VISUAL COGNITION, 2000, *7* (1/2/3), 377–396

Attenuated Change Blindness for Exogenously Attended Items in a Flicker Paradigm

Brian J. Scholl

Rutgers Center for Cognitive Science, Piscataway, NJ, USA

When two scenes are alternately displayed, separated by a mask, even large, repeated changes between the scenes often go unnoticed for surprisingly long durations. Change blindness of this sort is attenuated at "centres of interest" in the scenes, however, supporting a theory of change blindness in which *attention* is necessary to perceive such changes (Rensink, O'Regan, & Clark, 1997). Problems with this measure of attentional selection—via verbally described "centres of interest"—are discussed, including worries about describability and explanatory impotence. Other forms of attentional selection, not subject to these problems, are employed in a "flicker" experiment to test the attention-based theory of change detection. Attenuated change blindness is observed at attended items when attentional selection is realized via involuntary *exogenous capture* of visual attention—to late-onset items and colour singletons—even when these manipulations are uncorrelated with the loci of the changes, and are thus irrelevant to the change detection task. These demonstrations ground the attention-based theory of change blindness in a type of attentional selection which is understood more rigorously than are "centres of interest". At the same time, these results have important implications concerning the nature of exogenous attentional capture.

It often seems, from the perspective of vision science, that our perceptual systems work *too* well. They provide us with extremely strong intuitions about how perception works—about how and when we represent visual scenes—which often turn out to be misleading. When viewing natural scenes, for example, we have the sense that we are *seeing* the entire scenes in all of their

---

Brian Scholl is now at the Vision Sciences Laboratory in the Department of Psychology at Harvard University. Please address all correspondence to B.J. Scholl, Vision Sciences Laboratory, Department of Psychology, Harvard University, 33 Kirkland Street, Cambridge, MA02138. Email: scholl@wjh.harvard.edu

Thanks to James Intriligator, Zenon Pylyshyn, Dan Simons, Jan Theeuwes, and two anonymous reviewers for many helpful suggestions. Thanks also to Damien Henderson and Allan Kugel for help constructing stimuli and running subjects, and to Michael Tarr for making available the scans of the Snodgrass and Vanderwart stimuli.

detail. It is thus natural to predict that we would immediately detect sudden changes made to most parts of an actively viewed scene, such as an object suddenly changing colour, or being replaced by another object.

These intuitions, which suggest the existence of detailed internal representations of visual scenes, are deceptive. When changes like these are made to scenes during eye-movements, for instance, they often go unnoticed, even when they are repeated and expected. Grimes (1996), for example, observed this type of "change blindness" when, during observers' saccades, he changed the sizes or colours of objects in a scene, swapped the locations of two objects, removed various foreground elements from the scene, and so on. This inability to detect many kinds of changes made during saccades occurs with natural scenes, simple geometric stimuli, and even written text (e.g. Bridgeman, Hendry, & Stark, 1975; Currie, McConkie, Carson-Radvansky, & Irwin, 1995; Grimes, 1996; Henderson, 1997; McConkie & Currie, 1996; McConkie & Zola, 1979).

Studies like these, which explore change detection across different views of scenes, provide a window on the underlying processes of visual perception and attention, and suggest that our internal representations of scenes are far less detailed and complete than is suggested by our everyday visual experience. Such sparse representations and the corresponding change blindness may be crucial for our perception of a stable visual world across saccades: "change blindness supports the phenomenal experience of continuity by not preserving too much information from one view to the next" (Simons & Levin, 1997, p. 267).

CHANGE BLINDNESS IN THE FLICKER PARADIGM

Recent experiments by Ron Rensink and his colleagues, however, have suggested that blindness for scene changes made during saccades may not be due to anything specific about the saccadic system, but rather to the fact that saccades (among many other manipulations) mask the motion transients associated with scene changes (Rensink et al., 1997). Rensink and his colleagues thus predicted and observed change blindness in other situations where scene changes were *not* constrained to occur only during saccades.

The primary tool these researchers have used to investigate change blindness is the *flicker paradigm*, depicted in Figure 1.[1] In the flicker paradigm, two

[1] Although this paper focuses on the flicker paradigm, change blindness has also been observed when the motion transients of the scene changes are masked by blinks (O'Regan, Deubel, Clark, & Rensink, this issue), "mudsplashes" (O'Regan, Rensink, & Clark, 1999), film-cuts (Levin & Simons, 1997), and other manipulations (for reviews see Simons & Levin, 1997 and Simons, this issue). Change blindness has even been observed in real-world occlusion events (Simons & Levin, 1998), demonstrating that the basic phenomenon is not an artifact of the passive viewing of ecologically invalid 2D displays, as in the experiments reported in this paper.

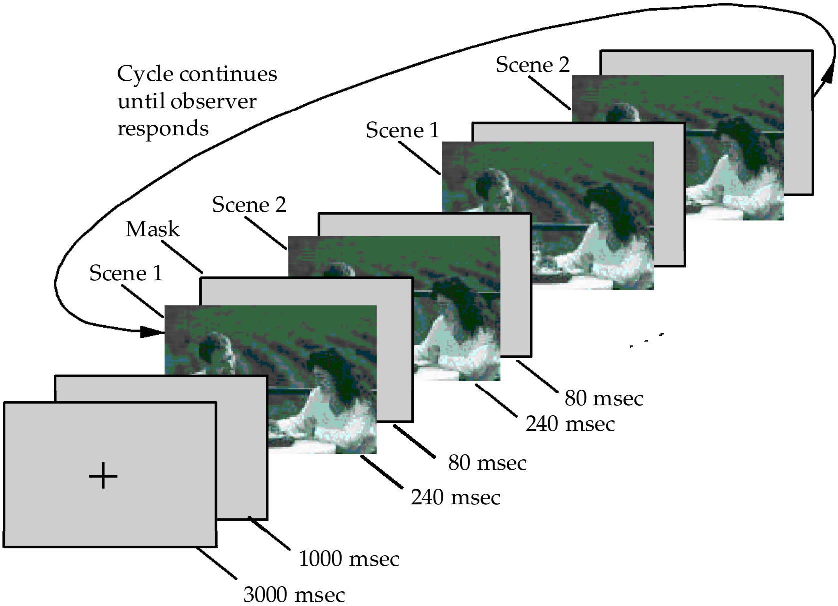

FIG. 1. A sample trial from a "flicker" experiment, after Rensink et al. (1997). After an initial fixation and grey mask, a flicker sequence begins in which two images are alternately displayed (for 240 msec each), separated by a grey mask (displayed for 80 msec). This flicker sequence is repeated until subjects detect the change between the two scenes—in this case, the movement of the railing. (Note that in some experiments, each scene is displayed twice before moving on to the other scene, in order to create some temporal uncertainty; this is not depicted here.) (Modified from Rensink et al., 1997; the depicted scenes were created and made available by Ron Rensink.)

