOR can be found so long as the search stimuli remain visible (Müller & von Mühlenen, 2000; Takeda & Yagi, 2000). Moreover, data indicate that IOR can be found not only at the most recently attended locus but at several previous loci (Danziger, Kingstone, & Snyder, 1998; Snyder & Kingstone, 2000; Wright, 1994; Wright & Richard, 1996).

Showing that IOR is present in visual search is not the same as showing that it has a functional role in visual search. The randomized search experiments, described previously, are hard to explain if IOR is effectively preventing redeployment of covert attention. Moreover, the experiments that show IOR for multiple locations still show it for only three–six locations. If IOR was to permit sampling without replacement in visual search, it would have to be possible to inhibit all rejected distractors even for set sizes much greater than six. Finally, the time course of IOR is slow compared to the inferred time course of covert attentional deployments in visual search. If attention is being deployed to one object at a time, then the visual search data suggest that those deployments occur at a rate of 20–40 per s. IOR, by contrast, takes 200 or more ms to develop (but see Danziger & Kingstone, 1999; Posner, Rafal, Choate, & Vaughan, 1985; Taylor & Klein, 1998).

This slower rate is reminiscent of the rate of voluntary eye movements. IOR exists for eye movements (e.g., Briand, Larrison, & Sereno, 2000). It has been argued that IOR is a product of oculomotor programming and, perhaps ‘‘reflects a motor response bias’’ (Taylor & Klein, 1998). At least, there appears to be a substantial oculomotor contribution to IOR (Kingstone & Pratt, 1999). Klein and MacInnes (1999) demonstrated that IOR for eye movements could serve as a ‘‘foraging facilitator’’ in visual search, biasing eye movements toward unexamined portions of the field.

Given this as background, it seemed possible that inhibitory tagging for eye movements might be found in the static condition of a randomized search experiment even if inhibitory tagging for covert attention was not found in the previous experiments. IOR for the past three–six saccades would be sufficient for it to serve as a foraging facilitator in visual search. To provide evidence for such facilitation, we conducted a version of the randomized search experiment that required observers to fixate individual items in order to determine if they were targets or distractors.

Method

Subjects. Sixteen subjects from our volunteer subject pool served as subjects in return for compensation of $10/hour. All subjects had normal or corrected to normal visual acuity and passed the Ishihara colour screen. All gave informed consent prior to their participation.

Stimuli. These were lower-case trigrams presented in 9-point Palatino font. Targets were ‘‘bab’’ and ‘‘hoh’’. Distractors were drawn from the set {‘‘beb’’, ‘‘bib’’, ‘‘bob’’, ‘‘bub’’, ‘‘hah’’, ‘‘heh’’, ‘‘hih’’, ‘‘huh’’, ‘‘kak’’, ‘‘kek’’, ‘‘kik’’, ‘‘kok’’, ‘‘kuk’’, ‘‘dad’’, ‘‘ded’’, ‘‘did’’, ‘‘dod’’, ‘‘dud’’}.

Procedure. On each trial, one of the two target trigrams was present. Subjects were asked to press the ‘‘a’’ key if they saw ‘‘ bab’’, and the quote key if they saw ‘‘ hoh’’; subjects were asked to respond as quickly and accurately as possible. There were 40 possible stimulus locations, arranged on four concentric circles corresponding to the inner four circles from Experiment 1. Viewing distance was again 57 cm. In the static condition, 9, 12, or 15 trigrams were presented until response. In the dynamic condition, the trigrams were randomly shuffled among the same set of locations (as in the fixed dynamic condition of Experiment 1) every 444.5 ms, until a response was made. In order to defeat the ‘‘sit-and-wait’’ strategy, the target trigram was only presented at a randomly chosen sub-set of 1/3 of the selected locations for that trial.

Results

As before, we analysed three measures of correct RT: Means, medians, and corrected medians (Townsend & Ashby, 1983). We also analysed the RT $\times$ Set size slopes based on each of these measures. Finally, we analysed the error rates.

Mean RT, plotted in Figure 6, increased with set size, $F(2,30) = 18.45$, $p < .0001$, and was 1134 ms longer for dynamic trials, $F(1,15) = 17.23$, $p < .001$. However, these two factors did not interact, $F(2,30) = 1.93$, n.s. Slopes (see Table 3) computed on mean RT did not vary by condition, $F(1,15) = 1.58$, n.s., nor did intercepts, $F(1,15) = 1.19$, n.s.

Similar results were obtained on median RT. There was a trend towards longer median RTs in the dynamic condition, though it did not reach significance, $F(1,15) = 3.46$, $p = .083$, and a main effect of set size, $F(1,15) =$

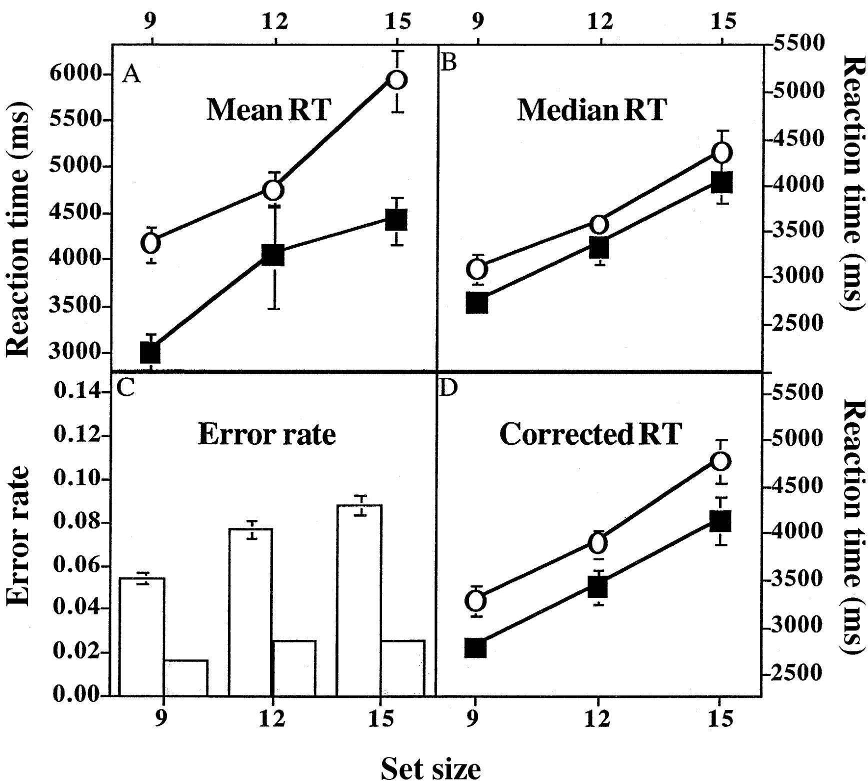

Figure 6. Results from Experiment 2. Mean RTs are shown in panel A, median RTs in panel B, error rates in panel C, and corrected median RTs in panel D. In panels A, B, and D, solid squares denote the static condition, and open circles the dynamic condition. In panel C, open bars represent data from the dynamic condition, and stippled bars the static condition.

