The role of representational volatility in recognizing pre- and postchange objects

MELISSA R. BECK and DANIEL T. LEVIN

Kent State University, Kent, Ohio

Theories relating attention to change blindness (CB) imply that representations of objects in the focus of attention are stable and coherent. However, CB occurs for objects in the focus of attention. Here, we explore this apparent contradiction and the possibility that changes can be detected without having a complete and stable representation of the prechange object. The first experiment required observers to recognize a prechange object and a postchange object after viewing arrays of various sizes in which the prechange object was replaced by the postchange object after a brief delay. Results indicated that the representation of the prechange object was strong enough to cue a change but not strong enough to support accurate recognition. The remaining experiments demonstrated that the representation of the prechange object is volatile in that a shift in the display or the presence of a postchange object can disrupt the representation. These findings add to current theories of attention and representations by showing that attention may result in volatile representations that can support change detection without supporting accurate recognition.

Recent research suggests that our visual system is not able to monitor every detail in our visual field. In particular, subjects fail to notice large changes to the location, properties, and identity of objects. This phenomenon, referred to as “change blindness” (CB), occurs in many different change detection tasks as long as some interruption (a flicker, a blink, an eye movement, or a movie cut) disrupts the transient that would generally cue a change (Blackmore, Brelstaff, Nelson, & Troscianko, 1995; Dennett, 1991; Grimes, 1996; Henderson, 1997; Irwin, Zacks, & Brown, 1990; Levin & Simons, 1997; McConkie & Currie, 1996; O’Regan, Rensink, & Clark, 1999; Pashler, 1988; Phillips, 1974; Rayner, McConkie, & Ehrlich, 1978; Rensink, 2000; Simons, 1996; Simons & Levin, 1997; for a review see Simons, 2000). These findings conflict with the impression that the coherent, richly detailed world that we see is associated with a similarly rich representation. If our visual representations are richly detailed, we should rarely miss changes to our visual world. Instead, visual representations are sparse and may be limited to a few objects or features that receive some attention (Becker, Pashler, & Anstis, 2000; Irwin, 1991; Levin & Simons, 1997; O’Regan, Deubel, Clark, & Rensink, 2000; Rensink, O’Regan, & Clark, 1997). Here, we explore these representations and their properties by testing the degree to which they are durable enough to support successful change de-

tection and recognition of changing objects. In particular, we test the possibility that a representation can allow for accurate change detection while failing to support accurate recognition.

Representational Failures Implied by Change Blindness

CB has been observed in a wide variety of tasks (see Rensink, 2002; Simons & Levin, 1997, for reviews). For example, Levin and Simons (1997) showed subjects videos in which a variety of changes occurred across cuts (for example, plates changed color, an actor’s scarf disappeared, and one actor was substituted for another) and found that very few subjects reported seeing any of the changes in the video, even when the change occurred in objects they were paying attention to (Levin & Simons, 1997). In addition, Simons and Levin (1998) examined changes in real-world interactions and found that fewer than half the subjects noticed when the identity of the person they were talking to changed midconversation. Although subjects in these experiments were not on the lookout for changes, similarly poor performance has been observed when subjects are actively searching for the changing objects (Hollingsworth & Henderson, 2000; Irwin, 1991; Irwin et al., 1990). For example, in the flicker paradigm (Rensink et al., 1997), brief blank fields were placed between alternating displays of an original and a modified scene. Typically, in these experiments, many alternations are required for subjects to see changes that are quite striking in retrospect.

Rensink (2000) introduced coherence theory to account for CB and to specify the requirements for successfully detecting a change. He argued that representations of objects in our visual field are in various states of

This research was presented in a poster at the 2000 annual Association for Research in Vision and Ophthalmology conference. We thank Bonnie Angelone for comments on an earlier version of this manuscript. We also thank Daniel Simons for providing data from his studies. Correspondence should be addressed to M. R. Beck, Department of Psychology, Kent State University, Kent, OH 44242-0001 (e-mail: mbeck@kent.edu).

volatility depending on whether they are the focus of attention or not. In particular, prior to the arrival of focal attention, object representations (proto-objects) are continually formed in parallel, have limited coherence in space and time, and are replaced when new visual stimuli appear. Once attention is focused on an object (or an object's features), its representation becomes coherent across space and time. When an object is no longer in focused attention, it becomes a proto-object once again with little or no aftereffect. Rensink (2000) argued that a change can be detected only if the object is the focus of attention when the change occurs. This might also imply that if a change is detected, representations of the changing object have been stabilized in both its pre- and postchange forms.