scenes are alternated back and forth, with a homogenous grey mask in between each, to mask the motion transients caused by changes between the scenes. Rensink et al. (1997) presented photographed natural scenes in this manner, with each scene displayed for 240 msec, followed by an 80 msec mask. The two scenes in each trial were identical, except for a single object or location which changed colour, changed location, or was deleted altogether from one scene to the next. (In some experiments, each image was "flickered" twice in a row before displaying the other image to create some temporal uncertainty about when the change was being made; this is not represented in Figure 1.) Observers freely viewed these alternating scenes, and were to press a key as soon as they could identify the change. (See Blackmore, Brelstaff, Nelson, & Troscianko, 1995, for a similar method.)

This flicker paradigm is similar in some respects to earlier studies which assessed change detection using only two individual scenes (presented only once each) separated by a mask or a blank field, with accuracy as a dependent measure (e.g. Pashler, 1988b; Phillips, 1974; cf. Simons, 1996). These earlier experiments, however, may not have displayed the scenes long enough for observers to build up the requisite representations. In the flicker paradigm, in contrast, the change and the scenes themselves are presented continuously until the observer responds.

Even so, observers took surprisingly long (on the order of tens of seconds) to identify (i.e. to *see*) even very large, repeated changes in this paradigm (Rensink et al., 1997). The phenomenology of this effect is quite striking: Once a change is seen, it becomes extremely obvious, and it seems almost inconceivable that it could have been missed only moments before. (Observers in the flicker experiments reported here often initially refused to believe that the change had been there from the beginning of the "flickering", suspecting instead that it had been added only at the moment in which they noticed it!) This change blindness persisted even with longer presentation times (e.g. 560 msec per scene). Change detection was much faster, however, when observers were simply told where to look for the changes (thus ruling out a simple explanation in terms of poor visibility), and also when the intervening grey mask fields were removed (thus reintroducing the motion transients).[2]

ATTENUATED CHANGE BLINDNESS AT CENTRES OF INTEREST: THREE PROBLEMS

Another central result from the Rensink et al. (1997) experiments was that change blindness was attenuated when the change was made at *attended* locations or objects. Rensink et al. (1997) interpret these results in terms of an "attentional gating theory", the primary claim of which is that *attention is necessary to perceive changes in scenes* (Rensink et al., 1997, pp. 368, 372; see also Rensink, this issue a):[3]

> [T]he visual perception of a change in a scene occurs only when focused attention is given to the part being changed. ... [T]hese results indicate that—even when sufficient viewing time has been given—an observer does not build up a representation of a scene that allows him or her to perceive change automatically. Rather, perception of change is mediated through a narrow attentional bottleneck. ... Attended items are loaded into a durable store and are perceived to undergo transformation whenever

[2] O'Regan et al. (1999) reported similar experiments in which the mask was presented in the form of "mudsplashes" which suddenly appeared on the screen. Change blindness was observed in this case even though the mask did not actually cover the locations of the changes, suggesting that change blindness in the flicker paradigm is not due simply to interrupted processing at the location of the change.

[3] Other non-attentional systems may still register the changes on this theory, and it may be possible to have some form of access to this information (i.e. *sensing* that there is a change, but not being able to identify or locate it), even without attention, before the change is actually *seen* (Fernandez-Duque & Thornton, this issue; Rensink, 1998).

they are changed, whereas unattended items are simply replaced by the arrival of new items, with no awareness that replacement has occurred.

The "attended items" in these experiments, however, were assessed only by a measure of the "centres of interest" in the scenes, defined as those objects or areas mentioned by at least three out of five additional observers who were simply asked to describe the (static) scenes. Changes to these centres of interest tended to be detected several seconds faster than changes to other objects or locations.

There are at least three problems associated with this measure of assessing attention. First, there is a worry that describability may not perfectly correlate with visual interest and attention. Verbal descriptions of scenes are mediated by several other factors, notably including the relative ease of description: Some objects or areas may be mentioned simply because they are easily described, whereas others may not be mentioned because they are difficult to put into words. Some popular change-blindness demonstrations, for instance, employ changes that can be described no more elegantly than: "A blob-like area of the field in the background, defined by the area inside a bit of a tree-branch, the horizon, and the side of the barn, which is changing colour". Such an area may be visually interesting, despite the fact that it eludes ready description. Factors such as describability may thus affect verbal reports of scenes in ways that have no impact on the control of visual attention.[4] Second, other researchers have doubted the explanatory utility of this notion of attention. Zelinsky (1998, p. 38), for instance, suggests that:

> [A]n appeal to attentional selection as the cause of change detection failure is tantamount to invoking the homunculus. In order to be a useful explanation, theorists must specify why attention should be preferentially directed to some objects but not others. ... Rensink ... answered [this question] by saying that attention is drawn to centers of interest in a scene. Note however that calling something a center-of-interest essentially just redefines an object of attention and adds little to the understanding of how this object becomes an attentional attractor.

Finally, this attention-based theory deserves to be investigated further because of an odd contrast with many other experiments in which observers exhibit change blindness even in those cases where it seems that they *would* be purposely attending to the locus of the change (e.g. Ballard, Hayhoe, & Pelz,

---

[4]Even more generally, both visual factors (e.g. size) and high-level factors (e.g. overall scene consistency) may contribute to centres of interest, and are not pulled apart in the Rensink et al. (1997) experiments. Hollingworth and Henderson (this issue) looked at semantic factors, and found that scene-inconsisten t changes were noticed faster than scene-consistent changes in a flicker paradigm. The implications of these results remain unclear, however, given that their effect magnitudes were extremely small (on the order of 50 msec when reaction time was the dependent measure).

1995; Levin & Simons, 1997; Simons, 1996; Simons & Levin, 1998). (As noted by Simons & Levin, 1997, however, this is not inconsistent with a theory that attention is necessary but not sufficient for change detection.)