TABLE 3
Slopes and intercepts of the RT $\times$ Set size functions from Experiment 2

Based on	*Measure*	*Condition*	
		Fixed	*Static*
Means	Slope	294.31	237.74
	Intercept	1410.56	955.35
Medians	Slope	211.06	217.88
	Intercept	1141.75	747.48
Corrected medians	Slope	251.58	227.19
	Intercept	956.00	716.85

67.06, $p < .0001$, again with no interaction, $F(2, 30) < 1$. Consistent with the median RT analysis, neither slope nor intercept of the RT $\times$ Set size function varied with condition, both $F(1, 15) < 1$.

Correcting median RTs for error does not change the outcome. The corrected median RTs do show a significant effect of condition, $F(1, 15) = 11.48$, $p < .005$, as well as set size, $F(2, 30) = 79.37$, $p < .0001$, but no interaction, $F(2, 30) < 1$, and again there was no effect on either slope or intercept, both $F(1, 15) < 1$.

Subjects were more error-prone in the dynamic condition, $F(1, 15) = 12.75$, $p < .005$, and as set size increased, $F(2, 30) = 8.48$, $p < .005$. As with the RTs, there was no interaction, $F(2, 30) = 2.19$, n.s.

Discussion

The results from Experiment 2 do not support the hypothesis that eye movements are sampling these displays without replacement. Were that the case, the slope of the RT $\times$ Set size function should have been markedly steeper in the dynamic condition than in the static condition, which it is not. In particular, the dynamic slope should be double the static slope when calculated using means, and 1.38 times the static slope when medians are used. The actual ratios are 1.24:1 for means and 0.96:1 for medians. It may be argued that dramatic performance differences are not observed because we focus our analysis on the slope of the RT $\times$ Set size function, and that the substantially higher error rate observed in the dynamic condition is a cue that different processes or strategies are in fact in operation in the two conditions (Shore & Klein, 2000). However, when we combine RT and error into a single measure, the results are substantially the same. Aside from an increase in mean corrected RT, the dynamic and static conditions produce the same results, and the slope ratio is 1.11:1.

Some insight can be obtained from the overall distribution of RTs from this experiment. These are shown in Figure 7.

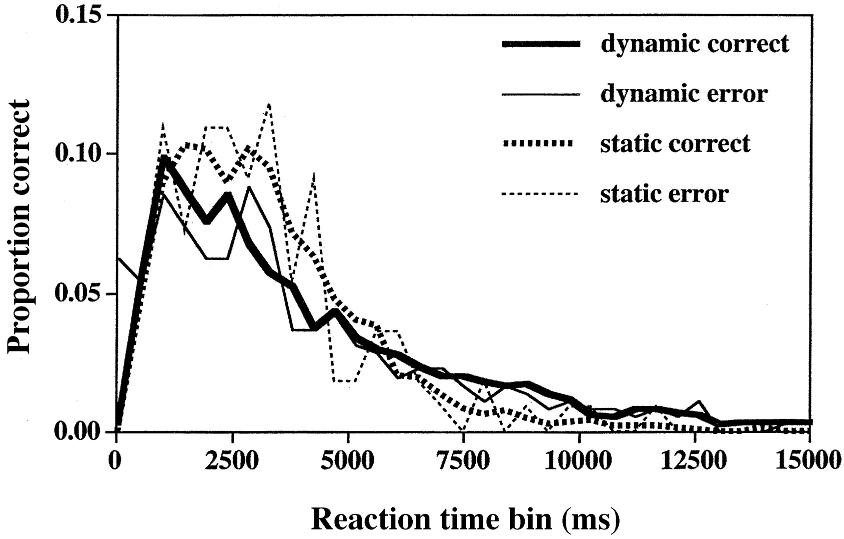

Figure 7. Histogram of RTs from the two conditions of Experiment 2, collapsed across subjects and set sizes. Thick lines indicate correct RTs, and thin lines the RTs on error trials. Solid lines represent RTs from the dynamic condition, dotted lines the static condition.

As in Experiment 1, the distributions are quite similar, and exponential in shape. If search were memory driven in this case, we should expect the static distribution to have a rectangular shape (presumably convolved with an ex-gaussian or similar distribution used to model RT in general; see Van Zandt & Ratcliff, 1995), noticeably lighter in the tail than the dynamic distribution. This is not what we observe.

Although the randomization manipulation has no noticeable effect on the efficiency of search, the intercept offset problem remains. The intercept difference was not statistically significant. However, a glance at Figure 6 indicates that the magnitude of the effect (250 ms in the corrected median data, larger in the other measurers) was substantial. It seems to be due to the presence of very long RTs in the dynamic condition that are not present in the static condition, which also accounts for the increased skew of the dynamic distribution.

It seems safe to conclude that these data reject the standard model that proposes that visual search can be characterized as sampling without replacement. However, it would be rash to claim that the dynamic and static conditions were identical in Experiment 2. The dynamic condition produces long RTs that are absent in the static case. We suspect that this reflects the existence of different quitting criteria in the two conditions. In the dynamic condition, subjects might have been willing to believe that it would take 20 s to track down the target. In the static condition, they might have been unwilling to believe this and might have terminated static searches more readily than dynamic. For this story to be correct, the errors in the static case would need to represent early search terminations, whereas the errors in the dynamic case would need to arise from some other cause such as the reduction in information caused by the disappearance of items. At present, this is merely hypothetical as these data cannot directly address this question. These very long RTs do raise questions about the common practice of removing the longest RTs from analysis in visual search experiments. At the very least, if search is amnesic, this practice could lead to overestimation of the efficiency of a search.

GENERAL DISCUSSION

The randomized search paradigm (Horowitz & Wolfe, 1998) was designed to directly test the common assumption that visual search is memory driven. We hypothesized that if search behaved as a serial processor utilizing sampling without replacement, then randomly replotting stimuli during a search trial should double the slopes of mean RT $\times$ Set size functions. The three experiments initially published in Horowitz and Wolfe (1998) disconfirmed this hypothesis. In Experiment 1 of the present paper, we returned to this question. The original paradigm was modified to take into account criticisms of the original work. We extended the range of set sizes, allowed displays to continue until response, added a same-location randomization condition, reduced the frame rate, and incorporated error data into the analysis of RT. Again, we found no difference in the mean RT $\times$ Set size slopes, and no evidence for other changes in the RT distribution that would be expected if the memory-driven search hypothesis were correct (e.g., differences in the slopes based on median RTs). Even when RT and error data were combined into a single measure, we found only a small increase in the slope of the dynamic condition, far smaller than predicted by the memory-driven search hypothesis. Experiment 2 again disconfirmed the predictions of the memory-driven search hypothesis, this time using stimuli that required fixations, suggesting that visual search may be amnesic about the overt deployments of the eyes as well as about the covert deployments of attention.

Conflicting data

Although our experiments using the randomized search method have produced data that contradict the memory-driven models of search, other laboratories have conducted experiments that sometimes produce different results. Experiment 1 was designed to address two important issues raised by Kristjánsson (2000).

