

Eye movements during change detection: Implications for search constraints, memory limitations, and scanning strategies

GREGORY J. ZELINSKY

State University of New York, Stony Brook, New York

Search, memory, and strategy constraints on change detection were analyzed in terms of oculomotor variables. Observers viewed a repeating sequence of three displays (Scene 1 → Mask → Scene 2 → Mask . . .) and indicated the presence-absence of a changing object between Scenes 1 and 2. Scenes depicted real-world objects arranged on a surface. Manipulations included set size (one, three, or nine items) and the orientation of the changing objects (similar or different). Eye movements increased with the number of potentially changing objects in the scene, with this set size effect suggesting a relationship between change detection and search. A preferential fixation analysis determined that memory constraints are better described by the operation comparing the pre- and postchange objects than as a capacity limitation, and a scanpath analysis revealed a change detection strategy relying on the peripheral encoding and comparison of display items. These findings support a signal-in-noise interpretation of change detection in which the signal varies with the similarity of the changing objects and the noise is determined by the distractor objects and scene background.

Building on Neisser's (1967) notion of a *perceptual cycle*, Parker (1978) described a simple four-stage model of visual change detection as it applies to line drawings of scenes. Observers first *extract* object information from across a study scene, then perform a series of *comparison* operations between these represented objects and items appearing in a subsequently presented test scene. The efficacy of each comparison is determined by both the magnitude of the change, if any, and the ability of observers to resolve the change in their peripheral vision. On the basis of the results of this comparison process, *gaze is guided* to the region of the picture most likely to contain the change; then, a *decision* is made as to whether the object in foveal view is the same object as that from the study picture. If the observer concludes that these objects differ, a positive change judgment can be indicated. However, if these objects are judged to be the same, gaze must be shifted to the next most likely change location, with the guidance and decision processes repeating until the picture is exhaustively inspected and a negative judgment can be rendered.

This early model by Parker (1978) highlights two aspects of change detection important to the present study.

First, Parker's proposed *comparison* and *guidance* operations nicely capture the dualistic nature of a change detection task. One component of this task is obviously rooted in memory. Objects or properties from a prechange scene must be remembered and *compared* with the corresponding objects or properties in a postchange scene. Memory is therefore required, not only for the object identifying information, but also for the location information needed to establish object correspondence. A second component of change detection seems very much like search. The comparison operation suggests regions of potential change in a scene; then, more focused processing is directed or *guided* to the suspect region in order to confirm or reject the presence of change. Depending on the outcome of this focal scrutiny, search would either terminate upon detection of the target or proceed to the next most promising change candidate, with these comparison and guidance processes repeating until the target is located or the display exhaustively inspected.

Second, Parker's (1978) model highlights the role of eye movements in a change detection task and the information about change detection strategy available from an eye movement analysis. Eye position is an integral part of his model, determining the efficacy of the comparison operation and, ultimately, change detection performance. When gaze is far from the changed object, the comparison process suffers, because the remembered object is now being compared with an object degraded by peripheral acuity. Gaze must therefore be guided to the changed object in order to make a confident change judgment. Parker quantified this strategic allocation of gaze by showing that observers frequently break from their preferred scanning of a scene to preferentially fixate the

This work was supported by a Beckman Institute postdoctoral fellowship. A preliminary account of this study was presented at the 1997 meeting of the Association for Research in Vision and Ophthalmology. I thank George McConkie for generously providing me with the use of his laboratory resources and Gary Wolverton for helping with the stimulus presentation and data analysis programming. I also thank John Henderson, Dan Levin, Mary Peterson, Ron Rensink, Dan Simons, and Jeremy Wolfe for comments on earlier versions of this manuscript. Correspondence regarding this article should be addressed to G. Zelinsky, State University of New York at Stony Brook, Department of Psychology, Psych B-240, Stony Brook, NY 11794-2500 (e-mail: gzelinsky@notes.cc.sunysb.edu).

changed object. He also found that observers may respond correctly without fixating the changed object when this change is very dramatic (such as cases in which an object is removed from a scene), presumably because the change was sufficiently large to overcome the acuity limitation.

The present study builds on this earlier work by Parker (1978), applying his approach to issues of visual search, memory, and strategy in a change detection task and extending his arguments to accommodate more recent ideas regarding the change detection process. With regard to search, I will show that the number of eye movements accompanying change detection increases with the set size of the scene and the orientation similarity of the objects undergoing change. With regard to memory, I will derive a gaze-based estimate of capacity limitation and argue that this limitation alone *cannot* account for the observed differential detection of change between oriented objects. I will contrast this memory capacity constraint with a limitation on the process that compares the pre- and postchange objects, an operation that *can* explain an orientation effect. Finally, by analyzing scanpaths on the pre- and postchange scenes, I will describe how observers were coordinating these search and memory processes to reach a change detection judgment. Like Parker, I argue that this strategy relies heavily on the nonfoveal encoding and comparison of objects, ending with a confirmatory eye movement to the peripherally detected target.

THE CHANGE DETECTION SEARCH PARADIGM

Recent change detection studies have shown that objects in realistic scenes can undergo a series of profound manipulations (substitutions, deletions, etc.), often for extended periods of time, without the observers' awareness of these changes (Grimes, 1996; Levin & Simons, 1997; McConkie & Currie, 1996; O'Regan, Deubel, Clark, & Rensink, 1997; O'Regan, Rensink, & Clark, 1999; Rensink, O'Regan, & Clark, 1997; Simons & Levin, 1997; Zelinsky, 1997). This *change blindness* is believed to be a direct result of an incomplete (O'Regan, 1992; Simons, 1996) or unavailable (Rensink, 1999, 2000) representation of objects in the change scenes and an indirect result of visual short-term memory (vSTM) limitations (Zelinsky & Loschky, 1998).¹ According to this explanation, our very limited ability to retain complex patterned information in vSTM (Phillips, 1974, 1983; Phillips & Christie, 1977), particularly when the task allows eye movement (Hayhoe, Bensinger, & Ballard, 1998; Irwin, 1991, 1996) and involves real-world objects (Zelinsky & Loschky, 1998), prevents some (or even most) of the objects in a briefly presented scene from being committed to memory. If the objects undergoing change happen not to be among these remembered items, the change will either go undetected until the relevant object is remembered in a subsequent change cycle,

or the observer will render an incorrect detection judgment (a miss).

To be sure, change detection is a memory task, and it is reasonable to assume that memory capacity limitations might affect change detection performance. But change detection also has a component of search, and search-related limitations might be expected to influence change detection difficulty (Rensink, 2000). Two characteristics of a change detection task also qualify it as a search task. First, the target's location in a change detection task is spatially uncertain. Just as spatial uncertainty requires observers in a search task to *search* for a complex target, participants in a change detection task must likewise *search* for the objects undergoing change. If an observer had foreknowledge of the target's location, the change detection task would degenerate into a single-item same-different task (with visual noise introduced by the irrelevant items), and search would degenerate into a simple detection task (again, with noise). Second, the target in both change detection and search is typically presented with distractors. Although distractors in a search task usually take the form of well-segmented items scattered randomly across a display, whereas the distractor set in a scene-based change detection task consists of other objects in the scene, in both cases the target must be isolated from irrelevant stimuli. Again, without distractor noise, both change detection and search become trivial.

The present study addresses both memory and search components of a change detection task by adding to a popular change detection paradigm two familiar search manipulations: set size and orientation similarity. We know from the search literature that detection performance for a complex and spatially uncertain target will worsen as objects are added to the display, a decrement known as a *set size effect*. One popular explanation for the set size effect posits that the featural primitives for each additional item must be correctly bound before an item can be compared with the target template (Treisman & Gelade, 1980; Treisman & Schmidt, 1982). If these binding or comparison operations are serial, or parallel but subject to capacity constraints (Townsend, 1976, 1990), search performance will degrade with increasing set size. An identical argument can be applied to scene-based change detection. As the number of potentially changed objects in the scene increases, so too will the number of time-consuming comparison operations needed to locate the change target. Given the fully realistic scenes used as stimuli in recent change detection studies, it is impossible to determine the contribution of set size to the detection difficulty demonstrated by observers. What is clear, however, is that the number of objects appearing in these scenes, although unspecified, was certainly large—easily ranging into the hundreds, depending on one's definition of an object. Such large set sizes might, therefore, have played a considerable role in the poor change detection performance reported in these studies. The present study uses pseudorealistic

scenes in which the number of objects can be more accurately specified, thereby allowing for the systematic manipulation of this variable and a meaningful interpretation of its role in change detection.