The experiment reported here attempts to test the attention-based theory of change blindness by determining whether attenuated change blindness at attended objects holds even for other types of *exogenous* attentional selection that are controlled more rigorously.

ENDOGENOUS CONTROL VS. EXOGENOUS CAPTURE OF ATTENTION

There is a central distinction in the study of visual attention between (a) voluntarily directing attention to some object or location, and (b) having one's attention be involuntarily captured by some salient aspect of a scene. In current parlance, this is the distinction between the *endogenous control* of attention and the *exogenous capture* of attention, respectively.[5] Among other differences, these two types of attentional selection have different timing profiles, with exogenous capture being relatively fast and transient, and endogenous control being relatively slow and sustained. For reviews of research on this fundamental distinction, see Egeth and Yantis (1997) and Theeuwes (1994). Previous studies of change blindness (e.g. Rensink et al., 1997) have exploited endogenous attentional selection, as described above. The experiment reported here, in contrast, employs two sorts of manipulations of exogenous attention: Sudden onsets and colour singletons.

Among the phenomena that exogenously capture attention are *sudden onsets*—that is, the sudden appearance of a stimulus where none was before (e.g. Burkell & Pylyshyn, 1997; Jonides, 1981; Jonides & Yantis, 1988; Theeuwes, 1991b; Yantis & Jonides, 1984). Yantis and Jonides (1984), for example, tested observers on a visual search task in which the display began as an array of figure-eights (as on a digital alarm clock). Each of the figure-eights then dropped various segments to become various letters; at the same time a new letter was displayed in a new location (i.e. a location not previously occupied by a figure-eight). This late-onset item was no more likely to be the target letter than the other items, and observers knew this. Nevertheless, even though there was no incentive to attend to the late-onset letter, response times to find the target letter when it did happen to be the late-onset item were fast and did not vary with the number of elements in the display.

[5]Also variously known as the distinction between *central* and *peripheral* cueing of attention; between *goal-driven* and *stimulus-driven* attentional selection; between *top-down* and *bottom-up* control of attention; between what William James (1890) called *active* and *passive* attention. This basic distinction has been implicitly heeded for the duration of scientific psychology, and probably as far back as the start of the previous millennium (see Hatfield, 1998).

Another phenomenon which has been thought to exogenously capture attention under some conditions is that of *featural singletons*—that is, the presence of a unique feature in a display, such as a red item in a field of black items. In addition to "popping out" of search displays when they serve as targets, the presence of colour singletons (and featural singletons more generally) has been shown to slow response times to other targets in various tasks even when the singletons were completely irrelevant, suggesting that they were involuntarily and necessarily attended (e.g. Folk, Remington, & Johnston, 1992; Pashler, 1988a, exp. 7; Theeuwes, 1991a, 1992; Theeuwes & Burger, 1998; Todd & Kramer, 1994). Theeuwes (1992), for example, employed displays consisting of arrays of oriented lines, each of which was inside either a diamond or a circle. The observers' task was simply to report the orientation (horizontal or vertical) of the single non-oblique line. The target was always inside a form singleton (i.e. the single square in a field of circles), and colour was completely irrelevant to the task. Nevertheless, response times were lengthened when the display contained a colour singleton—for example, a single green circle in a field of black shapes—even though the task was essentially to attend to the form singleton and ignore the colour singleton. This inability to ignore the colour singleton was quite robust, and occurred despite full knowledge of the irrelevancy of colour, and despite extensive practice.

The claim with regard to late onsets and featural singletons is that in certain circumstances they will exogenously capture visual attention—not that they must always do so, in all circumstances. It has been shown, for example, that exogenous capture by late onsets can be attenuated when, in anticipation of a target, observers tightly focus their attention on another distinct location (e.g. Yantis & Jonides, 1990), or featural dimension (e.g. waiting for a colour singleton; Folk et al., 1992). In addition, it may be that certain sorts of onsets do not capture attention when they do not mark the appearance of a new perceptual object (Yantis, 1993a), for example when an item reappears from behind an occluder (Scholl & Pylyshyn, 1999). Similarly, colour singletons are thought to exogenously capture attention only in certain circumstances. Capture by colour singletons in visual search paradigms is particularly attenuated when the target of the search does not involve a singleton of any kind (e.g. Hillstrom & Yantis, 1994; Jonides & Yantis, 1988). The conclusion from these studies seems to be that "when the target of search is a featural singleton, any salient singleton (even ones known to be irrelevant) captures attention, and when the target of search is not a singleton, irrelevant singletons do not capture attention" (Yantis, 1993a, p. 158; see also Bacon & Egeth, 1994). Capture by colour singletons is also attenuated during some types of serial search when observers know precisely the properties (e.g. the specific colours) of both the target and distractor singleton on every trial (Theeuwes & Burger, 1998).

Note that while exogenous capture by late onsets and featural singletons remains controversial in these sorts of circumstances, nearly all of the researchers in this area agree that capture by such stimuli occurs in *some* conditions. It seems best to characterize these conditions simply as the lack of attentional engagement elsewhere (or "elsehow", in the case of an "attentional control setting" for another featural dimension). In change detection experiments using the flicker paradigm, it seems plausible that observers do not immediately set out on a serial search of the display, but first try to "sense" the change across the display as a whole. Such circumstances seem ideal for exogenous attentional capture, which should be detectable as attenuated change blindness to exogenously attended items.

CHANGE BLINDNESS AND EXOGENOUS ATTENTION: THE CURRENT EXPERIMENT

In the flicker experiment reported here, one item in the first scene of each trial is manipulated to capture visual attention exogenously. This is done here in two ways, using arrays of line-drawings as (quite unnatural) scenes: (a) Attention is drawn to a late-onset item in the initial display, before the "flickering" begins; (b) Attention is drawn to a colour singleton in the initial display. These manipulations—a late-onset item or a colour singleton—are never useful as cues to the location of the change: The actual change (either a replacement or a change in orientation) is equally likely to occur anywhere, completely independently of the exogenously attended item, and observers know this. Nevertheless, it is predicted that change blindness will be attenuated when changes do happen to be made to the exogenously attended items. The idea is that attention is involuntarily captured by the late-onset items and colour singletons, even though such attentional capture is not useful. When a change is made to such an item, it will be attended, and therefore more readily seen. These manipulations will test the attention-based theory of change blindness (Rensink et al., 1997), using a more rigorously understood type of attentional selection.