Fixed locations. There were two dynamic search conditions in Experiment 1. In one, items were randomly replotted on every frame. In the other, items were randomly shuffled among a fixed set of locations. We found that the fixed location condition was slightly slower but that both dynamic conditions produced slopes that did not differ from the static condition. This is at variance with Kristjánsson's (2000) results. Kristjánsson replicated the first experiment of Horowitz and Wolfe (1998), with two modifications. In the static condition, stimuli changed orientation on each frame, but did not move. In the ''relocation'' condition, on each frame, the stimuli were shuffled among a fixed set of locations, as in our fixed condition of Experiment 1. He found that the (mean) RT $\times$ Set size slope in the relocation condition was substantially elevated compared to the static slope (37 ms/item versus 11 ms/item).[4] He concluded that the visual system does use inhibitory tagging to achieve sampling without replacement from the search array, and that the results of Horowitz and Wolfe (1998) may have been due to our subjects adopting an alternative to the serial self-terminating search strategy; he suggested serial exhaustive search (Sternberg, 1969) or parallel search (Eckstein, 1998; Palmer, 1995; Palmer & McLean, 1995).

In Experiment 1 of the present paper, replotting items at fixed locations did not produce an increase in the mean or median slopes, and only a modest increase in the slope of corrected RT. A critical difference between Kristjánsson's version of this experiment and ours is that his frame rate was 10 Hz (as in Horowitz & Wolfe, 1998) whereas the rate in our Experiment 1 was slowed to 2 Hz. When stimulus locations are fixed, a 10 Hz randomized search paradigm becomes a rapid serial visual presentation (RSVP) experiment, except instead of a single RSVP stream at fixation, Kristjánsson's experiment has multiple streams. This arrangement could produce masking and/or attentional blink effects (Brehaut, Enns, & Di Lollo, 1999; Breitmeyer, Ehrenstein, Pritchard, Hiscock, & Crisan, 1999; Chun & Potter, 1995; Giesbrecht & Di Lollo, 1998; Grandison, Ghirardelli, & Egeth, 1997; Raymond, Shapiro, & Arnell, 1992) that might degrade search performance. Although this would apply to both static and relocation conditions, the consequences are very different in the two conditions. A degraded target in the static condition might lead to a slower RT, but there would be no interaction with set size, because the target need not be re-acquired. However, if target recognition is significantly delayed or disrupted in the relocation condition, the target will be gone by the time the system recovers, requiring an additional search on some trials, which would lead to increased slopes. The slower, 2 Hz rate of our Experiment 1 would be less subject to this problem.

Extended set sizes. Kristjánsson's (2000) other important innovation was to increase the set size from 4 to 56 items. His data replicated our original results over the range of 4 to 16 items. However, from 24 to 56 items, RTs increased much faster with set size in the dynamic (or ''relocation'') condition than in the static condition. He explained this finding as due to the fact that, given a fixed

[4] Anne Treisman (personal communication) has reported similar data.

set of potential stimulus locations for all set sizes, a target is more likely to be placed in a previously inhibited location as set size increases. In other words, as set size increases, a free location condition more closely resembles a fixed location condition. In our Experiment 1, however, we observe no slope difference between the two dynamic conditions and the static condition, either at low set sizes or high set sizes. We believe the RSVP explanation may account for this discrepancy as well. Instead of increasing the probablity of a target being placed at an inhibited location, the probability is increased that the target will appear at a location occupied by a distractor on the previous frame. At the faster rates used by Kristjánsson, the possibilities for masking and/or attentional blink effects increase with set size in his experiment.

In our data (see Table 4), we do not see a significant increase in the dynamic slopes with set size on any measure. The trend, if anything, is for a slight increase in the static slope when larger set sizes are presented. The same search process seems to be operating whether the array has 9 or 60 items. However, it is possible that we might observe different results for very small set sizes. The increase in the slope with set size would be predicted if a small amount of memory permitted sampling without replacement for small set sizes but not for larger set sizes. This effect should be more prominent for the static condition, since randomization could disrupt even a small memory. Thus, examining the shape of the RT $\times$ Set size function for smaller set sizes (say, one to nine items) might be useful in gathering evidence for a limited capacity memory in search.

Converging evidence

While the randomized search experiments pose a strong challenge to memory-driven accounts of visual search, the method will always be imperfect because the dynamic and static conditions can never be perfectly equated. An attended

TABLE 4
Mean slopes (ms/item) from Experiment 1 computed on the lowest or highest set sizes

Based on	*Set size range*	*Condition*		
		Fixed	*Free*	*Static*
Means	9–18	54.15	52.74	48.42
	36–60	43.82	43.86	52.67
Medians	9–18	35.38	37.86	31.92
	36–60	29.35	26.55	39.67
Corrected medians	9–18	36.76	41.34	33.06
	36–60	37.45	38.91	42.72

item in the dynamic case can suddenly vanish. This will never happen in the static case. This inevitable degradation of the stimulus in the dynamic case makes it likely that error rates will be higher in dynamic search than in static search. The persistent difference in mean RT (and intercept) between the two conditions points to another insurmountable difference between dynamic and static search—perhaps a difference in the tolerance for the very long RTs that can be produced by amnesic search. Given these limitations, it is important to have converging evidence for the argument that search mechanisms do not keep track of rejected distractors. Here we review two such lines of evidence.

Multiple-target search. In the standard search paradigm, with one target per trial (on target-present trials), the only difference between the amnesic and memory-driven models in terms of the mean RTs is a scaling factor. If we observe, for instance, an RT $\times$ Set size slope of 25 ms/item on target present trials, we do not know if this is the result of a memory-driven system working at 50 ms/item and searching an average of $(N+1)/2$ items, or an amnesic system working at 25 ms/item searching an average of N items.

With multiple targets, however, the difference between the two models becomes clearer. As you increase the number of targets (holding the total number of items constant), any model must predict that RTs to the first target become faster, and this fact has been amply demonstrated in the literature (Grice & Canham, 1990; Mordkoff & Yantis, 1991; Mullin, Egeth, & Mordkoff, 1988; Ward & McClelland, 1989). However, what about RTs to subsequent targets? Unsurprisingly, the time to find a target is a function of the ratio of distractors to targets. The more targets there are, the faster you will find a target. The more distractors there are, the slower you will be to find a target. In a multiple target search, after you find each subsequent target, there are fewer targets left in the display, so RTs should slow down. However, in memory-driven search, there are also proportionately fewer distractors, because some fraction of the distractors were marked off in the course of the search for the first target. As a result, the distractor:target ratio remains roughly constant, and RT increases linearly with each successive target. In amnesic search, however, the number of distractors remains constant, because no distractors are ever marked off. The distractor:target ratio increases, and the additional time required to find each additional target increases with each successive target. The function relating RT to number of targets would accelerate rather than increasing linearly.

In Horowitz and Wolfe (2001), we set out to test whether this function would be linear, as predicted by the memory-driven model, or accelerating, as predicted by the amnesic model. We presented subjects with displays of 16 items containing from one to five targets. Targets were digits, and distractors were letters. In one experiment, all targets in a display were unique, in another experiment, they were identical. This made little difference to the conclusions. We wanted to find out how long subjects would take to find each target in a multiple target display. However, we did not want subjects to have to make a motor response to each target when they found it, because it seemed likely that the intervals between target detections would probably be smaller than the resolution of the motor system. Instead, we asked subjects to report (yes or no) whether there were at least *n* targets in the display, where *n* could vary from 2 to 5. Thus, in a display with five targets, for instance, if *n* were 2, they had to respond as soon as they found the second target, and if *n* were 3 they had to respond as soon as they found the third target, and so on. If *n* was 3 and the actual number of targets in the display was 2, that would be an example of a trial requiring a negative response. *N* was varied across blocks, and by combining data from different blocks, we could determine, for a display with five targets in it, how long it took subjects to find the second target, the third target, and so on.