This study also characterizes change detection in terms of object similarity. Borrowing again from the search literature, we know that search performance varies with the distinctiveness of the target in the display (Duncan & Humphreys, 1989, 1992; Treisman & Sato, 1990; Wolfe, 1994; Wolfe, Cave, & Franzel, 1989). When the target is very distinct, as would be the case when it is dissimilar from the distractors, search performance is good. However, if this distinctiveness is lessened, either by increasing the similarity between the target and the distractors or by making the distractor set very heterogeneous, search becomes more difficult. The reason for this effect of similarity on search performance is again linked to the comparison operation: The more similar a distractor is to the target, the more time-consuming the act of rejecting the item. A similar dynamic might characterize comparisons in a change detection task. Changes between dissimilar objects may be relatively easy to detect, whereas changes between similar objects may result in prolonged detection times and/or high miss rates. Because the objects undergoing change in previous scene-based change detection studies often had an unspecified similarity relationship, it is difficult to know how this variable affected performance. The present study manipulates the orientation similarity of changing real-world objects in order to better understand this relationship.

METHOD

Participants

Six experimentally naive observers were paid \$8/h for their participation in this experiment. None of the participants required visual correction, and all had normal color vision.

Stimuli

The stimuli were pseudorealistic color scenes depicting common real-world objects (tools or toys) arranged on an appropriate background surface (a workbench or a crib). The objects and backgrounds composing these scenes were part of a custom-created image database, the details of which are described more fully in Zelinsky (1999b). The scenes used in this experiment were assembled on line from objects in this database, using a 486 PC clone, and presented to observers on a 21-in. Conrac color monitor. Despite their on-line generation, the compositions of these scenes were completely determined prior to the experiment, including the identity of the background surface, the number and identity of the objects appearing in the image, and the locations of these objects in the scene. The final images viewed by observers were 756×486 pixels in dimension (subtending $18^\circ \times 11.6^\circ$ of visual angle) and were of near-photographic quality. Figure 1 (top) shows a grayscale version of one of these scenes depicting an arrangement of nine toys on a background crib surface. All of the objects could fit inside a 2.4° bounding box, and their display locations were constrained to 18 positions, illustrated in Figure 1 (bottom) by the white squares superimposed over an empty workbench background surface. These placement constraints forced objects to appear at five visual eccentricities relative to the center of the image—with 4 of the 18 object positions each at 2.4° , 3.4° , 5.3° , and 7.2° from the center,

and the remaining two positions having a 4.8° eccentricity. These constraints also produced a minimum and maximum center-to-center separation between objects of 2.4° and 14.5° , respectively.

Procedure

A trial can be conceptualized as a repeating sequence of four 300-msec images (Figure 2). The first and third images depicted one of the above-described crib or workbench scenes. These scenes were either (1) identical (a target-absent no-change trial) or (2) different with respect to a single changed object (a target-present change trial). An example of a change is shown in Figure 2 between the trumpet in Image 1 and the baby bottle in Image 3. Note that this type of change is a simple substitution that preserves the configuration of items in the display, with the substituted object not appearing elsewhere in either of the two scenes. The second and fourth images were a full-screen color noise mask. The noise mask performed two functions. First, it disrupted any apparent motion percept that might accompany the sequential presentation of similar images. Second, it removed the sudden visual transient that would otherwise make the detection of change between Images 1 and 3 trivial (Rensink et al., 1997).

This sequence of four images was repeated until the trial was terminated by a buttonpress. The observer's task was to indicate the presence or absence of a change as quickly as possible without sacrificing accuracy. Judgments were indicated by two hand-held buttons, with the target-present response mapped to the button held in the dominant hand. Preceding each trial was a central fixation cross (an \times in a box) that would position initial gaze at the center of Image 1. The fixation cross remained visible until the observer initiated the trial by pressing one of the buttons. The participants were allowed 24 practice trials providing examples from each of the experimental levels (described below), using object configurations not appearing in the data trials.

Apparatus

A Fourward Technologies Generation V dual-Purkinje-image eyetracker was used to record the position of the right eye every millisecond from the initial onset of Image 1 to the termination of the trial. The spatial precision of this tracker is better than ± 1 min of visual angle. A dental impression bite-bar was constructed for each participant, and the observer was asked to remain on this bite-bar throughout the experiment until scheduled breaks following every block of 50 trials. A 9-point calibration was performed prior to starting the experiment and again before every block of trials. Other than the instructions required to carry out this calibration and a reminder to look at the fixation cross whenever it appeared on the screen, the observers were not told how to direct their gaze in this study. Saccadic eye movements were extracted off line, using a velocity-based algorithm implementing a roughly 12.5 %/sec detection threshold.² Once a saccade was detected, more conservative velocity thresholds were used to determine exactly when the eye began to move and when it eventually stopped. Fixation duration was defined as the period between the offset of one saccade and the onset of the next, with the term *initial saccade latency* reserved for the time needed to execute the first eye movement of a trial.

Design

Scenes depicted one, three, or nine objects on an otherwise empty crib or workbench surface. Objects were selected at random from a set of 10 appropriate to each scene type and were arranged into positions randomly selected from the 18 locations illustrated in Figure 1 (bottom). Because the selection of objects was without replacement, an object was never duplicated in a given scene. The 300 trials per observer were evenly divided into change (target-present) and no-change (target-absent) conditions and three set sizes, with both manipulations randomly interleaved throughout the experiment. The 50 change trials per set size were in turn divided




Figure 1. Top: A grayscale version of one of the nine-item scenes used as stimuli in this study. Not shown are one- and three-item set size displays. Bottom: The 18 allowable object positions (white squares) and their initial visual eccentricities. Object locations are relative to the center of an imaginary 2.4° bounding box enclosing each item.

into 25 similar-orientation and 25 different-orientation trials, a manipulation referring to the orientation of the objects making up the target. Half of the object change pairs had a similar orientation (the major axes for both objects were rotated roughly 45° clockwise or counterclockwise from vertical), and half had a different orientation (one object was rotated 45° clockwise, whereas the other object was rotated 45° counterclockwise). All of the changes reported here occurred within a given superordinate category, meaning that a tool

would change with another tool or a toy would change with another toy. As in the case of the set size and target present-absent manipulations, orientation condition was interleaved throughout the experiment.

Figure 3 illustrates three representative trials from the 12 experimental levels discussed above (set size [3] × orientation [2] × target present-absent [2]). All three examples depict target-present trials, with the left panel indicating the prechange image and the




Figure 2. Repeating sequence of images making up a representative change trial.

right panel indicating the postchange image. The top scenario shows a one-item trial in which the pre- and postchange objects have a similar orientation. The hammer in the left workbench scene changes with the clamp in the right workbench scene. The middle scenario shows a three-item crib display, now with the trumpet target in the prechange scene having a different orientation than the dinosaur target in the postchange scene. The bottom scenario again shows a workbench scene, this time with nine items and the change occurring between the paintbrush and the hammer objects. Note that in all three examples, the background remains the same between the pre- and the postchange scenes, and the distractors, when they are present, are likewise unchanged and always consistent with the scene type.

Target-absent trials, although of less theoretical interest in this study, provide a necessary baseline against which change detection performance must be gauged, particularly with regard to eye movement variables. Because even a statically presented object will be fixated with some probability, it is difficult to interpret a change manipulation involving this object without knowing its no-change base rate. Target-absent trials were created by “splitting apart” the

two scenes making up each change trial, then alternating each scene with itself (i.e., Images 1 and 3 would be identical). This method ensured that a dedicated no-change scene would exist for every scene (pre- and postchange) used in the change trials.