The notion of a coherence field makes sense as an account for the transition from the volatility of preattentive vision to a more complete and stable representation of attended objects. However, coherence theory does not specify the completeness of the representations that underlie change detection, only that they are stabilized while the object changes. This allows coherence theory to accommodate failures to detect changes to objects within the focus of attention, even in intentional change detection paradigms (O'Regan et al., 2000) under the assumption that in some cases only a subset of a given object's features is stabilized. Therefore, attention to an object may allow a representation that is in some sense stable, but it does not guarantee a complete representation and comparison across views. Instead, attention to the specific features that change may be necessary for successful change detection (Levin & Simons, 1997; O'Regan et al., 2000).

Accordingly, to understand change detection, it may be necessary to supplement existing theories with a more detailed account of the representations that support change detection. In particular, the coherence that allows change detection might be associated with a relatively sparse and/or volatile representation of an object, or it may cause a deeper encoding and consolidation process to occur. In the former case, change detection would result from the kind of identification and on-line comparison that is hypothesized to occur in conceptual short-term memory (Chun & Potter, 1995; Potter, 1999) but that would not allow later recognition. In the latter case, a more durable representation would be associated with change detection. This might include either a detailed and permanent perceptual representation or a durable verbal code (Simons, 1996; Zelinsky & Murphy, 2000).

There are a number of circumstances under which a change can be detected using incomplete or volatile representations (Becker et al., 2000) that might be subject to overwriting (see Simons, 2000, for review). Consider a task in which subjects view an array of objects that briefly disappears and then reappears with one of the objects changed. Several studies have shown that change detection performance is less than perfect for arrays of more than four objects (Pashler, 1988; Phillips, 1974; Simons, 1996). The question is, what underlies change detection

on the subset of trials (for arrays of more than four objects) where subjects are successful? According to coherence theory, the change was detected because attention arrived at the prechange object, making it coherent while it changed into the postchange object. The next question is, does coherence entail stability of the features that allowed change detection? If so, representations of both pre- and postchange objects could be both strong and equivalent. Such representations might allow subjects to recognize both the pre- and postchange forms of the changing object with equal success. This is particularly true in an intentional task, where subjects know that they will have to identify both the new object and the object that was replaced. The moment subjects see the change, they need only represent the pre- and postchange objects, and if even minimally durable representations of both are available, subjects should have access to them, or at least recode them in more durable form (e.g., by labeling them). If, on the other hand, change detection can be based on more volatile representations, recognition may be worse for the prechange object if the representation is not durable enough to be remembered or recoded once the subject knows it is a prechange object.

Thus far, one study we know of has systematically tested subjects' ability to report information about pre- and postchange objects in an intentional task. Brawn, Snowden, and Wolfe (1999) showed subjects displays containing 32 red or green filled circles. One of the objects changed luminance across a 500–1,000 msec interstimulus interval (ISI), and on half of the luminance-change trials, the object also changed color. The subjects' task was to report whether a luminance change occurred, and if it did, to report what color the circle had been prior to the change. In this task, subjects were above chance at detecting the luminance change, but could not report the color of the object before the change. Brawn et al. argued that in this case, the representation of the postchange object overwrote the representation of the prechange object. Accordingly, it is possible that the representation of the prechange object was weak and transitory despite the fact that it was sufficient to allow change detection. However, there are two limits to this study. First, the changing features and the reported features were complementary and probably processed through different physiological channels. Subjects had to attend to luminance information to detect the change and then switch to color information to describe the prechange object. Second, the change detection task was based on a fairly diffuse spreading of visual attention. Subjects were 79% accurate at detecting the change among 32 objects, which far outstrips the capacity of visual short-term memory (VSTM), implying that the changing object may have been processed only as part of a larger object formed via grouping. Accordingly, representations of the color of the prechange object were at a considerable disadvantage because the object may have been embedded in a larger object, and what attention it did receive was limited to luminance, which was not the feature that subjects had to report.