We found that the RT $\times$ *n* functions were clearly quite accelerated, rejecting the memory-driven model. Note that these stimuli are entirely static and not subject to the objections that could be raised to randomized search. Note also that, in order to do this task at all, subjects must have had memory for the location of the targets (Gibson et al., 2000). It is memory for the distractors that is lacking.

Attentional reaction time methods. So far, we have described the difference between amnesic and memory-driven models of search in terms of sampling regimes, with and without replacement, respectively. Another way of framing the difference is in terms of the function describing the momentary probability of detecting the target over time. For memory-driven models, this probability must increase over time, whereas for amnesic models, it must remain constant. As a result, the cumulative probability of finding the target in a memory-driven model increases linearly to 1.0 in *N* steps (where *N* is the set size), whereas for an amnesic model, it will increase as a negative exponential and will never quite reach 1.0 in a finite number of steps (hence the possibility of very long RTs).

We have employed two methods, both inspired by the attentional reaction time (ART) paradigm developed by Reeves and Sperling (1986; Sperling & Weichselgartner, 1995) to test whether the momentary probability of detecting the target is constant or increasing with time. Standard RT measures the motor output that includes the time to find the target, but also includes other components like the time to move the finger. The principle behind the ART is to infer the moment at which a subject actually detected a target, rather than (or in addition to) recording the time they pressed a key. In the original experiments of Reeves and Sperling (1986), subjects were presented with two time-varying stimulus streams, a peripheral stream of letters, and a central stream of digits. After detecting a target letter in the peripheral stream, subjects had to shift their attention to report the next four digits in the central stream. Since the experimenter knows when the target letter was shown and at what time each digit was presented, he can infer how long it took for subjects to shift their attention from one stream to the other.

In our first ART experiment, we combined this logic with the randomized search procedure we developed in Horowitz and Wolfe (1998). The search array itself became the time-varying stream. Subjects were asked to search for a consonant among vowel distractors. On each frame, targets and distractors were randomly replotted at new locations. Frames were changed every 100 ms. On target-present trials, there was a different consonant on each frame. In the critical condition, subjects were asked to report the consonant that they found. Since a different consonant was associated with each frame, the reported target identified the frame on which the target was found. If the probability of *detecting* the target was identical on each frame, then the probability of *reporting* the target would decrease in a purely exponential fashion. However, if the probability of detecting the target increased over time, then the function would deviate from the pure exponential. In two versions of the experiment, we found pure exponential decay—consistent with the predictions of amnesic search.

In a simplified version of this experiment, we changed the target (and distractors) on only one frame. In this experiment, subjects were asked to look for a mirror-reversed letter among normal letters. The letters were coloured, and at some point during the trial (the transition time), all the letters changed colour simultaneously. The subjects' task was simply to report the colour of the target. By varying the transition time, the cumulative probability of reporting the target by a given time can be computed. For memory-driven models, this function should be linear, with a slope determined by the hypothetical dwell time. For an amnesic model, however, the function should be exponential. As in the prior experiment, we observe the exponential result consistent with amnesic search.

How amnesic is visual search?

The consistent results of several experiments from several paradigms appear to falsify the standard, serial, self-terminating account of visual search because the data are not consistent with sampling without replacement. As a working hypothesis, we propose that visual search is amnesic and samples the stimuli with replacement and without *any* memory for the previous course of search. The full-memory and amnesic models represent two extremes on a continuum of possible models. Largely ignored in the mainstream cognitive literature on visual search have been several studies in the human factors literature which have looked at the question of whether searches that require saccades are memory-driven ''systematic'' or ''random'' (memory-free). A number of studies concluded in favour of a random model of search (Bloomfield, 1972; Engel, 1977) on the basis of search time distributions, though there is also some evidence for systematic search (Gupta & Geyer, 1981). Arani et al. (1984) proposed a model of search in which the amount of memory was an explicit parameter. Memory for prior fixations was assumed to degrade as a power function, governed by two free parameters: ϕ, the (constant) probability of recalling an item on the i^{th} fixation; and θ, the probability of encoding the item in the first place. The probability of recalling on the i^{th} fixation given the fact that the area was previously fixated on the k^{th} fixation is then given by:

$$P_{i,k} = \theta\phi^{i-k} \tag{5}$$

If either of the two parameters is set to 0, the model becomes memory-free. If they are both 1, then the model exhibits perfect memory-driven behaviour. Arani et al. (1984) suggest that a mixed model more accurately reflects actual search behaviour, and propose that θ and ϕ can be measured directly in eye-movement studies, though this experiment does not appear to have been done. The mixed model also receives theoretical support from the work of Courtney and Guan (1996), who proposed that the ideal eye movement strategy in search is to have some modicum of memory which allows random fixations with some fixed overlap in the useful field of view (UFOV; Ball, Beard, Roenker, Miller, & Griggs, 1988; Ball, Owsley, Sloane, Roenker, & Bruni, 1992) afforded by successive fixations. Such a model is not incompatible with the evidence provided by Klein and MacInnes (1999), who simply demonstrated an inertia-like bias against saccades which reversed the direction of the last saccade. Such biases may serve to reduce the overlap between visual lobes, without necessarily requiring an inhibitory map of the entire search array. Furthermore, the sort of mixed or minimal-memory model proposed by Arani et al. (1984) might not produce dramatic changes in the RT $\times$ Set size slopes in the experiments reported here.

Thus, while our randomized search experiments appear to falsify models that propose perfect or near-perfect memory for the course of visual search, they do not falsify models that propose some modest amount of memory. As an example, consider the experiments that attempt to measure the number of previously attended items that can be inhibited in an IOR paradigm. The most optimistic of these propose IOR for three–six previous items. Although that would produce memory-driven search for very small set sizes, it would provide proportionally less memory as set size increased (e.g., to 60 items in our Experiment 1). These ''*n*-back'' models of search are worth further investigation. The randomized search technique is not well suited to looking for limited memory of this sort.

Other approaches also suggest that any memory for rejected distractors in visual search is fairly minimal. For example, Gilchrist and Harvey (2000) recorded eye movements during a visual search task on three observers. They then calculated the a priori probabilities of returning to a previously fixated item based on either sampling with replacement or sampling without replacement. Using the observed frequency of refixating, they were able to calculate a memory index, ranging from 0 (no memory) to 1 (perfect memory). Values of the index ranged from 0.12 to 0.60 (mean = 0.40). However, as the authors point out, the assumptions underlying the computations probably inflated the memory index estimates, suggesting that the true values are probably even lower (see also Findlay, Brown, & Gilchrist, 2001).

Is amnesic search plausible?