CHANGE DETECTION AS SEARCH Effects of Set Size and Orientation Similarity

Discarded Data

Of the 1,800 trials contributing to this study, 89 trials (about 5%) were discarded owing to an unacceptable loss of the eye position signal by the tracker. An unacceptable track loss was defined as (1) a cumulative loss of signal for a period greater than 750 msec for any reason or (2) a single loss of signal for a period exceeding 150 msec. Applying these rejection criteria resulted in the removal of trials in which the observer blinked excessively and cases in which the tracker did not immedi-




Figure 3. Top: A one-item similar-orientation change in a workbench scene. Middle: A three-item different-orientation change in a crib scene. Bottom: A nine-item similar-orientation change in a workbench scene. Not shown are the other three change conditions or examples of the no-change trials.

ately relock onto the eye following a blink. Either of these sources of track loss might compromise interpretation of the oculomotor data.

Manual Responses

If change detection is related to set size, observers might be expected to miss changes as objects are added to the display. Figure 4 (left) indicates such a relationship. Error rates in the one- and three-item data were uniformly low

(<3.5%), both in the no-change trials (false alarms) and in the orientation change conditions (misses). However, striking differences were observed when nine items appeared in the displays. When the pre- and postchange objects had a similar orientation, observers responded incorrectly on 38% of the trials. Recall that chance response would have resulted in 50% errors. Change detection accuracy improved when the pre- and postchange objects differed in orientation, but errors still remained high at 16%. Main ef-

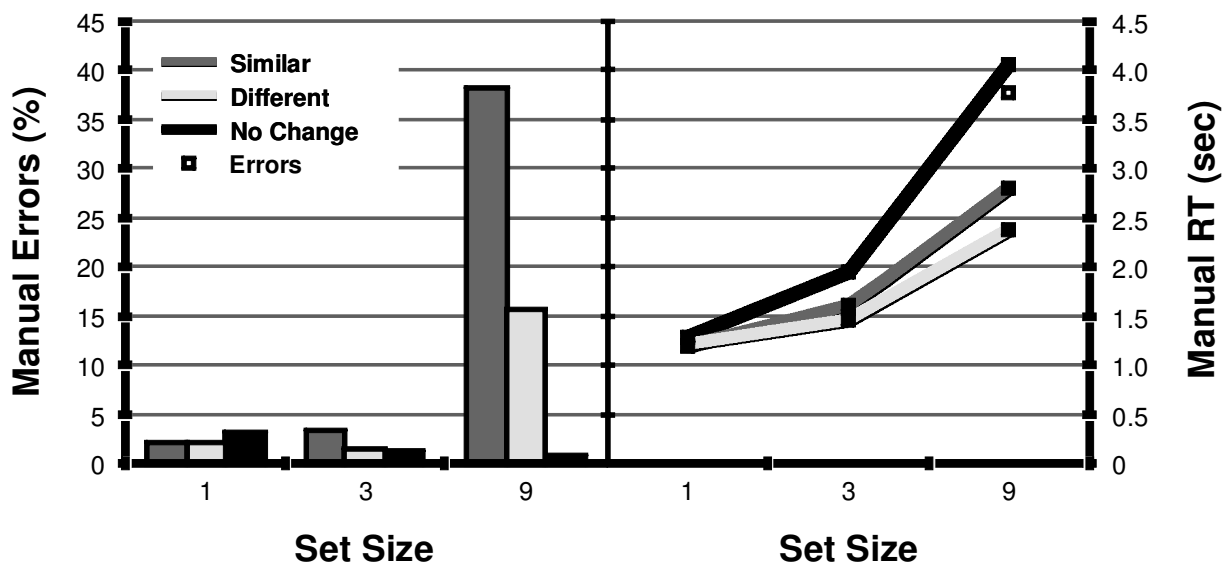


Figure 4. Left panel: Buttonpress errors (misses) as a function of set size and object change orientation. Target-absent data (false alarms) are shown as black bars. Note that misses increase with set size and that more errors occur when the changing objects have a similar orientation. Right panel: Manual reaction times (RTs) again plotted as a function of set size and object change orientation. Target-absent data are shown by the black line, and RTs for misses in the nine-item similar-orientation condition are indicated by the unfilled marker (the paucity of errors in the other conditions prevented the reporting of these data). Note that these RT data replicate the patterns of manual errors, thereby ruling out a speed–accuracy tradeoff.

fects of both set size and orientation were highly significant [$F(2,10) \geq 20.03, p < .001$], as was the interaction between these variables [$F(2,10) = 16.09, p = .001$].

Such high miss rates, although not usually observed in a search task, are nevertheless consistent with a straightforward search-based interpretation of these data—adding distractors decreased sensitivity to the change signal. Because there is an inherent bias to answer “no” under conditions of low sensitivity in a search task (which is why false alarm rates are typically small or nonexistent), the participants likely started to feel time pressure after searching unsuccessfully for the change target and reasoned that they would have found the target if, indeed, it was present in the display (Chun & Wolfe, 1996). They therefore frequently answered “target-absent,” despite the presence of a change target.

Because trials were response terminated in this study, differences in change detection performance might be expected in manual reaction times (RTs), as well as in error rates. More specifically, if both dependent measures vary with change detection difficulty, rather than simply offset each other in the form of a speed–accuracy tradeoff, RT search slopes should be steeper in the similar-orientation condition, relative to the different-orientation condition. This is precisely the pattern indicated in Figure 4 (right). Consistent with a relationship between set size and object change similarity, search slopes (and their standard errors) were 201 ± 38 msec/item in the similar-orientation change condition and only 148 ± 38 msec/item in the different-orientation trials [$F(2,10) =$

$16.51, p = .001$], a 53-msec/item slope difference that parallels the profound change detection deficit observed in the corresponding manual errors.³

Oculomotor Responses

Eye movements during visual search vary as a function of set size and target–distractor discriminability (Engel, 1977; Findlay, 1997; Gould, 1973; Hooge & Erkelens, 1998; Jacobs, 1986; Motter & Belky, 1998; Rayner & Fisher, 1987; Viviani, 1990; D. Williams, Reingold, Moscovitch, & Behrmann, 1997; L. Williams, 1967; Zelinsky, 1996; Zelinsky, Rao, Hayhoe, & Ballard, 1997; Zelinsky & Sheinberg, 1995, 1997). To the extent that change detection is also a search task (one in which the objects undergoing change are the targets), effects of these manipulations on oculomotor variables would also be expected. Figure 5 (left) indicates clear evidence for set size and orientation similarity effects on eye movements in this task. The number of saccades initiated before the change detection judgment increased with set size [$F(2,10) = 20.77, p < .001$] and the orientation similarity of the changing objects [$F(1,5) = 15.94, p = .010$], with both variables again reliably interacting [$F(2,10) = 17.94, p < .001$]. The time taken to launch the initial saccade of a trial also increased with set size (Figure 5, right). Initial saccade latency rose from approximately 200 msec at a set size of one to 300 msec at a set size of nine [$F(2,10) = 6.44, p = .016$], although this effect failed to interact with the orientation manipulation [$F(2,10) = 1.37, p = .299$].




Figure 5. The number of saccades (left) and the latency of the initial saccadic eye movements (right) plotted as a function of set size and object change orientation.

Discussion

Framed in terms of search, change detection performance degraded with increasing target spatial uncertainty (set size) and decreasing target conspicuity (the orientation manipulation). Specifically, changes either were missed or took longer to detect as objects were added to a display, with this set size effect attenuated when the object change violated orientation. These findings replicate recent reports of change detection performances varying with the number of potentially changing items and the similarity relationships between the changing objects (Rensink, 2000; Scott-Brown & Orbach, 1999; P. Williams & Simons, 2000). The present experiment (see also Zelinsky, 1997, 1998) extends this growing body of work in two directions. First, the set size and similarity relationships reported here were obtained using real-world objects in pseudorealistic scenes. Second, I provide an initial documentation of oculomotor behavior in a variable set size change detection task. Observers made more eye movements and took longer to launch their initial saccades as objects were added to the display. As was argued in Zelinsky and Sheinberg (1997), delays associated with these oculomotor patterns would be expected to lengthen the time needed to reach a search decision, particularly if the number of eye movements is large, as would be the case when the target is embedded in a realistic scene.