Here, we tested explicit representations of pre- and postchange objects under conditions more conducive to strong representations of both objects. First, we used a task in which whole objects changed. Subjects viewed arrays of natural objects that disappeared for 350 msec and then were replaced by new arrays in which one of the objects was different. In addition, the new arrays were randomly displaced to the left or right. The arrays contained from 3 to 16 objects. Therefore, the number of objects in the smaller arrays did not exceed the capacity of VSTM, thus allowing subjects to focus attention on all of the objects. Performance for these arrays should reach ceiling and then decline for the larger arrays. Therefore, we can examine the effects of outstripping VSTM capacity on change detection performance. After viewing the changing arrays, subjects completed recognition tests for the pre- and postchange objects. The recognition test was a four-alternative forced-choice test (4AFC) in which four objects were presented (the target, a nonchanging object from the array, and two distractors). The subjects chose which of the four was the pre- or postchange object.

The 4AFC recognition test was our sole dependent measure of change detection. Therefore, unlike Brawn et al.'s (1999) subjects, subjects in the present study did not have to divide their attention between two tasks (reporting *whether* a change occurred and reporting *what* change occurred). In the present change detection tasks a change occurred in every trial, and the subjects had to recognize only the pre- and postchange objects. Although it can be argued that detecting a change and describing what has changed are not identical tasks, researchers have found similar levels of change detection and change identification and have therefore examined change identification only (Mondy & Coltheart, 2000). In addition, our 4AFC was designed to control for responses based simply on familiarity for an object. The recognition test included three distractors, one of which was a randomly chosen nonchanging object from the array (the lure). Therefore, subjects could not simply pick the object that was the most familiar (an effective strategy if all distractors were not from the array) or the object that was least familiar (an effective strategy if all distractors were nonchanging objects from the array). Accuracy on the test was determined by subtracting lure false alarms from hits on target. Therefore, accurate responding on the 4AFC is based on detection of the change, not on familiarity or a general memory for any or all objects in the array.

The present experiments will provide further clarification of the representations that underlie change detection. In particular, it may be possible to detect changes using volatile representations that are then unable to support accurate recognition of the changing objects. If stable representations are required for change detection, then we should expect equal recognition performance for pre- and postchange objects. If, on the other hand, volatile representations can support change detection, then subjects should show poorer recognition for prechange objects than for postchange objects.

EXPERIMENT 1

Method

Subjects. Thirty-three introductory psychology students at Kent State University participated in exchange for class credit. For 18 subjects, the prechange 4AFC was presented before the postchange 4AFC (pre-first condition), and for 15 subjects the postchange 4AFC was presented before the prechange 4AFC (post-first condition).

Materials. Arrays were created using the 260 images from the Snodgrass and Vanderwart (1980) stimulus set. Each object was digitized in 256-level gray scale and sized to a maximum of 120 pixels on either dimension, depending on the shape of the particular objects. Individual images were presented at 72 dpi and were therefore a maximum of approximately 4.5 cm wide or high. Subjects viewed the screen from an approximate distance of 50 cm.

Pairs of pre- and postchange arrays varied in size from 3 to 16 objects. Within each pair, all objects were the same except that 1 of the objects in the postchange array (postchange object) was different from the object in corresponding location of the prechange array (prechange object). Objects were randomly chosen without replacement within each array pair and with replacement across array pairs. The postchange object was chosen randomly from all images except those already in the prechange array. The objects were randomly placed in 1 of 16 locations in a 4×4 grid (containing 4 columns and 4 rows). The entire grid was 18×22 cm, with each cell measuring 4.5×5.5 cm. The objects were separated by approximately 1.5 cm horizontally and .5 cm vertically. The arrays were presented on a 21×28 cm screen and subtended approximately 25° of visual angle. The objects were black and placed on a white background with a luminance of 61.58 fL (211 cd/m^2). The ISI presented between the pre- and postchange scenes was the all-white background with no objects present.

Subjects saw two 4AFC recognition tests (a prechange test and a postchange test) consisting of the postchange object or the prechange object, a lure (a nonchanging object from the corresponding pair of arrays), and two objects that were not present in the corresponding pair of arrays (Figure 1).