Part of the attraction of the standard model of search is that it is so reasonable. It seems obvious that search should proceed without replacement. Why reexamine an item that you have already rejected? By contrast, the amnesic account can seem far-fetched. In addition to the seemingly undesirable matter of revisiting rejected distractors, amnesic search would seem to be subject to severe problems with perseveration. If you do not keep track of rejected distractors, how do you avoid continually revisiting particularly salient distractors? Moreover, as discussed later, an amnesic model requires very fast attentional deployments in order to explain the RT $\times$ Set size slopes observed in typical search tasks.

The perseveration problem. The first objection we refer to as the ‘‘perseveration problem’’. Several models of visual search (e.g., Treisman & Sato, 1990; Wolfe, 1994) propose that visual search is conducted in two stages. In the first stage, information from various feature maps is combined in parallel to produce a salience map (Koch & Ullman, 1984, 1985) indicating the distribution of likely targets in the visual field. Attention is then deployed to items in order of their salience. Deploying in order of salience implies a memory driven search. The salience is determined by bottom-up factors (luminance, sharp discontinuities, etc.), top-down factors (match to the preattentive feature of the designated target), and a certain amount of noise. What happens when a distractor happens to have a higher salience than the target? If we assume sampling without replacement, the salient distractor is attended to first, marked off, and then attention moves to the next most salient item, eventually reaching the target. In a sampling with replacement regime, however, if the system simply directs attention to the most salient item at that time, attention could be drawn to the salient distractor, like a moth to a flame. This problem would be exacerbated for peripheral targets, given that foveal items seem to be intrinsically more salient (Wolfe et al., 1998).

The perseveration problem is primarily a problem in the cases where no preattentive information differentiates targets and distractors. In these cases, dynamic noise is an answer to this problem. If the salience of items fluctuates over time, and if all items are otherwise equal, then no one item will arbitrarily attract attention. In the cases where an item is salient because of the presence of a non-target feature, then the attentional ‘‘set’’ of the subject will be adjusted to reduce the effectiveness of that feature. For example, Yantis and Egeth (1999) have shown that salient cues need not capture attention if they are known to be invalid.

We have developed a new implementation of the Guided Search model (Wolfe, 2001) that works without memory. In Guided Search, preattentive information about basic features is used to guide attention to likely targets (Wolfe, 1994). Thus, in a search for a red ‘‘T’’ among red and green ‘‘L’’s, deployment of attention would be restricted to red items, even without memory for the previous deployments of attention. Guided Search generates a ‘‘salience map’’ in which each item is assigned an activation level depending on how different it is from its neighbours and how many features it shares with the target. This activation level is also noisy, so that within a set of roughly equivalent stimuli (the red letters in this example), attention is deployed at random. Perseveration is avoided by assuming that noise in the activation levels is dynamic. As a consequence, an item that is the most attractive to attention at one moment may be less attractive at the next. It is possible to set a noise level that does not mask the guiding effects of preattentive features or eccentricity (Carrasco et al., 1995; Wolfe et al., 1998) while preventing attention from becoming ‘‘stuck’’ on one or a few distractor items. In this manner, it is possible to accurately simulate human RT data without assuming that rejected distractors are marked in any way.

Rate of deployment. If we discard the assumption of sampling without replacement, we must change our estimate of the rate at which attention can be deployed from item to item in a search. In search tasks, it has long been assumed that the rate is best estimated by the target-absent search slope, because it is assumed that each item must be attended in order to be sure that there is no target in the display. On target-present trials, subjects are assumed to scan, on average, roughly half of the items in the display, as predicted by equation 1.[5] Thus the target-present slope should be half the rate. So twice the target-present slope is the common yardstick for cost of each additional item in a search display. This method yields estimates of 40–60 ms/item for typical spatial configuration searches (Schneider & Shiffrin, 1977; Wolfe, 1998b), believed to be serial in nature. Of course, the ratio of target-absent slopes to target-present slopes is not always 2:1 (Wolfe, 1998b). Furthermore, target-absent RTs seem to be governed by a variety of quitting strategies (Chun & Wolfe, 1996).

Other measures of time that attention spends on a single item range from 10 ms/item (Sperling, Budiansky, Spivak, & Johnson, 1971) to 300–500 ms/item (Duncan, Ward, & Shapiro, 1994; Chun & Potter, 1995; Cobo, Pinilla, & Valdes-Sosa, 1999; Ward, Duncan, & Shapiro, 1996, 1997), depending on the methodology. Although the very fast (10 ms/item) estimates are near the lower limit of plausibility, most of the estimates cluster around the high end, and reconciling these numbers with the much faster estimates from search tasks

[5] The expected mean number of items examined will actually be $(n+1)/2$ rather than $n/2$, as shown in equation 1, but this does not change the slope.

remains a theoretical problem (Moore, Egeth, Berglan, & Luck, 1996; Moore & Wolfe, 2000). In an amnesic search model, the target-present slope itself becomes the measure of the dwell time because the average number of attended items is equal to the set size in an amnesic search. This changes the estimated rate to one item every 20–30 ms/item (somewhat slower if we correct for errors). Such fast rates are often seen as a problem for serial models of search because it seems unlikely that a stimulus can go from light on the retina to identified object in 25 ms. This, however, is a misunderstanding of the meaning of the rate. The fact that search may proceed at 25 ms/item does not imply that it takes only 25 ms to identify the item. A conveyor belt provides a useful metaphor. One could have items going on and off the conveyor every 25 ms but an item might take a second to travel the length of the belt. The rate would still be one every 25 ms if we assume that many items could be on the belt at the same time (Moore & Wolfe, 2000).

Is memory worth the trouble?

Even though it sounds quite reasonable, one can question whether memory for the course of a visual search is worth the investment. The exact route of attention is not likely to be a useful piece of information. Moreover, given that the eyes can move around during the course of a search, the inhibition would need to be maintained in world-centred (or object-centred) coordinates. The time that it takes to generate the memory could offset the savings in deployments compared to amnesic search. It may be more adaptive to save memory for more useful information such as the probable location of undiscovered targets (Chun & Jiang, 1998) or the actual location of targets already found (Gibson et al., 2000).

Keeping track of the course of a search or imposing an order on the course of search is worthwhile in search through complex scenes, particularly when the cost of a missed target is high (Is there a tumour in that X-ray image? Is there a missile launcher in that satellite image?). Under these conditions, observers may adopt a regular ‘‘scan-path’’ in an effort to search systematically (Noton & Stark, 1971). This order comes at a cost. Although covert deployments of attention may proceed at a rate of 20–30 items/s, volitional, controlled deployments occur at a rate nearly an order of magnitude slower (Wolfe, Alvarez, & Horowitz, 2000).

Implications for parallel models of visual search

The arguments thus far have been framed in terms of a serial model of search in which the focus of attention is shifted around the display from one object to the next, one object at a time. Although we think that this is a plausible account (e.g., Woodman & Luck, 1999), there are credible limited- and unlimited-capacity parallel models of search. These generally assume a different sort of memory. In these models, information accumulates for each item, eventually identifying the item as target or distractor. For instance, take the diffusion model proposed by Ward and McClelland (Ratcliff, 1978; Ward & McClelland, 1989). In this model, there is a detector at every location in the visual field, which accumulates information about whether or not the stimulus at that location is a target. This information accumulates in a random-walk fashion, until the certainty of the stimulus being a target (or distractor) reaches some threshold. When either a target signal or enough distractor signals reach the threshold, a response is made. Such a model would be disrupted by the randomization procedure as severely as the serial search without replacement model.