Collectively, these findings suggest that statements of change detection difficulty must be qualified by the set size of the scene and the orientation similarity of the changing objects. The possibility that change detection difficulty may arise owing to very large set sizes and potentially high similarity relationships between the changing objects can no longer be ignored. In fact,

once a search-based explanation is considered, the surprising demonstrations of poor change detection performance recently introduced into the literature seem far less remarkable—and in some sense, even predictable. Failure to articulate these qualifications and control for their effects will make it impossible to evaluate alternative explanations of change detection difficulty, especially when the stimuli are real-world scenes depicting objects varying widely over these dimensions.

CHANGE DETECTION AS MEMORY: Evidence for Comparison Constraints

The present paradigm required the observers to compare properties of two scenes separated in time by a 300-msec interval, and it is this temporal separation between the target-defining objects that makes change detection a memory task. To the extent that the participants were successful in performing this task, they were therefore remembering objects from the nonvisible scene pair and comparing these objects with those in the visible scene. It makes sense, then, to ask the number of objects that can be remembered and compared on a given change cycle—a question of capacity.

The Conveyor Belt Conception of Memory Capacity

Common to every system is the notion of capacity, a term specifying the rate relationship between a system's input and its output (Broadbent, 1958). A large-diameter pipe allows water to flow through faster than does a small-diameter pipe and, therefore, has a greater capacity. In memory systems, this rate term is usually factored out, with capacity being expressed as the number of

items available to immediate memory at any given time (Miller, 1956). Capacity limitations are observed when the number of items requiring memorial representation exceed this availability. Although capacity constraints can take many forms (see Shiffrin, 1976, for an excellent treatment of this topic), one definition of capacity refers to a limitation on storage. This is probably the most common conception of a memory capacity limitation, with modern roots reaching back to the modal memory models (Atkinson & Shiffrin, 1968; Murdock, 1974; Peterson & Peterson, 1959; Waugh & Norman, 1965). A storage capacity limitation can be idealized as a sort of fixed-length conveyor belt. Only a limited number of items can fit on the belt at any one time. If the belt space is fully occupied and a new item is added, an old item must drop off. Using such a conveyor belt model, Rensink (2000) varied the display time in a repeating change detection task and estimated capacity to be approximately five oriented-bar items, a figure in line with other estimates of vSTM in the literature (Luck & Vogel, 1997; Pylyshyn & Storm, 1988).

Eye-Movement-Based Estimates of Memory Capacity

Eye movements have long been used to study visual memory phenomena (e.g., Friedman, 1979; Gould, 1973; Just & Carpenter, 1980; Locher & Nodine, 1974; Loftus, 1972; Noton & Stark, 1971; Parker, 1978; see Rayner, 1978, 1998, for comprehensive reviews), including those investigations into the hypothetical integrative perceptual buffer that spawned the renewed interest in scene-based change detection research (Grimes, 1996; McConkie, 1991; McConkie & Currie, 1996). Despite this legion of memory-related eye movement studies, relatively few attempts have been made to estimate memory capacity by using eye movement variables. The best known estimate of this sort comes from the work of Irwin (1991, 1992, 1996) on transsaccadic memory. By asking observers to make an eye movement during a standard partial-report task, Irwin was able to distinguish between performance benefits arising from *iconic* memory and benefits owing to vSTM (vSTM can survive an eye movement, iconic memory cannot). Using this technique and varying array size between 6 and 10 letters, Irwin (1992, Experiment 2) found that only about 3.6 items could be remembered across an eye movement, regardless of array size—a finding leading him to conclude that vSTM has a fixed capacity of 3–4 items. A recent study by Zelinsky and Loschky (1998), also using an eye contingent display change paradigm, obtained a slightly smaller estimate of visual memory capacity. They had observers study a 9-item scene in preparation for a spatially cued recognition test and counted the number of objects visited by gaze after fixation of the target object (the identity of which was unknown at that point in the trial). Rather than finding evidence for a conveyor belt conception of capacity, Zelinsky and Loschky observed a steady decrease in recognition per-

formance with each object fixated after the target. This performance decrease asymptoted into the long-term memory baseline after three intervening objects, leading these authors to conclude that only 1–3 items can be held in vSTM. By far the smallest (and most naturalistic) oculomotor estimate of visual memory capacity was reported by Hayhoe, Ballard, and colleagues (Ballard, Hayhoe, Li, & Whitehead, 1992; Ballard, Hayhoe, & Pelz, 1995; Hayhoe et al., 1998), using a block copying task. They found that observers often devoted one eye movement to retrieving a property of the target object (e.g., color), followed by a second eye movement to establish the location of the target in the display—making their estimate of vSTM capacity to be less than a single object.

The present study uses a preferential fixation technique to estimate memory capacity in a change detection task. This technique is based on the observation that observers overwhelmingly elect to fixate a changing object in preparation for their manual judgment. At the time of the buttonpress response, the observers were fixating the target object on 94% of the similar-orientation trials and 91% of the different-orientation trials.⁴ These frequencies reflect only correct responses and cases in which gaze was within the 2.4° bounding box enclosing each item. When the target fixation criterion was relaxed to a 3.0° box, these frequencies increased to 97% and 95%, respectively, suggesting that gaze was about to converge on the change objects (following a corrective saccade) even on those trials in which the target was not accurately fixated.

The tendency for observers to fixate an object change upon detection in a free-viewing task can be used to estimate memory capacity for objects in the first change cycle (i.e., before the reappearance of Image 1 in Figure 2). The two extreme scenarios help to illustrate this point. If observers were able to remember all nine objects from the initial prechange scene and accurately compare them with every object in the initial postchange scene, an eye movement should be directed to the target object (or its location in the following mask display) on 100% of the change trials. Likewise, if observers had no memory of the prechange objects, the frequency of target fixation in the first display cycle should be no greater than the no-change baseline rate (reflecting chance target fixation). Capacity is therefore defined as $S \times F$, where S is the set size and F is the cumulative probability of fixating the target location in the first change cycle.

Figure 6 shows the cumulative probability of fixating the target (or its location in the mask display) during the initial change cycle (within 1,200 msec of trial onset) as a function of set size.⁵ A target was considered fixated if the eye landed within 1° of the center of the bounding box enclosing either object of the change pair. There are two general points to note about these functions. First, they all decrease with set size [$F(2, 10) \geq 15.87, p < .001$]. Given that the observers would elect to fixate the object change if it




Figure 6. Cumulative probability of gaze reaching the change target in the different (light gray line) and similar (dark gray line) change conditions as a function of set size. The no-change baseline probability is indicated by the black line.

was detected, these sloped functions suggest a decrease in detection probability with the number of items in the display. Second, the slopes of these functions for both the similar- and the different-orientation change conditions were shallower than the no-change baseline slope [$F(2,10) \geq 4.59, p < .039$]. Because the no-change condition describes a random probability of object fixation (one that controls for object-specific saliency differences), finding preferential fixation of the changed objects implies the perception of change. These two observations define the boundary conditions of the capacity estimate. The observers were not remembering and comparing every object from the prechange scene, or else the set size functions would be flat and close to a probability of 1.0. Likewise, the observers were clearly remembering some of the prechange object properties, because they were fixating the changed objects more frequently than would be expected by chance inspection of the display.

How many objects were the observers remembering from the prechange scene? According to the $S \times F$ measure of capacity, the .92 probability of target fixation in the single-item trials indicates a memory capacity of about one object—an estimate reflecting a ceiling imposed by the number of objects in the display (the fact that the observers occasionally responded in this trivial change detection condition before accurately fixating the target made a 1.0 probability unattainable). The observers fixated the changing objects in the three-item different-orientation displays with a probability of .62, resulting in a capacity estimate of just under two items,

and targets in the nine-item different-orientation condition were fixated with a probability of .54, suggesting a capacity of almost five objects. Because memory capacity would not be expected to change as a function of set size, it must be assumed that the observers were capable of holding five objects in memory and that the lower three-item estimate reflects normal trial-to-trial memory variability (sometimes two objects were stored, sometimes five, etc.), possibly as a result of momentary lapses of attentional set. Note also that this five-object estimate is in near-perfect agreement with the orientation change capacity estimate derived by Rensink (2000), using an entirely different technique and class of materials.