In this experiment, "centres of interest" are factored out by employing scenes composed of arrays of small line-drawings, each depicting a relatively long and narrow object, oriented diagonally (e.g. a carrot, a pen, a broom, a fork). A scene change always consisted of either one object being replaced by another (of the same orientation)—see Figure 2a—or a flip in the orientation of the object (maintaining the overall axis of orientation)—see Figure 2b. Other recent experiments confirm that change blindness is observed in such scenes, composed of arrays of simple geometric shapes (Intriligator, He, & Barton, 1998) or rendered photographs of long, thin, diagonally oriented items (Zelinsky, 1998). This use of such contrived, unnatural scenes has the effect of lowering overall response times to detect changes, because the possible change locations are confined to a small discrete number of simple items, rather than the thousands of possible locations and objects in a natural scene. This method is crucial in the current context, however, in that it factors out centres of interest

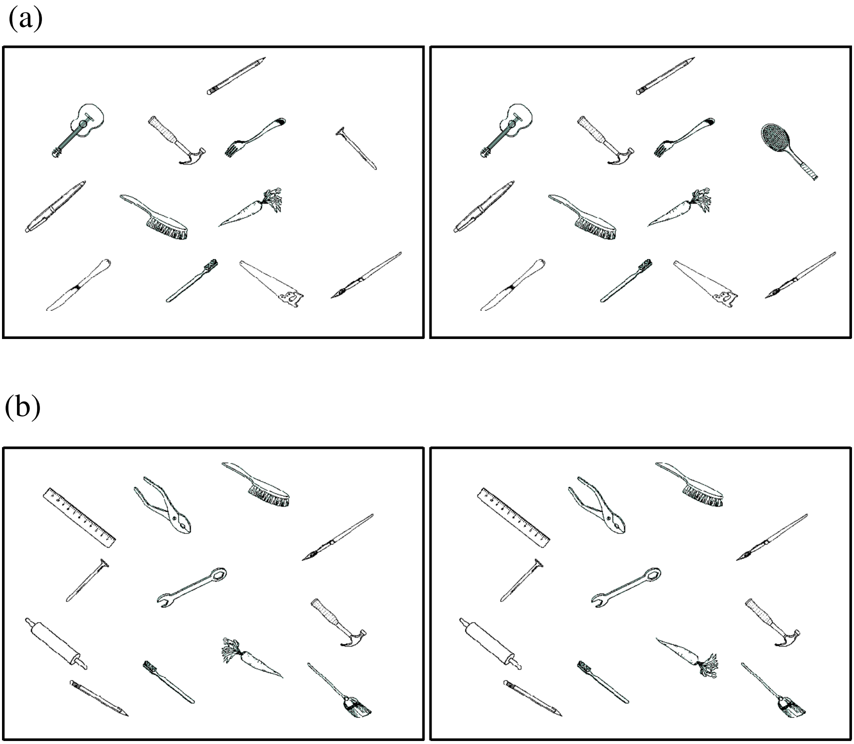

FIG. 2. Examples of the two types of scene changes. (a) Depicts a *replacement change*: A nail changes into a tennis-racket. (b) Depicts a *flip change*: The carrot's orientation is flipped both vertically and horizontally.

(i.e. endogenous attention) and other high-level semantic factors that affect change blindness (e.g. Friedman, 1979; Hollingworth & Henderson, this issue; Rensink et al., 1997), by eliminating any meaningful scene composition.

This experiment employs three main stimulus conditions (each seen by a different group of observers). In the *late-onset condition*, attention is exogenously captured via late onsets, as described previously. Rensink et al.'s (1997) flicker paradigm was altered such that the initial frame appeared in two steps, with one item being displayed 200 msec later than the other 11 items. See Figure 3 for a depiction of this method. It is predicted that these late-onset items will capture visual attention, and thus that response times to detect changes to these late-onset items will be faster, even though the changes are no more likely to involve these items than the others, and observers know this.

The *colour-singleton condition* is identical in all respects to the late-onset condition (using all the same trials), except that the previously late-onset item in each trial (which is now drawn along with the others) is now made into a

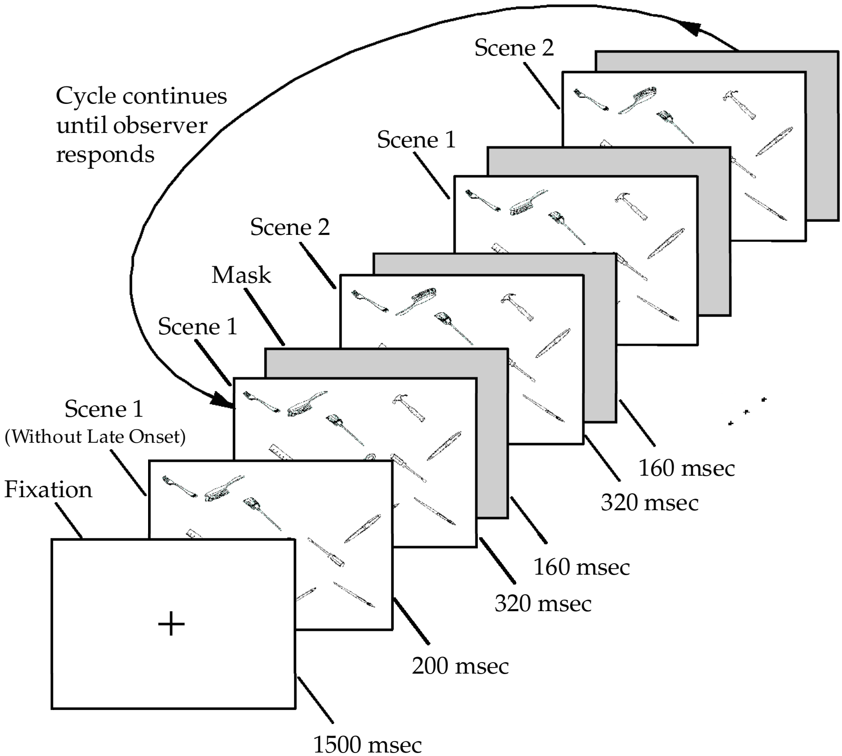

FIG. 3. A sample trial from the "late onset" condition. After an initial fixation, the first scene is presented—missing one item. After 200 msec, the missing item (in this case a hammer) is onset, thus beginning a flicker sequence in which two scenes are alternately displayed (for 320 msec each), separated by a grey mask (displayed for 80 msec). This flicker sequence is repeated until subjects detect the change between the two scenes— in this case a brush that is "flipping".

colour singleton (i.e. it is drawn in blue or red, in a field of black items). This colour singleton was no more likely to be the changed item than any of the others (and observers knew this), but it did happen to be the changed item exactly 1/12 of the time (where the displays again consisted of 12 items). Nevertheless, attenuated change blindness—that is, lower response times to detect changes—is predicted in trials in which a colour singleton was the changed item.