There is one class of model that can accommodate the equivalence of static and dynamic search efficiency while maintaining some accumulation of information over time. Memory-driven search assumptions can be maintained if the information is assumed to be global in nature, rather than location based. That is, the system may slowly accumulate evidence that there is a target present in the field, without knowing *where* the target is. When this evidence reaches some criterion, a response can be issued without the subject having seen any particular instance of the target. Of course, this does not accord with the phenomenology of the task. Even in dynamic search displays, subjects claim to see the target letter in a particular place at a particular time. However, introspection may not be a reliable guide in this case. Fortunately, our attentional reaction time experiments (described earlier) were designed to test exactly this hypothesis. If information is accumulating in a global fashion, then the probability of detecting the target should increase over time, which is not what we observe.

Therefore, our general claim is that search does not rely on building up information about the scene during the course of a search. Instead, search can best be modelled as a series of successive judgements of the momentary or immediate probability that there is a target. Unsuccessful searches are terminated not when all items have been exhaustively examined, but rather when subjects have made a subjective judgement that they have searched ''long enough'' (Chun & Wolfe, 1996).

In some ways, we are proposing major reconceptualization of visual search, abandoning a widely held assumption, and bringing theories of search into line with the ''just-in-time vision'' literature (Giesbrecht & Di Lollo, 1998; Irwin & Gordon, 1998; O'Regan, Rensink, & Clark, 1999; Rensink, 2000; Rensink, O'Regan, & Clark, 2000; Schneider, 1999; Simons, 2000; Simons & Levin, 1998). On the other hand, we do not believe that we are creating a major workload for theorists of visual search. As noted previously, Guided Search can be made to operate successfully without assuming memory for rejected distractors. There is no reason to suppose that other models cannot be similarly modified.

CONCLUSION

The data presented and reviewed here falsify models that propose that subjects keep track of the course of a visual search in a manner that permits that search to be a case of sampling without replacement from the set of search stimuli. It is in this sense that it can be claimed that ''visual search has no memory'' (Horowitz & Wolfe, 1998). Although complete memory was a common assumption, in some ways it should come as no surprise that it is false. No one had ever produced data showing that all items in a search display could be marked during the course of a search. IOR, the most likely candidate mechanism, has only been demonstrated to work for a few items at a time.

It is important to reiterate that this is not a claim that visual *searchers* have no memory. Observers can explicitly remember what they have found and can show implicit effects of one search trial on subsequent search trials. The amnesia we propose is quite specific. The mechanism of search does not remember which items it has previously examined. It is possible that even that claim is too strong and that the search mechanism does recall the last N items where N is some small number. We have on-going experiments that search for this minimal memory in visual search. So far we have not found it, but it remains possible that it exists.

REFERENCES

Alvarez, G.A., Horowitz, T.S., & Wolfe, J.M. (2000). Multielement tracking and visual search use independent resources. *Investigative Ophthalmology and Visual Science*, *41*(4).

Arani, T., Karwan, M.H., & Drury, C.G. (1984). A variable-memory model of search. *Human Factors*, *26*(6), 631–639.

Backer, A., & Preal, E. (1999). *Does visual search have memory? A mathematical framework suggests current experimental paradigm is not conclusive*. Paper presented at the third annual Vision Research conference, Fort Lauderdale, FL.

Ball, K., Owsley, C., Sloane, M.E., Roenker, D.L., & Bruni, J.R. (1992). Visual attention problems as a predictor of vehicle crashes among older drivers. *Investigative Ophthalmology and Visual Science*, *33*(4), 1419.

Ball, K.K., Beard, B.L., Roenker, D.L., Miller, R.L., & Griggs, D.S. (1988). Age and visual search: Expanding the useful field of view. *Journal of the Optical Society of America—A*, *5*(12), 2210–2219.

Bloomfield, J.R. (1972). Visual search in complex fields: Size differences between target disc and surrounding discs. *Human Factors*, *14*(2), 139–148.

Brady, N., & Lawrie, M. (1999). Looking for a face in a crowd. *Perception*, *28*(Suppl.), 482.

Brainard, D.H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 443–446.

Brehaut, J.C., Enns, J.T., & Di Lollo, V. (1999). Visual masking plays two roles in the attentional blink. *Perception and Psychophysics*, *61*(7), 1436–1448.

Breitmeyer, B.G., Ehrenstein, A., Pritchard, K., Hiscock, M., & Crisan, J. (1999). The roles of location specificity and masking mechanisms in the attentional blink. *Perception and Psychophysics*, *61*(5), 798–809.

Briand, K.A., Larrison, A.L., & Sereno, A.B. (2000). Inhibition of return in manual and saccadic response systems. *Perception and Psychophysics*, *62*(8), 1512–1524.

Broadbent, D.E. (1957). A mechanical model for human attention and immediate memory. *Psychological Review*, *64*, 205–215.

Carrasco, M., Evert, D.L., Chang, I., & Katz, S.M. (1995). The eccentricity effect: Target eccentricity affects performance on conjunction searches. *Perception and Psychophysics*, *57*(8), 1241–1261.

Cave, K.R., & Wolfe, J.M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, *22*(2), 225–271.

Chan, A.H., & Courtney, A.J. (1998). Revising and validating the random search model for competitive search. *Perceptual and Motor Skills*, *87*(1), 251–260.

Cherry, E.C. (1953). Some experiments on the recognition of speech, with one and two ears. *Journal of the Acoustical Society of America*, *25*, 975–979.

Chun, M.M., & Jiang, Y. (1998). Contextual cuing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*, 28–71.

Chun, M.M., & Potter, M.C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, *21*(1), 109–127.

Chun, M.M., & Wolfe, J.M. (1996). Just say no: How are visual searches terminated when there is no target present? *Cognitive Psychology*, *30*(1).

Chun, M.M., & Wolfe, J.M. (in press). Visual attention. In E.B. Goldstein (Ed.), *Blackwell's handbook of perception*. Oxford, UK: Blackwell.

Cobo, A., Pinilla, T., & Valdes-Sosa, M. (1999). Attention to surfaces defined by transparent motion: Measuring dwell time. *Brain and Cognition*, *40*(1), 85–90.

Cohen, E., & Ruppin, E. (1999). From parallel to serial processing: A computational study of visual search. *Perception and Psychophysics*, *61*(7), 1449–1461.

Courtney, A.J., & Guan, L. (1996). Visual search performance with complete and incomplete visual lobe areas. *Ergonomics*, *39*(5), 749–756.

Danziger, S., & Kingstone, A. (1999). Unmasking the inhibition of return phenomenon. *Perception and Psychophysics*, *61*(6), 1024–1037.

Danziger, S., Kingstone, A., & Snyder, J.J. (1998). Inhibition of return to successively stimulated locations in a sequential visual search paradigm. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(5), 1467–1475.

Duncan, J., & Humphreys, G. (1992). Beyond the search surface: Visual search and attentional engagement. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(2), 578–588.