Unplugging the Conveyor Belt

Despite its agreement with other estimates of vSTM capacity, the above-described fixed-capacity approach—and in fact, any estimate appealing to a conveyor belt metaphor of memory—cannot explain the observed effect of orientation change on fixation probability. As is indicated in Figure 6, object changes violating orientation were more likely to be fixated than targets having a similar orientation [$F(2,10) = 6.25, p < .017$]. If a fixed number of objects in the prechange scene (say, five) were randomly selected for representation in memory, why then would differently oriented targets enjoy such a detection advantage? It could not be the case that the observers chose to encode five objects in the nine-item different-orientation condition but only three objects when the changing items had a similar orientation. Be-

cause the similar- and different-orientation conditions were not defined until the initial postchange scene appeared, this manipulation could not have affected the number of items encoded during viewing of the prechange display. The observers would simply not have known at the time of the initial prechange scene whether the change would violate orientation. According to a fixed-capacity model, the number of prechange objects encoded into memory at a given set size should have been constant over trials and determined by the difficulty of the easiest detection decision. In other words, the observers should have committed about five objects to memory regardless of the orientation condition, making the obtained effect of orientation difficult to explain with a capacity argument.

Discussion

Change detection is a vSTM memory task, and as such it is tempting to ask how many objects can be remembered from the prechange scene. For better or for worse, answering this question often starts rolling the conveyor belt memory model. I report a memory capacity estimate based on preferential fixation that is not easily reconciled with such a conveyor belt conception of change detection. A five-object estimate of capacity in the nine-item different-orientation condition appeared to change when the target objects were made visually more similar—an observation violating the assumption of fixed capacity in the conveyor belt metaphor.

The present data are better described by a limitation on the operation comparing the pre- and the postchange objects. Note that this suggestion effectively redefines change detection as a signal-in-noise problem (see Zelinsky, 1998, for a computational model partially implementing this approach). The “signal” in this model is taken to be the magnitude of the change between the target objects. When the changing objects are similar, as would be the case in the similar-orientation condition, the signal would be small, and the comparison operation difficult. When the target objects are dissimilar, as in the different-orientation condition, the signal would be larger, and detection performance would improve. The noise term would vary with the number of distractor objects in the display. Recent models of search have demonstrated how distractor noise might produce set size effects by creating opportunities for false alarms (Eckstein, 1998; Palmer, Ames, & Lindsey, 1993); a similar dynamic might apply to change detection.

The comparison-constrained model outlined here extends current conceptions of change detection in two respects. First, it accounts for data patterns that are not easily explained by a conveyor belt memory model, yet it does not demand that this model be rejected (it is easy to imagine the above-described dynamics at work within a fixed-size subset of display objects). In this sense, it supplements, rather than supplants, existing theory. Second, change detection is a complex behavior having close ties

to visual search (as was argued above) and low-level perception and, as such, would benefit from bridges to these larger bodies of work. Conceptualizing change detection in terms of a signal-in-noise problem should help to build these bridges and allow the application of perception and search theory to the study of change phenomena.

CHANGE DETECTION AS STRATEGY: Evidence for Peripheral Encoding and Comparison

Any complex task amenable to a great deal of high-level control is also likely to have a strategic component, and eye movements might be used to better understand the influence of this component on behavior. A case in point is visual search. In a 1960 symposium on visual search techniques, Boynton outlined a strategy to help aircraft pilots better locate military targets on the ground. In this address, Boynton pointed out a critical relationship between eye movements and search efficiency: Good searchers make many short fixations, whereas poorer searchers make fewer, longer duration fixations. In addition to his more specific recommendations about optimal search altitude and aircraft velocity, Boynton advised pilots to make briefer fixations and more eye movements (see also Megaw & Richardson, 1979). In a more recent study designed specifically to address oculomotor scanning strategies during search, Scinto, Pillalamarri, and Karsh (1986) made two observations. First, the observers often made a systematic quadrant-by-quadrant search of a complex display, presumably in a strategic attempt not to revisit previously inspected regions.⁶ Second, fixation durations tended to increase, and saccade amplitudes tended to decrease, throughout the course of a prolonged search. They summarized this latter tendency with the rule: “If you can’t find what you’re looking for, look more carefully.” The final eye movement analyses of the present study look for a similar strategic component in a change detection task.

Evidence Against an Object-by-Object Comparison Strategy

The efficiency of a change detection scanning strategy is likely to vary with the operation comparing the component scenes. This point is nicely made by Just and Carpenter (1976) in a study using a three-dimensional (3-D) same-different task. The observers viewed two simultaneously presented geometric objects and had to judge whether the objects were different or the same except for a 3-D rotation. They found that performance in this task depended on the adoption of a very conservative comparison strategy, with the observers systematically selecting and comparing a feature on one of the objects with the corresponding feature on the other object and concluding that these objects were the same only after a thorough inspection of these corresponding features.

Eye movement analyses were used to reveal this feature-by-feature comparison process.

Observers might use an analogous comparison operation in a change detection task. An object in the prechange scene might be selected and fixated and gaze held at that location until the corresponding object appeared with the postchange scene. These two objects might then be compared, with this process being repeated for a different object pair if the comparison failed to reveal a change. Such an object-by-object comparison strategy would leave a telling oculomotor signature. Specifically, gaze would remain at a particular display location until both the pre- and the postchange objects were inspected. To look for this signature oculomotor pattern, object gaze durations were divided into three categories: (1) those gaze durations that spanned all three images in a change cycle (Scene 1 → Mask → Scene 2), (2) durations that spanned one of the scene images and the following mask (Scene 1 → Mask →, or Scene 2 → Mask →), and (3) those durations spanning only a single image (Scene 1, Mask, or Scene 2). If the observers were routinely maintaining gaze on a single object throughout a complete change cycle, gaze at an object location should frequently span at least three consecutive images in a trial. Figure 7 shows that this pattern did not often occur. Plotted are the relative frequencies of object gaze durations spanning one, two, or three frames for each of the change conditions. The observers held object gaze throughout a change cycle on only 12% of the similar-orientation change trials and 9% of the different-orientation trials. Furthermore, in the majority of cases in which the observers did appear to adopt this strategy, their gaze was already on the target objects—most likely, to confirm that the ob-

jects were indeed changing. When these final confirmatory durations were excluded from the analysis, the frequency of gaze durations spanning an entire change cycle dropped to less than 3%. By far the most common pattern (about 73%) was to look at one or two objects while they were visible, then to shift the gaze to a new location at the moment when the scene was replaced by the mask (an event counted here as a single image span). This finding suggests that the observers in the present change detection task were not using the conservative point-by-point fixation strategy reported by Just and Carpenter (1976), even though such a strategy would likely have resulted in accurate detection.

Evidence for a Peripheral Discrimination Strategy

If the observers were not performing an item-by-item inspection of potentially changed objects, then what strategy were they using? Figure 8 shows two representative change detection scanpaths that partly answer this question. The three leftmost images (a–c) represent the sequence of displays (top to bottom) viewed by an observer on one trial; the three rightmost images (d–f) represent this sequence for a second trial. In both cases, the observer responded correctly before a fourth scene appeared. The black arrows indicate the observer's eye movements on these scenes, and the dotted arrows indicate eye movements during the mask interval following each scene presentation (not shown). The changing objects in the left trial were the baby doll (a, c) and the baby bottle (b); the trumpet (d, f) and the panda bear (e) were changing in the right trial.

Turning first to the left column, the observer devoted two brief fixations to objects in the prechange scene (a),




Figure 7. Frequency of gaze durations spanning one, two, or three of the repeating image frames as a function of the orientation similarity of the object change. Note that gaze duration on an object rarely spanned two (a scene and the following mask) or three (scene, mask, scene) frames, with the most common behavior being an eye movement following each display change.