Finally, a *control condition*, employing trials that do not contain late onsets or colour singletons but are otherwise identical, is used to confirm that attenuated change blindness in the other conditions is due to the late-onset and colour-singleton manipulations, and not to any other haphazard differences between the two classes of trials.

METHOD

Participants

Thirty naive observers (10 per condition) participated in one 40-min session. All observers had normal or corrected-to-normal acuity.

Apparatus

The displays were presented on a monitor controlled by a Power Macintosh microcomputer. Observers were positioned approximately 46 cm from the monitor, such that the display subtended approximately 39 × 28 degrees of visual angle. The displays were presented using the *RSVP* experimentation package (Williams & Tarr, 1998).

Materials

Scenes. Each condition employed the same set of 120 scenes. Each scene was composed of an array of line drawings of simple objects. Each item subtended approximately 6 degrees, was relatively thin, oriented diagonally, and was presented in black on a white background. The actual line drawings were taken from the set presented in Snodgrass and Vanderwart (1980).[6] The initial scene of each trial consisted of 12 items (randomly chosen from the set of 20 listed in note 6), randomly placed about the screen independently on each trial. The diagonal orientation of each item in each trial was chosen randomly from a possible set of four directions (facing "northeast", "northwest", "southeast", or "southwest").

Changes. The second scene on each trial was constructed by replacing or "flipping" one of the 12 items, chosen at random. Replacement objects were drawn randomly from the unused set of eight line drawings for each trial, and appeared in the same orientation as the replaced object. "Flips" were accomplished simply by reflecting the changed item about both its vertical and horizontal axes (thus preserving the overall axis of orientation). Note that these changes were designed to be fairly subtle, in order that response times be fairly high even though this experiment does not use natural scenes. See Figure 2 for examples of these changes.

Attentional Selection Manipulations. In the late-onset condition, one item on each trial was late-onset as described next. This late-onset item was chosen at random, independently of the changed item; as such, the late-onset item was

[6]The 20 line drawings from Snodgrass and Vanderwart (1980) used in this experiment were: The axe, broom, brush, carrot, fork, guitar, hammer, knife, nail, paintbrush, pen, pencil, pliers, rolling-pin, ruler, saw, screwdriver, tennis racket, toothbrush, and wrench.

the changed item on exactly 10—1/12—of the 120 total trials. In the colour-singleton condition, this same item (which is now drawn along with the others) was made into a colour singleton: It was displayed in bright red on half of the trials and in bright blue on the other half. (All other items in the experiment were displayed in black.) Since the same set of scenes was used in each condition, the colour singleton and the changed item again corresponded on exactly 10 of 120 trials. The control condition employed the same set of trials as in the other two conditions, except that there were no late onsets or colour singletons.

Procedure

A single trial in the late-onset condition proceeded as follows (see Figure 3). Observers initiated each trial by pressing a key, which blanked the screen. After 500 msec, a fixation cross appeared for 1500 msec. Eleven of the twelve items in the first scene then appeared, followed 200 msec later by the appearance of the late-onset item; 320 msec later, the "flickering" began. This consisted of repeating the following sequence of events: (a) Presenting the uniform grey mask for 160 msec, (b) presenting the second (changed) scene for 320 msec, (c) presenting the grey mask for 160 msec, and (d) presenting the original scene for 320 msec. This sequence of events was repeated until observers responded. Trials in the colour-singleton and control conditions were presented using this same method, except that there were no late onsets: All of the items in a scene appeared simultaneously, and the flicker began immediately, after the initial display of the first scene for 320 msec (see Figure 1). Observers were to press a key on the keyboard as soon as they detected the change; while pressing the key, observers also said aloud the name of the changed item, which the experimenter checked for accuracy. The observers' response times were recorded on every trial, as indexed by the keypress; the response timer began upon the first presentation of the second (i.e. changed) scene.

At the beginning of an experimental session, observers were instructed as to the types of changes possible, and were familiarized with seven practice trials. Observers in the late-onset and colour-singleton conditions were informed that the late-onset items or colour singletons were no more likely to be the changed items than any of the others (in other words, that the two were completely uncorrelated), such that the existence of the late-onset item or colour singleton was completely irrelevant to the task. The actual experiment followed, consisting of 120 trials. Sixty of these trials were "replacement" changes, and 60 were "flip" changes (randomly ordered).

Results

Since observers knew that a change was made on each trial, errors consisted only of misidentifying the item that was changed. As in other change blindness studies employing the flicker paradigm, these error rates were extremely low—on average, 0.05% across all three conditions. These accuracy data were used to ensure that observers were correctly engaging the task, and were not analysed further. The few trials on which errors were made were not included in the response time analyses below.

Mean response times are presented in Figure 4 for all three conditions, broken down by whether or not the changed item was an attentionally selected item (i.e. a late-onset item or a colour singleton) or a black normal-onset item. (Of course, such categories have no meaning for the control condition. Nevertheless, since exactly the same set of trials were used, these sub-sets were compared just as in the other conditions, to ensure that any differences were actually due to the late-onset and colour-singleton manipulations.) Performance in trials of each individual change type (flips and replacements) is discussed later; these