Duncan, J., Ward, R., & Shapiro, K.L. (1994). Direct measurement of attentional dwell time in human vision. *Nature*, *369*(6478), 313–315.

Eckstein, M.P. (1998). The lower visual search efficiency for conjunctions is due to noise and not serial attentional processing. *Psychological Science*, *9*(2), 111–118.

Engel, F.L. (1977). Visual conspicuity, visual search and fixation tendencies of the eye. *Vision Research*, *17*, 95–108.

Fairclough, S., & Maternaghan, M. (1993). Changes in drivers' visual behavior due to the introduction of complex versus simple route navigation information. In D.E.G. Brogan & A. Gale (Eds.), *Visual search* (Vol. 2, pp. 419–431). London: Taylor & Francis.

Findlay, J.M., Brown, V., & Gilchrist, I.D. (2001). Saccade target selection in visual search: The effect of information from the previous fixation. *Vision Research*, *41*(1), 87–95.

Gale, A.G., & Walker, G.E. (1993). Visual search in breast cancer screening. In D.E.G. Brogan & R. Carr (Eds.), *Visual search* (Vol. 2, pp. 231–238). London: Taylor & Francis.

Gibson, B.S., Li, L., Skow, E., Salvagni, K., & Cooke, L. (2000). Memory-based tagging of targets during visual search for one versus two identical targets. *Psychological Science*, *11*(4), 324–328.

Giesbrecht, B., & Di Lollo, V. (1998). Beyond the attentional blink: Visual masking by object substitution. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(5), 1454–1466.

Gilchrist, I.D., & Harvey, M. (2000). Refixation frequency and memory mechanisms in visual search. *Current Biology*, *10*(19), 1209–1212.

Grandison, T.D., Ghirardelli, T.G., & Egeth, H.E. (1997). Beyond similarity: Masking of the target is sufficient to cause the attentional blink. *Perception and Psychophysics*, *59*(2), 266–274.

Grice, G.R., & Canham, L. (1990). Redundancy phenomena are affected by response requirments. *Perception and Psychophysics*, *48*(3), 209–213.

Grossberg, S., Mingolla, E., & Ross, W.D. (1994). A neural theory of attentive visual search: Interactions of boundary, surface, spatial, and object representations. *Psychological Review*, *101*(3), 470–489.

Gupta, S.M., & Geyer, L.H. (1981). *Visual search is systematic*. Paper presented at the Human Factors Society 25th annual meeting, Santa Monica, CA.

Helsen, W., & Pauwels, J.M. (1993). A cognitive approach to visual search in sport. In D.E.G. Brogan & A. Gale (Eds.), *Visual search* (Vol. 2, pp. 379–388). London: Taylor & Francis.

Hillstrom, A.P. (2000). Repetition effects in visual search. *Perception and Psychophysics*, *62*(4), 800–817.

Horowitz, T.S., & Wolfe, J.M. (1998). Visual search has no memory. *Nature*, *394*(6693), 575–577.

Horowitz, T.S., & Wolfe, J.M. (2001). Search for multiple targets: Remember the targets, forget the search. *Perception and Psychophysics*, *63*(2), 272–285.

Humphreys, G.W., & Muller, H. (1993). Search via recursive rejection (SERR): A connectionist model of visual search. *Cognitive Psychology*, *25*, 43–110.

Irwin, D.E., & Gordon, R.D. (1998). Eye movements, attention and trans-saccadic memory. *Visual Cognition*, *5*(1–2), 127–155.

Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*(10–12), 1489–1506.

Johnson, N.L., & Kotz, S. (1977). *Urn models and their applications*. New York: John Wiley & Sons.

Kinchla, R.A. (1974). Detecting targets in multi-element arrays: A confusability model. *Perception and Psychophysics*, *15*, 149–158.

Kingsley, H.C. (1932). An experimental study of ''search''. *American Journal of Psychology*, *44*, 314–318.

Kingsley, H.C. (1934). The influence of instruction and context upon perceptual search. *American Journal of Psychology*, *46*, 437–442.

Kingstone, A., & Pratt, J. (1999). Inhibition of return is composed of attentional and oculomotor processes. *Perception and Psychophysics*, *61*(6), 1046–1054.

Klein, R.M. (1988). Inhibitory tagging system facilitates visual search. *Nature*, *334*, 430–431.

Klein, R.M. (2000). Inhibition of return. *Trends in Cognitive Sciences*, *4*(4), 138–147.

Klein, R.M., & MacInnes, W.J. (1999). Inhibition of return is a foraging facilitator in visual search. *Psychological Science*, *10*(4), 346–352.

Koch, C., & Ullman, S. (1984). *Selecting one among the many: A simple network implementing shifts in selective visual attention* (AI Memo 770). Cambridge, MA: Massachussetts Institute of Technology Artificial Intelligence Laboratory.

Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, *4*, 219–227.

Kontsevich, L.L. (2001). *Visual search is certainly serial*. Paper presented at the meeting of the Vision Sciences Society, Sarasota, FL.

Kristjánsson, A. (2000). In search of remembrance: Evidence for memory in visual search. *Psychological Science*, *11*(4), 328–332.

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

Maljkovic, V., & Nakayama, K. (1996). Priming of pop-out: II. Role of position. *Perception and Psychophysics*, *58*(7), 977–991.

McElree, B., & Carrasco, M. (in press). The temporal dynamics of visual search: Evidence for parallel processing in feature and conjunction searches. *Journal of Experimental Psychology: Human Perception and Performance*.

Moore, C.M., Egeth, H., Berglan, L., & Luck, S.J. (1996). Are attentional dwell times inconsistent with serial visual search? *Psychonomic Bulletin and Review*, *3*(3), 360–365.

Moore, C.M., & Wolfe, J.M. (2000). Getting beyond the serial/parallel debate in visual search: A hybrid approach. In K. Shapiro (Ed.), *The limits of attention: Temporal constraints on human information processing*. Oxford, UK: Oxford University Press.

Mordkoff, J.T., & Yantis, S. (1991). An interactive race model of divided attention. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(2), 520–538.

Mullin, P., Egeth, H., & Mordkoff, T. (1988). Redundant-target detection and processing capacity: The problem of positional preferences. *Perception and Psychophysics*, *43*(6), 607–610.

Müller, H., & von Mühlenen, A. (2000). Probing distractor inhibition in visual search: Inhibition of return. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(5), 1591–1605.

Noton, D., & Stark, L. (1971). Eye movements and visual perception. In R. Held & W. Richards (Eds.), *Perception: Mechanisms and models—Readings from Scientific American* (pp. 218–227). San Francisco: W.H. Freeman.

O'Regan, J.K., Rensink, R.A., & Clark, J.J. (1999). Change-blindness as a result of ''mudsplashes''. *Nature*, *398*(6722), 34.

Palmer, J. (1995). Attention in visual search: Distinguishing four causes of a set-size effect. *Current Directions in Psychological Science*, *4*(4), 118–123.

Palmer, J., & McLean, J. (1995). *Imperfect, unlimited-capacity, parallel search yields large set-size effects*. Irvine, CA: Society of Mathematical Psychology.

Pashler, H. (1997). *The psychology of attention*. Cambridge, MA: MIT Press.