Figure 8. Scanpaths from two representative change detection trials (a–c, d–f). The top two images (a, d) show the initial prechange scenes for both of these trials, the middle two images (b, e) show the initial postchange scenes, and the bottom two images (c, f) are identical to the top two scenes and mark the beginning of the second change cycle. Solid arrows superimposed on a scene represent eye movements occurring during the 300-msec viewing of that scene; dotted arrows represent eye movements during the 300-msec masking interval following scene offset (not shown).

neither of which was the target, then made a single gaze shift during the following masking interval. Upon presentation of the initial postchange scene (b), the observer fixated the target but was apparently unaware of the object change (as is indicated by the three eye movements following target fixation). Rather, the observer continued his inspection of the display, almost as if executing a preprogrammed oculomotor inspection pattern to a statically presented scene. This inspection was abruptly halted upon re-presentation of the initial scene (c), an event resulting in the almost immediate refixation of the target. Gaze then remained at the target location until the manual response 1,742 msec after trial onset. The ob-

server therefore fixated the target once during scanning, then again immediately before the buttonpress. I will refer to this as a two-fixation pattern in reference to the two separate fixations on the target objects (with the word *separate* here meaning that at least one nontarget object was fixated between the two target views). Note, however, that the two target fixations in a two-fixation trial probably serve two very different functions. The first target fixation, the one occurring during scene inspection, likely served the standard function of encoding object properties into memory. The second target fixation, because it represented a fairly abrupt break in what was, up to that point, an unremarkable scanpath,

probably played a more confirmatory role. The most likely scenario is that the observers detected the change on the fixation preceding the final targeting saccade, then looked at the target in order to confirm that the objects were indeed changing. In this sense, the final target fixation may have contributed to the confidence of the change detection judgment but was probably not critical to the actual detection of change. This two-fixation pattern occurred on approximately 41% of the correct similar-orientation change trials and on 33% of the different-orientation trials. Interestingly, of the incorrect nine-item similar-orientation trials, only 7% showed this two-fixation pattern, with this drop from the 41% frequency being attributed to the absence of a final confirmatory targeting saccade. When the observers missed a change, they typically only fixated the target once and were not looking at the target at the time of the button-press.⁷

From the two-fixation pattern indicated in Figure 8 (left), as well as from the Parker (1978) study and the above-reported preferential fixation analysis, it is safe to conclude that change detection is very often also a peripheral discrimination task. Objects appearing in the visual periphery were compared with the corresponding objects encoded into memory, with a confirmatory eye movement directed to the target when this comparison indicated a change. It would also be tempting to conclude from Figure 8 (left) that the observers needed to fixate the target object and encode it clearly into memory before that object could be peripherally detected as changing, but such a conclusion would be in error, as is demonstrated by the scanpath appearing in Figure 8 (right). As in the two-fixation trial, the observer once again made two brief fixations to nontarget objects in the prechange scene (d). However, when the postchange scene appeared (e), the observer now directed a saccade very near to the changed object, followed by a second smaller eye movement during the masking interval that brought the gaze even closer to the target. Upon re-presentation of the initial scene (f), a corrective saccade acquired the target almost immediately, with the manual response occurring soon after (1,404 msec from trial onset). I refer to this pattern as a one-fixation scanpath to highlight the fact that the target was fixated only once throughout the entire change detection trial. One-fixation patterns were observed on approximately 51% of the similar-orientation trials and 57% of the different-orientation trials.⁸ It is also important to note that these single-target fixations likely served the same confirmatory role as the final targeting saccades in the two-fixation cases. The eye movement toward the target in Figure 8e was probably not coincidental. If indeed these targeting saccades in the one-fixation trials were confirmatory, it becomes possible to conclude that the perception of change in this study was largely independent of fixation on the changing objects.

Discussion

Although it is perhaps unwise to examine too closely the behavior from two trials, the scanpaths illustrated in Figure 8 suggest patterns that help shed light on a change detection strategy. One pattern is relatively clear and fails to support the conservative point-by-point fixation strategy reported by Just and Carpenter (1976). The observers almost never fixated a single object throughout an entire change cycle during their inspection of the scenes. Rather, their strategy apparently was to fixate objects in these scenes much as they would fixate items in a static search display. Occasionally the target would be fixated during this inspection; more often, it would not. If a target was detected, the gaze would be sent directly to the target location in order to confirm the presence of object change and to synchronize eye position with the manual response—but the detection event itself did not require target fixation.

The scanning strategy adopted by the observers in this study highlights the dualistic nature of eye movements in a change detection task. Eye movements can be either an information-gathering behavior affecting the detection process or a confirmatory behavior driven by the perception of change. Although both characterizations are undoubtedly true in part, it is probably more accurate to say that change detection influences gaze direction more than gaze direction influences change detection. Certainly, a fortuitous fixation on a target object might allow its properties to be better encoded and later used in the comparison operation underlying detection, and in this sense eye position can be said to affect change detection performance. However, given the preferential fixation tendency indicated in Figure 6 and the scanpaths in Figure 8, it appears that gaze is also very often driven over the display by the perception of change.

Because it was unnecessary to be looking at an object in order to detect that it had changed, the signal indicating change must have originated in the visual periphery. According to a recent computational model detailing the nature of this change signal (Zelinsky, 1998, 1999a), the two alternating scenes (the entire scenes, not just the scene objects) in a repeating change detection task are represented by arrays of linear filter responses, and these two arrays are compared with each new scene presentation. The outcome of this comparison operation is a map indicating the difference or change between the two scenes. Changing objects would result in a salient region on this difference map, which could then be used to drive the confirmatory eye movements illustrated in Figure 8. According to this model, change detection is most often accomplished by comparing peripherally encoded representations with peripherally viewed representations, thereby making this behavior relatively insensitive to target fixation and explaining all of the oculomotor scanning and preferential fixation patterns reported in this study.

A change detection strategy relying on a scene-wide difference signal, however, is not without its limitations. Indeed, it is far less conservative than an object-by-object fixation strategy, depending much more on peripheral discrimination and perceptual memory. An object-by-object strategy makes minimal demands on immediate memory, requiring only a single object to be maintained per change cycle. The strategy indicated here requires the maintenance of multiple objects (or some subset of their properties) in memory and involves a comparison between all of these objects and their corresponding peripheral descriptions in the currently viewed display. A detection strategy using this generalized difference signal will therefore be more error prone than an exacting object-by-object strategy—but it will also probably require far less time. The observers apparently elected to trust their perceptual memories and peripheral discrimination abilities to detect change as quickly as possible in this task. However, as the high error rates in the similar-orientation condition indicate, this trust was often misplaced. At the very least, it appears that the observers were overestimating the capacity of the above-described comparison operation to generate a detectable change signal.

It is unclear whether the strategy adopted by the observers in this study should be characterized as suboptimal. Certainly, if the observers had adopted an object-by-object comparison strategy, their errors would have decreased, and in this sense these errors can be rightly attributed to an overly optimistic comparison strategy. However, it is equally correct to say that the adopted strategy was a reasonably efficient method of detecting change between differently oriented objects. Characterizations of the adopted strategy as either inefficient or efficient must, therefore, be qualified by the specific objects undergoing change. When these objects are very similar, the resulting difference signal may go undetected when using a strategy relying on peripheral discrimination; when these objects are dissimilar and give rise to a large difference signal, this same strategy would efficiently yield reliable evidence for change. Given the variety of stimuli currently being used to study change detection (ranging from simple geometric shapes to video sequences), it is probably prudent to consider how the salience of these various change signals might interact with an observer's change detection strategy.

GENERAL DISCUSSION

There is no such thing as a unitary change detection behavior. Like most complex behaviors, it is a composite of several more elemental cognitive processes acting in concert. Among its many component operations, the task of change detection includes (1) selecting and encoding one or more objects or properties into memory, (2) systematically comparing these objects or properties with corresponding representations in the visible scene, and (3) coordinating these selection, encoding, and comparison operations to eventually reach a change

judgment. In this sense, change detection is a memory task, a search task, and a task requiring the strategic combination of these memory and search components. The "task" of understanding change detection must therefore be to address all of these aspects. Failure to adopt such a holistic perspective will inevitably leave us with an incomplete or distorted picture of change detection behavior.