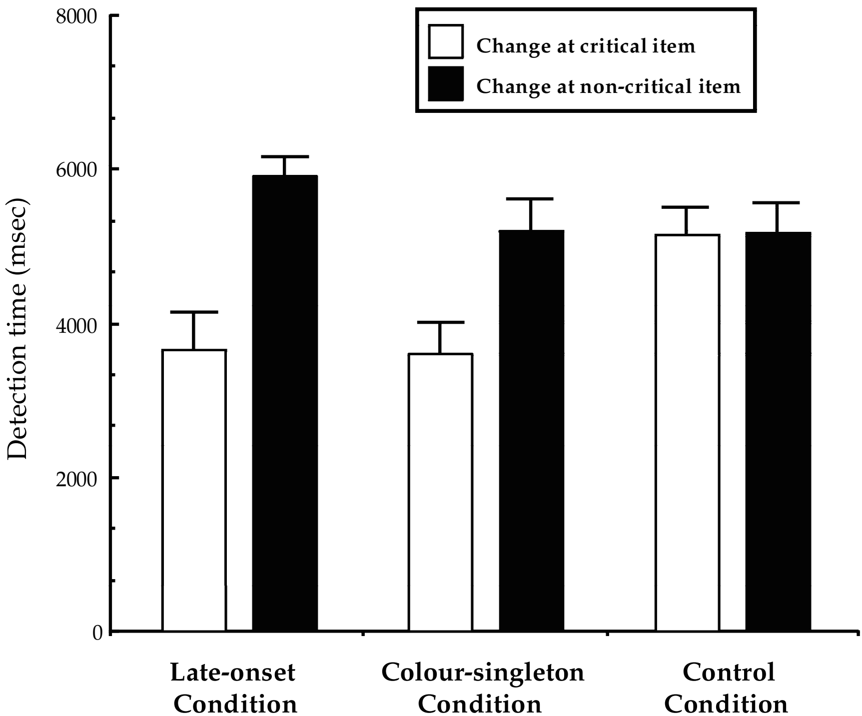

FIG. 4. Mean detection times in msec for all conditions, with standard errors. In the late-onset condition, the "critical items" were the late-onset items; in the colour-singleton condition, they were the singletons. Such categories have no meaning for the control condition. Nevertheless, since exactly the same trials were used, these sub-sets were compared just as in the other conditions, to ensure that any differences were actually due to the late-onset and colour-singleton manipulations.

trials individually could not be directly compared, however, since discriminability for these types of changes was not equated.

A two-way mixed-design analysis of variance (ANOVA) was performed on the mean correct response times with condition type (late-onset vs. colour-singleton vs. control) as the within-subjects factor and trial type (normal item change vs. late-onset or singleton item change) as the between-subjects factor. There was a significant main effect of trial type, $F(1,27) = 27.14$, $p < .01$, but not of condition type, $F(2,27) = 1.20$, $p > .3$. In addition, there was a significant trial type by condition type interaction, $F(2,27) = 7.03$, $p < .01$.

Planned comparisons were then performed within each condition. In the late-onset condition, response times were significantly faster on the trials in which the changed item was the late-onset item, compared to those trials on which the changed item was a normal-onset item, $t(9) = 4.11$, $p < .01$. This pattern of results was replicated for both types of changes individually, and was the case for 9 of the 10 observers. The magnitude of this difference was large: Late-onset changes were detected 2.2 sec faster, on average—nearly half the overall mean response time. The fact that the mean response time for the late-onset change trials (3.6 sec) was faster than the overall mean response time in the control condition (5.1 sec) suggests that the difference in change detection between the late-onset items and the other items was due to a speed-up of detection for late onset items, and not a slowing of detection for normal-onset items.

In the colour-singleton condition, response times were significantly faster on the trials in which the changed item was a colour singleton, compared to those trials on which the changed item was a normal-coloured item, $t(9) = 3.43$, $p < .01$. This pattern of results was replicated for both types of changes individually, and was the case for 9 of the 10 observers. This effect magnitude was also large: Changes at colour singletons were detected an average of 1.6 sec faster. Again, the fact that the mean response time for the colour-singleton change trials (3.6 sec) was significantly faster than the overall mean response time in the control condition (5.1 sec) suggests that the difference in change detection between colour singletons and other items is due to a speed-up of detection for singletons, and not a slowing of detection for normal items.

Comparing these identical sub-sets in the control experiment revealed no significant difference, $t(9) = .14$, $p > .5$, ensuring that the differences in the late-onset and colour singleton conditions were due to the attentional selection manipulations, and not to any other haphazard differences between these subsets of trials. This pattern was replicated for "flips" by themselves, $t(9) = 1.93$, $p > .05$, but replacements were detected significantly slower for those trials in the control condition which were late-onset or colour singletons in other conditions, compared to the other trials, $t(9) = 3.55$, $p < .05$. This suggests that the effect sizes in the other conditions may even be underestimates of the actual benefits of exogenous attentional selection.

DISCUSSION

In other experimental paradigms, late onsets and colour singletons have been shown to exogenously capture visual attention in certain conditions, in some cases even when the singleton or late onset is irrelevant to the task (Burkell & Pylyshyn, 1997; Folk & Remington, 1998; Folk et al., 1992; Jonides, 1981; Jonides & Yantis, 1988; Pashler, 1988a, exp. 7; Theeuwes, 1991a, 1991b, 1992; Theeuwes & Burger, 1998; Todd & Kramer, 1994; Yantis & Jonides, 1984). In this experiment, change blindness was attenuated when the changed item was late onset or was a colour singleton, even though such manipulations were completely uncorrelated with the location of the change. This suggests that these manipulations captured attention in the flicker paradigm (Rensink et al., 1997) and that changes to late-onset items and colour singletons were noticed faster because they were being attended.[7] Before discussing the implications of these results, I address in more detail the interplay between exogenous and endogenous attention in this experiment.

IMPLICATIONS FOR THE NATURE OF EXOGENOUS AND ENDOGENOUS ATTENTIONAL SELECTION

Note that it is possible that the exogenous capture of visual attention was only very briefly in effect in this experiment, followed immediately by endogenous attentional inspection of the scenes. This would be consistent with the relatively fast and transient timing profile of exogenous attentional capture, and also with the relatively slow response times obtained here. (It is not possible to draw solid implications from the overall response times themselves, however, since it is unclear how to factor them into their component stages of attentional selection, change detection, change verification, naming, response selection,