Pashler, H., & Badgio, P.C. (1985). Visual attention and stimulus identification. *Journal of Experimental Psychology: Human Perception and Performance*, *11*(2), 105–121.

Peterson, M.S., Kramer, A.F., Wang, R.F., Irwin, D.E., & McCarley, J.S. (2001). Visual search has memory. *Psychol Sci*, *12*(4), 287–292.

Posner, M.I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*, 3–25.

Posner, M.I., & Cohen, Y. (1984). Components of attention. In H. Bouma & D.G. Bouwhuis (Eds.), *Attention and performance X: Control of language processes* (pp. 55–66). Hove, UK: Lawrence Erlbaum Associates Ltd.

Posner, M.I., Rafal, R.D., Choate, L.S., & Vaughan, J. (1985). Inhibition of return: Neural basis and function. *Cognitive Neuropsychology*, *2*(3), 211–228.

Pylyshyn, Z. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition*, *32*(1), 65–97.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.

Raymond, J.E., Shapiro, K.L., & Arnell, K.M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? [see comments]. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(3), 849–860.

Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review*, *93*(2), 180–206.

Rensink, R.A. (2000). Seeing, sensing, and scrutinizing. *Vision Research*, *40*(10–12), 1469–1487.

Rensink, R.A., O'Regan, J.K., & Clark, J.J. (2000). On the failure to detect changes in scenes across brief interruptions. *Visual Cognition*, *7*(1–3), 127–145.

Roberts, S., & Sternberg, S. (1993). The meaning of additive reaction-time effects: Tests of three alternatives. In D.E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 611–653). Cambridge, MA: MIT Press.

Ruckmick, C.A. (1926). On overlooking familiar objects. *American Journal of Psychology*, *37*, 631–632.

Schneider, W., & Shiffrin, R.M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, *84*(1), 1–66.

Schneider, W.X. (1999). Visual-spatial working memory, attention, and scene representation: A neuro-cognitive theory. *Psychological Research*, *62*(2–3), 220–236.

Shore, D.I., & Klein, R.M. (2000). On the manifestations of memory in visual search. *Spatial Vision*, *14*, 59–75.

Simons, D.J. (2000). Current approaches to change blindness. *Visual Cognition*, *7*(1–3), 1–15.

Simons, D.J., & Levin, D.T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin and Review*, *5*(4), 644–649.

Snyder, J., & Kingstone, A. (2000). Inhibition of return and visual search: How many separate loci are inhibited? *Perception and Psychophysics*, *62*(3), 452–458.

Sperling, G., Budiansky, J., Spivak, J.G., & Johnson, M.C. (1971). Extremely rapid visual search: The maximum rate of scanning letters for the presence of a numeral. *Science*, *174*(6), 307–311.

Sperling, G., & Weichselgartner, E. (1995). Episodic theory of the dynamics of spatial attention. *Psychological Review*, *102*(3), 503–532.

Sternberg, S. (1969). High-speed scanning in human memory. *Science*, *153*, 652–654.

Takeda, Y., & Yagi, A. (2000). Inhibitory tagging in visual search can be found if search stimuli remain visible. *Perception and Psychophysics*, *62*(5), 927–934.

Taylor, T., & Klein, R.M. (1998). On the causes and effects of inhibition of return. *Psychonomic Bulletin and Review*, *5*(4), 625–643.

Townsend, J.T., & Ashby, F.G. (1978). Methods of modeling capacity in simple processing systems. In J.N. John Castellan & F. Restle (Eds.), *Cognitive theory* (Vol. 3, pp. 199–239). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Townsend, J.T., & Ashby, F.G. (1983). *The stochastic modeling of elementary psychological processes*. Cambridge, UK: Cambridge University Press.

Treisman, A. (1992). Perceiving and re-perceiving objects. *American Psychologist*, *47*(7), 862–875.

Treisman, A. (1993). The perception of features and objects. In A.D. Baddeley (Ed.), *Attention: Selection, awareness, and control—a tribute to Donald Broadbent* (pp. 5–35). Oxford, UK: Clarendon Press/Oxford University Press.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*, 97–136.

Treisman, A., & Sato, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance*, *16*(3), 459–478.

Treisman, A.M. (1964). Selective attention in man. *British Medical Bulletin*, *20*(1), 12–16.

Van Zandt, T., & Ratcliff, R. (1995). Statistical mimicking of reaction time data: Single process models, parameter variability, and mixtures. *Psychonomic Bulletin and Review*, *2*(1), 2.

Ward, R., Duncan, J., & Shapiro, K. (1996). The slow time-course of visual attention. *Cognitive Psychology*, *30*(1), 79–109.

Ward, R., Duncan, J., & Shapiro, K. (1997). Effects of similarity, difficulty, and nontarget presentation on the time course of visual attention. *Perception and Psychophysics*, *59*(4), 593–600.

Ward, R., & McClelland, J.L. (1989). Conjunctive search for one and two identical targets. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(4), 664–672.

Wolfe, J.M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review*, *1*(2), 202–238.

Wolfe, J.M. (1996). Extending Guided Search: Why Guided Search needs a preattentive ''item map''. In A.F. Kramer & M.G.H. Coles (Eds.), *Converging operations in the study of visual selective attention*. Washington, DC: American Psychological Association.

Wolfe, J.M. (1998a). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13–73). Hove, UK: Psychology Press.

Wolfe, J.M. (1998b). What can 1 million trials tell us about visual search? *Psychological Science*, *9*(1).

Wolfe, J.M. (2001). *Guided Search 4.0: A Guided Search model that does not require memory for rejected distractors*. Paper presented at the meeting of the Vision Sciences Society, Sarasota, FL.

Wolfe, J.M., Alvarez, G.A., & Horowitz, T.S. (2000). Attention is fast but volition is slow. *Nature*, *406*(6797), 691.

Wolfe, J.M., Cave, K.R., & Franzel, S.L. (1989). Guided Search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(3), 419–433.

Wolfe, J.M., O'Neill, P., & Bennett, S.C. (1998). Why are there eccentricity effects in visual search? Visual and attentional hypotheses. *Perception and Psychophysics*, *60*(1).

Wolfe, J.M., & Pokorny, C.W. (1990). Inhibitory tagging in visual search: A failure to replicate. *Perception and Psychophysics*, *48*(4), 357–362.

Woodman, G.F., & Luck, S.J. (1999). Electrophysiological measurement of rapid shifts of attention. *Nature*, *400*, 867–869.

Wright, R.D. (1994). Shifts of visual attention to multiple simultaneous location cues. *Canadian Journal of Experimental Psychology*, *48*(2), 205–217.

Wright, R.D., & Richard, C.M. (1996). Inhibition of return at multiple locations in visual space. *Canadian Journal of Experimental Psychology*, *50*(3), 324–327.

Yantis, S., & Egeth, H.E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *25*(3), 661–676.

Yantis, S., & Johnston, J.C. (1990). On the locus of visual selection: Evidence from focused attention tasks. *Journal of Experimental Psychology: Human Perception and Performance*, *16*(1), 135–149.

Zenger, B., & Fahle, M. (1997). Missed targets are more frequent than false alarms: A model for error rates in visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *23*(6), 1783–1791.

Manuscript received April 2001
Revised manuscript received February 2002