The memory aspect of change detection is fairly obvious, and indeed it is this aspect that has attracted the most research (Grimes, 1996; Levin & Simons, 1997; O'Regan et al., 1999; Rensink et al., 1997; Simons, 1996; Simons & Levin, 1997). Almost as obvious is the search aspect, although comparatively few studies have been devoted to this topic (Rensink, 2000; Zelinsky, 1997). Simply put, a search process must be suspected in any task requiring a judgment in response to a spatially uncertain object, particularly if that object is featurally complex and is presented simultaneously with multiple distractors. To the extent that a change detection task shares these defining search characteristics, it is also a search task. The strategic aspect of change detection is less obvious, probably because the manual dependent measures typically used to study detection behavior make discerning strategic factors difficult.

The present study exploited the spatiotemporal resolution of eye movements to investigate these three aspects of change detection. To summarize, (1) a scanpath analysis suggested a clear strategy for detecting change, one relying on peripheral processing to both encode and compare objects between the change scenes, (2) a preferential fixation analysis revealed a memory constraint imposed by the orientation similarity of the changing objects, and (3) a search analysis demonstrated effects of set size and target discriminability on the number of eye movements and the latency of initial saccades in a change detection task. All three of these constraints were explained in terms of a difference signal resulting from a comparison of the scenes undergoing change, with performance depending on the signal's strength (i.e., the similarity relationship between the changing objects) and the discriminability of this signal, given background noise.

In a behavior as complex and multifaceted as change detection, one of the greatest obstacles to understanding is learning the contributions of its various components. Eye movements provide a means of teasing apart these components and determining the limitations imposed by each on change detection success and failure.

REFERENCES

- ATKINSON, R., & SHIFFRIN, R. (1968). Human memory: A proposed system and its control processes. In K. Spence & J. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, pp. 89-195). New York: Academic Press.
- BALLARD, D., HAYHOE, M., LI, F., & WHITEHEAD, S. (1992). Hand-eye coordination during sequential tasks. *Philosophical Transactions of the Royal Society of London: Series B*, *337*, 331-339.
- BALLARD, D., HAYHOE, M., & PELZ, J. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, *7*, 66-80.

- BOYNTON, R. (1960). *Visual search techniques* (NAS-NRC Publication 712). NRC Vision Committee Symposium.
- BROADBENT, D. (1958). *Perception and communication*. New York: Pergamon.
- CHUN, M., & WOLFE, J. (1996). Just say no: How are visual searches terminated when there is no target present? *Cognitive Psychology*, **30**, 39-70.
- DUNCAN, J., & HUMPHREYS, G. (1989). Visual search and stimulus similarity. *Psychological Review*, **96**, 433-458.
- DUNCAN, J., & HUMPHREYS, G. (1992). Beyond the search surface: Visual search and attentional engagement. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 578-588.
- ECKSTEIN, M. (1998). The lower visual search efficiency for conjunctions is due to noise and not serial attentional processing. *Psychological Science*, **9**, 111-118.
- ENGEL, F. (1977). Visual conspicuity, visual search and fixation tendencies of the eye. *Vision Research*, **17**, 95-108.
- FINDLAY, J. (1997). Saccade target selection during visual search. *Vision Research*, **37**, 617-631.
- FINDLAY, J., & GILCHRIST, I. (1998). Eye guidance and visual search. In G. Underwood (Ed.), *Eye guidance in reading and scene perception* (pp. 295-312). Amsterdam: Elsevier.
- FRIEDMAN, A. (1979). Framing Pictures: The role of knowledge in automatized encoding and memory for gist. *Journal of Experimental Psychology: General*, **108**, 316-355.
- GOULD, J. (1973). Eye movements during visual search and memory search. *Journal of Experimental Psychology*, **98**, 184-195.
- GRIMES, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.), *Perception* (pp. 89-110). New York: Oxford University Press.
- HAYHOE, M., BENSINGER, D., & BALLARD, D. (1998). Task constraints in visual working memory. *Vision Research*, **38**, 125-137.
- HOOGE, I., & ERKELENS, C. (1998). Adjustment of fixation duration during visual search. *Vision Research*, **38**, 1295-1302.
- HOROWITZ, T., & WOLFE, J. (1998). Visual search has no memory. *Nature*, **394**, 575-577.
- IRWIN, D. (1991). Information integration across saccadic eye movements. *Cognitive Psychology*, **23**, 420-456.
- IRWIN, D. (1992). Memory for position and identity across eye movements. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 307-317.
- IRWIN, D. (1996). Integrating information across saccadic eye movements. *Current Directions in Psychological Science*, **5**, 94-99.
- JACOBS, A. M. (1986). Eye-movement control in visual search: How direct is visual span control? *Perception & Psychophysics*, **39**, 47-58.
- JUST, M., & CARPENTER, P. (1976). Eye fixations and cognitive processes. *Cognitive Psychology*, **8**, 441-480.
- JUST, M., & CARPENTER, P. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, **87**, 329-354.
- LEVIN, D. T., & SIMONS, D. J. (1997). Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin & Review*, **4**, 501-506.
- LOCHER, P. J., & NODINE, C. F. (1974). The role of scanpaths in the recognition of random shapes. *Perception & Psychophysics*, **15**, 308-314.
- LOFTUS, G. (1972). Eye fixations and recognition memory for pictures. *Cognitive Psychology*, **3**, 525-551.
- LUCK, S., & VOGEL, E. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, **390**, 279-280.
- MCCONKIE, G. (1991). Perceiving a stable visual world. In G. d'Ydewalle & J. Van Rensbergen (Eds.), *Perception and cognition* (pp. 333-348). Amsterdam: Elsevier.
- MCCONKIE, G., & CURRIE, C. (1996). Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 563-581.
- MEGAW, E., & RICHARDSON, J. (1979). Target uncertainty and visual scanning strategies. *Human Factors*, **21**, 303-315.
- MILLER, G. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, **63**, 81-97.
- MOTTER, B., & BELKY, E. (1998). The guidance of eye movements during active visual search. *Vision Research*, **38**, 1805-1815.
- MURDOCK, B. (1974). *Human memory: Theory and data*. Hillsdale, NJ: Erlbaum.
- NEISSER, U. (1967). *Cognitive psychology*. Englewood Cliffs, NJ: Prentice-Hall.
- NOTON, D., & STARK, L. (1971). Scan paths in saccadic eye movements while viewing and recognizing patterns. *Vision Research*, **11**, 929-942.
- O'REGAN, K. (1992). Solving the 'Real' mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, **46**, 461-488.
- O'REGAN, K., DEUBEL, H., CLARK, J., & RENSINK, R. (1997). Picture changes during blinks: Not seeing where you look and seeing where you don't look [Abstract]. *Investigative Ophthalmology & Visual Science*, **38**, S707.
- O'REGAN, K., RENSINK, R., & CLARK, J. (1999). Change-blindness as a result of "mudsplashes." *Nature*, **398**, 34.
- PALMER, J., AMES, C., & LINDSEY, D. (1993). Measuring the effect of attention on simple visual search. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 108-130.
- PARKER, R. (1978). Picture processing during recognition. *Journal of Experimental Psychology: Human Perception & Performance*, **4**, 284-293.
- PETERSON, L., & PETERSON, M. (1959). Short-term retention of individual items. *Journal of Experimental Psychology*, **61**, 12-21.
- PHILLIPS, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics*, **16**, 283-290.
- PHILLIPS, W. A. (1983). Short-term visual memory. *Philosophical Transactions of the Royal Society of London B*, **302**, 295-309.
- PHILLIPS, W. A., & CHRISTIE, D. (1977). Components of visual memory. *Quarterly Journal of Experimental Psychology*, **29**, 117-133.
- PYLYSHYN, Z., & STORM, R. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, **3**, 179-197.
- RAYNER, K. (1978). Eye movements in reading and information processing. *Psychological Bulletin*, **85**, 618-660.
- RAYNER, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, **124**, 372-422.
- RAYNER, K., & FISHER, D. L. (1987). Letter processing during eye fixations in visual search. *Perception & Psychophysics*, **42**, 87-100.
- RENSINK, R. (1999, May). *Seeing, sensing and scrutinizing*. Paper presented at the 3rd Annual Vision Research Conference, Fort Lauderdale, FL.
- RENSINK, R. (2000). Visual search for change: A probe into the nature of attentional processing. *Visual Cognition*, **7**, 345-376.
- RENSINK, R., O'REGAN, K., & CLARK, J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, **8**, 368-373.
- SCINTO, L., PILLALAMARRI, R., & KARSH, R. (1986). Cognitive strategies for visual search. *Acta Psychologica*, **62**, 263-292.
- SCOTT-BROWN, K., & ORBACH, H. (1999). Contrast discrimination, non-uniform patterns and change blindness. *Proceedings of the Royal Society of London: Series B*, **265**, 2159-2166.
- SHIFFRIN, R. (1976). Capacity limitations in information processing, attention, and memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Memory processes* (Vol. 4, pp. 177-236). Hillsdale, NJ: Erlbaum.
- SIMONS, D. (1996). In sight, out of mind: When object representations fail. *Psychological Science*, **7**, 301-305.
- SIMONS, D., & LEVIN, D. (1997). Change blindness. *Trends in Cognitive Sciences*, **1**, 261-267.
- TOWNSEND, J. (1976). Serial and within-stage independent parallel model equivalence on the minimum completion time. *Journal of Mathematical Psychology*, **14**, 219-238.
- TOWNSEND, J. (1990). Serial vs. parallel processing: Sometimes they look like tweedledum and tweedledee but they can (and should) be distinguished. *Psychological Science*, **1**, 46-54.
- 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 & Performance*, **16**, 459-478.