[7] One might also ask: Where does attention "go" after it leaves the late-onset item or colour singleton, in those many cases where the change is elsewhere? It seems plausible that the subsequent "attentional scan path" might travel from such items to whichever additional items are in close proximity. In other experiments we address such questions fully (Scholl, Pylyshyn, & Chen, in preparation), but a preliminary analysis here is suggestive. The 120 trials were first divided up into three groups: (a) Those in which the changed item was attentionally selected (10 trials, analysed previously), (b) *adjacent trials*, in which the changed item was roughly immediately adjacent to the attentionally selected item (36 trials), and (c) *far trials*, in which the changed item was neither the attentionally selected item nor adjacent to it. Response times to the adjacent trials were significantly faster than to the far trials in both the late-onset condition, $t(9) = 3.22$, $p < .05$, and the colour-singleton condition, $t(9) = 3.99$, $p < .05$, but not in the control condition, $t(9) = 1.41$, $p > .1$. These effect magnitudes (i.e. the far–adjacent differences) were noticeably smaller than the effect magnitudes in the main experiment (i.e. the attentionally-selected vs. unselected differences)—by 1393 msec for late-onsets and 692 msec for colour singletons—hinting at an attentional scan path rather than a "spotlight" of attention that simply spanned more than one item. See Scholl et al. (in preparation) for further discussion.

etc. The response times observed here may also have been lengthened when subjects verbally identified the change before pressing the response key; cf. Scholl et al., in preparation, where this was controlled.) In this scenario, the exogenous attentional capture would serve as an initial trigger for endogenous attention, essentially determining the beginning of the endogenous scan path (cf. Todd & Kramer, 1994). (As in other studies of exogenous capture, it seems unlikely that the late-onset items and colour singletons received much unmediated endogenous attention throughout the 120 trials, given their known irrelevancy to the task.)[8]

Rensink (this issue b) has stressed the usefulness of the flicker paradigm as a tool which can be used not just to demonstrate change blindness, but also to explore the mechanics of visual attention (see also note 7). Attenuated change blindness was taken in this experiment as an operationalization of exogenous attentional capture, and the results of this experiment, in turn, have implications for our understanding of how such attentional capture works.[9] In this vein, note that one aspect of the present results seems to be at odds with the received wisdom regarding exogenous attentional capture by colour singletons. Experiments involving a standard visual search paradigm have suggested that colour singletons will exogenously capture attention only when the target of the search is a singleton (of any type; for reviews, see Egeth & Yantis, 1997, and Theeuwes, 1994). In other words: "There is no evidence that any attribute other than onset will capture attention in the absence of a deliberate attentional set for singletons" (Yantis, 1993b, p. 679). The present experiment, however, constitutes just this type of evidence! In the flicker experiment reported here, colour singletons captured attention and attenuated change blindness even though this change detection paradigm did not engender an attentional set for singletons of

---

[8]An anonymous reviewer suggests that an increasing role for endogenous attention throughout the experiment might be manifest in faster detection of the critical item changes on early trials (i.e. the first five trials on which the changed item is the late-onset item or colour singleton) compared to late trials (i.e. the second and last five). This was the case in neither the late-onset, colour-singleton, or control conditions (all *p*s > .4). This also suggests that associated confounds involving practice effects and/or fatigue effects did not affect these results.

[9]This has been seen by some as courting circularity: I am arguing that onsets attract attention and thus affect a change detection experiment by attenuating change blindness; but my only evidence that attention is drawn to the onsets *is* attenuated change blindness! Such a relation is intrinsic to the process of operationalization, however. Indeed, the same concern could be articulated with regard to the analogous visual search experiments (e.g. Jonides & Yantis, 1988; Yantis & Jonides, 1984): "You are arguing that onsets attract attention and thus affect a visual search experiment by speeding search or flattening search slopes; but your only evidence that attention is drawn to the onsets *is* the speeded search or the flattened search slopes!" The converging evidence from these paradigms (and others; e.g. Burkell & Pylyshyn, 1997) strongly supports the operationalization of attentional capture. In the current experiment, the attenuated change blindness must itself be explained! The attention-based theory does explain these results (Rensink et al., 1997), and no other explanation seems forthcoming.

any sort. One possibility here is that the traditional notion of a featural singleton must be expanded to encompass the "flips" and "replacements" at work as changes in this experiment. Another possibility is that exogenous capture by featural singletons works differently in different tasks, and that the previous results will only generalize to standard visual search and spatial cueing tasks.

SUMMARY

The experiment reported here demonstrates attenuated change blindness for exogenously attended items. Rensink et al.'s (1997) flicker paradigm was used with scenes composed of arrays of simple line drawings, where one item on every trial was either replaced with another item, or was flipped about both its horizontal and vertical axes. One item on each trial also exogenously captured attention (via a colour singleton or a late onset), but this item was no more likely to be the changed item than any other, and observers knew this. Nevertheless, changes to these late-onset items and colour singletons were reported significantly faster than other items, by an average of 2.2 sec for late onsets and 1.6 sec for colour singletons. These differences are striking, given that the overall mean response time in the control condition was only 5.1 sec. These results are explained in terms of the exogenous capture of visual attention by the colour singletons and late-onset items (despite their irrelevancy to the task), along with the attention-based theory of change blindness (Rensink et al., 1997), in which only attended items are represented in a short-term visual store, and so can be compared from one scene to another.[10]

Unlike the "centre of interest" method of assessing attentional selection (O'Regan et al., this issue; Rensink et al., 1997), however, this method of *exogenous attentional capture* does not engender possible confounds with the ease by which scenes can be verbally described. Also unlike the "centre of interest" methods, exogenous capture does not raise worries of explanatory impotence. Zelinsky's (1998) complaints, for example—that "theorists must specify why attention should be preferentially directed to some objects but not others" and that "calling something a center-of-interest essentially just redefines an object of attention and adds little to the understanding of how this object becomes an attentional attractor" (p. 38)—are defused in the experiment reported here. By employing exogenous capture, an operationalization of attentional selection is linked to a rich body of experiment and theory that *does* specify why and when attention is preferentially directed to certain objects and locations.

This experiment thus supports the attention-based theory of change blindness, using a form of attentional selection that is more rigorously controlled and

[10] I have replicated this pattern of results in other recent pilot experiments using other experimental paradigms, such as a single display of each scene separated by a blank screen, using accuracy as a dependent measure (as, for example, in Simons, 1996).

understood. The advantages of this method of attentional selection can be harnessed in many other ways to refine our characterization of how visual attention interacts with change detection and change blindness. In other experiments not reported here, for example (Scholl et al., in preparation), this method is used to explore (a) whether change blindness is attenuated at *multiple* exogenously attended items, (b) whether attenuation of change blindness by exogenous capture *detracts* from attenuation by *endogenous* factors, and (c) whether the attenuation of change blindness at an exogenously attended item in a natural scene will *spread* to nearby locations or objects.

REFERENCES

Bacon, W., & Egeth, H. (1994). Overriding stimulus driven attentional capture. *Perception and Psychophysics, 55,* 485–496.

Ballard, D.H., Hayhoe, M.M., & Pelz, J.B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience, 7,* 66–80.

Blackmore, S.J., Brelstaff, G., Nelson, K., & Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception, 24,* 1075–1081.

Bridgeman, B., Hendry, D., & Stark, L. (1975). Failure to detect displacements of the visual world during saccadic eye movements. *Vision Research, 15,* 719–722.

Burkell, J., & Pylyshyn, Z.W. (1997). Searching through selected subsets of visual displays: A test of the FINST indexing hypothesis. *Spatial Vision, 11,* 225–258.

Currie, C., McConkie, G., Carlson-Radvansky, L., & Irwin, D. (1995). *Maintaining visual stability across saccades: Role of the saccade target object* (Tech. Rep. No. UIUC-BI-HPP-95-01). Champaign, IL: University of Illinois, Beckman Institute.

Egeth, H.E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology, 48,* 269–297.

Fernandez-Duque, D., & Thornton, I.M. (this issue). Change detection without awareness: Do explicit reports underestimate the representation of change in the visual system? *Visual Cognition, 7,* 323–344.

Folk, C.L., & Remington, R. (1998). Selectivity in distraction by irrelevant featural singletons: Evidence for two forms of attentional capture. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 847–858.

Folk, C.L., Remington, R.W., & Johnston, J.C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 1030–1044.

Friedman, A. (1979). Framing pictures: The role of knowledge in automatized encoding and memory for gist. *Journal of Experimental Psychology: General, 108,* 316–355.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.), *Perception* (pp. 89–109). New York: Oxford University Press.

Hatfield, G. (1998). Attention in early scientific psychology. In R.D. Wright (Ed.), *Visual attention* (pp. 3–25). New York: Oxford University Press.

Henderson, J.M. (1997). Transsaccadic memory and integration during real-world object perception. *Psychological Science, 8,* 51–55.

Hillstrom, A., & Yantis, S. (1994). Visual motion and attentional capture. *Perception and Psychophysics, 55,* 399–411.

Hollingworth, A., & Henderson, J.M. (this issue). Semantic informativeness mediates the detection of changes in natural scenes. *Visual Cognition, 7,* 213–235.

Intriligator, J., He. S., & Barton, J. (1998). Detection of object-change is more accurate in the lower visual field. *Investigative Ophthalmology and Visual Science, 39,* S852.

James, W. (1890). *The principles of psychology*. New York: Holt.

Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. In J.B. Long & A.D. Baddeley (Eds.), *Attention and performance IX* (pp. 187–203). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Jonides, J., & Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention. *Perception and Psychophysics, 43,* 346–354.

Levin, D.T., & Simons, D.J. (1997). Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin and Review, 4,* 501–506.

McConkie, G., & Currie, C. (1996). Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 563–581.

McConkie, G., & Zola, D. (1979). Is visual information integrated across successive fixations in reading? *Perception and Psychophysics, 25,* 221–224.

O'Regan, J.K., Deubel, H., Clark, J.J., & Rensink, R.A. (this issue). Picture changes during blinks: Looking without seeing and seeing without looking. *Visual Cognition, 7,* 191–211.

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

Pashler, H. (1988a). Cross-dimensional interaction and texture segregation. *Perception and Psychophysics, 43,* 307–318.

Pashler, H. (1988b). Familiarity and visual change detection. *Perception and Psychophysics, 44,* 369–378.

Phillips, W. (1974). On the distinction between sensory storage and short-term visual memory. *Perception and Psychophysics, 41,* 191–201.

Rensink, R.A. (1998). Mindsight: Visual sensing without seeing. *Investigative Ophthalmology and Visual Science,* 39, 631.

Rensink, R.A. (this issue). The dynamic representation of scenes. *Visual Cognition, 7,* 17–42.

Rensink, R.A. (this issue). Visual search for change: A probe into the nature of attentional processing. *Visual Cognition, 7,* 345–376.

Rensink, R.A., O'Regan, J.K., & Clark, J.J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8*(5), 368–373.

Scholl, B.J., & Pylyshyn, Z.W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology, 38,* 259–290.

Scholl, B.J., Pylyshyn, Z.W., & Chen, V. (in preparation). *The nature of attentional allocation: Evidence from change detection.* Manuscript in preparation.

Simons, D.J. (1996). In sight, out of mind: When object representations fail. *Psychological Science, 7,* 301–305.

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

Simons, D.J., & Levin, D.T. (1997). Change blindness. *Trends in Cognitive Sciences, 1*(7), 261–267.

Simons, D.J., & Levin, D.T. (1999). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin and Review, 5,* 644–649.

Snodgrass, J.G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 174–215.

Theeuwes, J. (1991a). Cross-dimensional perceptual selectivity. *Perception and Psychophysics, 50,* 184–193.

Theeuwes, J. (1991b). Exogenous and endogenous control of attention: The effect of visual onsets and offsets. *Perception and Psychophysics, 49,* 83–90.

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics, 51,* 599–606.

Theeuwes, J. (1994). Endogenous and exogenous control of visual selection. *Perception, 23,* 429–440.

Theeuwes, J., & Burger, R. (1998). Attentional control during visual search: The effect of irrelevant singletons. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 1342–1353.

Todd, S., & Kramer, A. (1994). Attentional misguidance in visual search. *Perception and Psychophysics, 56, 198–210.*

Williams, P. & Tarr, M.J. (1998). *RSVP: Experimental control software for MacOS, Version 4.0.3.* [Computer program]. Available: http://psych.umb.edu/rsvp/ [1998, June 17].

Yantis, S. (1993a). Stimulus-driven attentional capture. *Current Directions in Psychological Science, 2,* 156–161.

Yantis, S. (1993b). Stimulus-driven attentional capture and attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 676–681.

Yantis, S., & Jonides, J. (1984). Abrupt visual onsets and selective attention: Evidence from visual search. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 601–621.

Yantis, S., & Jonides, J. (1990). Abrupt visual onsets and selective attention: Voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 121–134.

Zelinsky, G. J. (1998). *Detecting changes between scenes: A similarity-based theory using iconic representations* . (Tech. Rep. No. CNS-98-01). Champaign, IL: University of Illinois, Cognitive Neuroscience Group.