- TREISMAN, A., & SCHMIDT, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, **14**, 107-141.
- VIVIANI, P. (1990). Eye movements in visual search: Cognitive, perceptual and motor control aspects. In E. Kowler (Ed.), *Eye movements and their role in visual and cognitive processes* (pp. 353-393). Amsterdam: Elsevier.
- WAUGH, N., & NORMAN, D. (1965). Primary memory. *Psychological Review*, **72**, 89-104.
- WILLIAMS, D., REINGOLD, E., MOSCOVITCH, M., & BEHRMANN, M. (1997). Patterns of eye movements during parallel and serial visual search tasks. *Canadian Journal of Experimental Psychology*, **51**, 151-164.
- WILLIAMS, L. (1967). The effects of target specification on objects fixated during visual search. *Acta Psychologica*, **27**, 355-360.
- WILLIAMS, P., & SIMONS, D. (2000). Detecting changes in novel 3D objects: Effects of change magnitude, spatiotemporal continuity, and stimulus familiarity. *Visual Cognition*, **7**, 297-322.
- WOLFE, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, **1**, 202-238.
- 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 & Performance*, **15**, 419-433.
- ZELINSKY, G. (1996). Using eye saccades to assess the selectivity of search movements. *Vision Research*, **36**, 2177-2187.
- ZELINSKY, G. (1997). Eye movements during a change detection search task [Abstract]. *Investigative Ophthalmology & Visual Science*, **38**, S373.
- ZELINSKY, G. (1998). *Detecting changes between scenes: A similarity-based theory using iconic representations* (Tech. Rep. CNS-98-01). Urbana, IL: Beckman Institute, Cognitive Neuroscience Group.
- ZELINSKY, G. (1999a). Effects of visual similarity on change detection: A behavioral and computational analysis [Abstract]. *Investigative Ophthalmology & Visual Science*, **40**, S50.
- ZELINSKY, G. (1999b). Precuing target location in a variable set size "nonsearch" task: Dissociating search-based and interference-based explanations for set size effects. *Journal of Experimental Psychology: Human Perception & Performance*, **25**, 875-903.
- ZELINSKY, G., & LOSCHKY, L. (1998). Toward a more realistic assessment of visual short-term memory [Abstract]. *Investigative Ophthalmology & Visual Science*, **39**, S224.
- ZELINSKY, G., RAO, R., HAYHOE, M., & BALLARD, D. (1997). Eye movements reveal the spatio-temporal dynamics of visual search. *Psychological Science*, **8**, 448-453.
- ZELINSKY, G., & SHEINBERG, D. (1995). Why some search tasks take longer than others: Using eye movements to redefine reaction times. In J. Findlay, R. Kentridge, & R. Walker (Eds.), *Eye movement research: Mechanisms, processes and applications* (pp. 325-336). Amsterdam: Elsevier.
- ZELINSKY, G., & SHEINBERG, D. (1997). Eye movements during parallel-serial visual search. *Journal of Experimental Psychology: Human Perception & Performance*, **23**, 244-262.
- plung frequency) and (2) the direction of this movement remained unchanged for four additional consecutive samples. Once this algorithm detected a saccade, stricter criteria were used to determine the exact onset of the eye movement. Starting at the sample in which the saccade was detected (t), two moving windows then compared the mean of t , $t - 1$, and $t - 2$ with the mean of $t - 3$, $t - 4$, and $t - 5$. If this comparison yielded a difference of less than 0.025° , in both the horizontal and the vertical signals, the eye was presumed to be in a fixation. If this difference failed to meet the 0.025° criterion, as it almost certainly would at time t , then t was decremented by 1 msec and the process was repeated. Allowing this algorithm to continue would eventually yield a time in which the eye was stable before the saccade was launched. Applying this identical algorithm in the forward time direction would similarly yield the time in which the eye stabilized after the saccade was completed.
3. Note that slopes calculated from these RTs would likely be far steeper than comparable slopes obtained from a static search task using these stimuli. This slope difference is due partly to the masking intervals inflating the RT measure. Recall that the change target is not even defined until after the first mask (600 msec into the trial) and, depending on the RT, several additional masks may appear before the button-press. Although correcting for the initial masking interval could be accomplished by simply subtracting 300 msec from the RT, adjusting for the subsequent masks is less straightforward, because it is not known whether processing associated with signal differentiation occurred during these intervals. Because the present study makes only relative comparisons between the change detection slopes (rather than inferring processing on the basis of the absolute slopes; see Rensink, 2000), no attempt was made to correct for these masking delays.
4. Although small, this trend toward fewer target fixations in the different-orientation condition is consistent with an observation by Parker (1978). He found fewer eye movements directed to deleted objects, relative to other forms of change, and attributed this finding to differential peripheral discriminability. If a change between differently oriented objects was obvious and could be confidently detected using peripheral vision, even a confirmatory saccade might be unnecessary.
5. Because the capacity estimate derived from the preferential fixation technique is specific to a single change cycle, such data segregation was needed to prevent information from subsequent display cycles from inflating this estimate.
6. Note that a quadrant-by-quadrant oculomotor search pattern might also be interpreted as evidence for a covert search process' failing to keep track of rejected distractors (Horowitz & Wolfe, 1998) or even for the nonexistence of covert search movement when gaze is free to move (Findlay & Gilchrist, 1998).
7. On the handful of nine-item similar-orientation error trials in which the target was fixated twice, the observer obtained a high-resolution foveal description of the target on two separate occasions, yet still responded incorrectly (see O'Regan et al., 1997, for a previous demonstration of change blindness for fixated objects). However, given the infrequent occurrence of this event in the present data (only 4 out of 54 total trials), it is probably not wise to draw conclusions about memory limitations on the basis of this pattern.
8. The percentages of two- and one-fixation similar- and different-orientation trials do not add to 100%, owing to a handful of zero-fixation (6% in the similar-orientation data, 9% in the different-orientation data) and three-fixation (2% in the similar-orientation data, 1% in the different-orientation data) patterns. These scanpaths, although potentially suggesting additional change detection strategies, were extremely infrequent and therefore not considered further.

NOTES

1. Rensink (2000) makes the useful distinction between objects that were never represented (i.e., their features never unitized into a whole) and objects that were at one time well represented but now have lost their featural coherence. Because in either event, the object would no longer be available as a distinct object to the processes serving change detection, I refer here to both possibilities as a limitation of object availability.

2. More specifically, a movement was labeled a saccade if (1) the horizontal or vertical eye position sampled at time t deviated by 0.05° of arc from another sample taken at $t - 4$ msec (assuming a 1000-Hz sam-

(Manuscript received April 6, 1999;
revision accepted for publication June 5, 2000.)