Four sets of pre- and postchange arrays and 4AFC questions for Array Sizes 3–16 were generated on each of four subject-running computers. Thus each computer produced 56 unique pre- and postchange arrays and 56 pre- and postchange 4AFC recognition tests. Each subject completed a total of 56 trials, 4 at each array size (each consisting of a prechange array, a postchange array, a pre–4AFC recognition test, and a post–4AFC recognition test).

Apparatus. Images were presented on 15-in. monitors set at a resolution of 600×800 (72 dpi) and at 16-bit color depth (thousands of colors). Power Computing and Motorola MacOS computers controlled the experiment, recorded subjects' responses, and created the object arrays.

Procedure. Before the change detection trials, the experimenter read instructions to the subjects. They were told that on each trial they would see an array of objects for 2 sec, and then it would be replaced by a second array that would be the same except one object would be replaced with another object. They were instructed that their job was to find the objects that were different between the arrays and then to identify them in recognition tests. Subjects in the pre-first condition were told they would first identify the prechange object and then the postchange object and subjects in the post-first condition were told the opposite.

On each trial, subjects saw the prechange object array for 2 sec and then, after a 350-msec ISI, the postchange object array for 2 sec. The postchange array shifted a half column (2.25 cm; recall that objects were arranged in four columns on the screen) either to the left or to the right. After the postchange array, subjects in pre-first condition completed a 4AFC recognition test first for the prechange object and then for the postchange object. Both recogni-

Figure 1. Example materials and trial sequence from Experiment 1.

tion tests contained four objects aligned horizontally and numbered 1 through 4. Beneath the objects was a question, “Which object got replaced? (1–4)” for the prechange 4AFC and “Which object was the new object? (1–4)” for the postchange 4AFC. Subjects responded by pressing 1–4, indicating which object was the target. The trials were exactly the same for subjects in the post-first condition except the 4AFC questions were presented in the opposite order (post first, then pre). The questions were response terminated with a 500-msec ISI between questions. One second after the second 4AFC question was answered, the next trial began.

Subjects completed the experiment in small groups ranging in size from 1 to 4 and were tested on individual computers in the same room. Alternate groups of subjects were assigned to the two experimental conditions. Subjects completed the experiment in approximately 15 min.

Results

Primary analyses were conducted using the percentage of hits on target (the target was correctly chosen) minus

the percentage of lure false alarms (the lure was incorrectly chosen). Figure 2 illustrates the percentage of hits on the target minus the percentage of lure false alarms at each array size combined across conditions. The graph shows that performance decreased as a function of array size and that accuracy for the postchange object was consistently higher than accuracy for the prechange object. The graph also shows that performance was high for Array Sizes 3 and 4 and then began to decline.

Hits on target minus lure false alarms. On the basis of pilot experiments indicating that performance levels out after Array Size 9, accuracy (the percentage of hits on the target minus the percentage of lure false alarms) for pre- and postchange objects across small array sizes (3–9) was entered into a two-factor analysis of variance (ANOVA) with type of recognition question (pre, post) as a within-subjects factor and condition (pre-first, post-

Figure 2. Percentage of hits on target minus percentage of lure false alarms (accuracy) for all conditions in Experiment 1.

first) as a between-subjects factor. The ANOVA revealed a significant main effect for recognition question (pre vs. post) [$F(1,31) = 39.134, p < .001$]. In particular, prechange accuracy ($M = 37\%$) was significantly lower than postchange accuracy ($M = 59\%$). However, there was no main effect for condition; pre-first accuracy ($M = 50\%$) was not significantly different from post-first accuracy ($M = 46\%$) [$F(1,31) = .721, p = .402$]. In addition, there was a significant interaction between condition and type of recognition question [$F(1,31) = 8.162, p = .008$]. These results indicate that the order of the recognition questions affected accuracy. However, there was still a strong effect for prechange accuracy versus postchange accuracy. Planned comparisons of accuracy at small array sizes for each condition revealed that in the pre-first condition the difference between prechange accuracy ($M = 44\%$) and postchange accuracy ($M = 57\%$) was significant [$t(17) = -2.609, p = .018$], and in the post-first condition, the prechange ($M = 29\%$) and postchange ($M = 62\%$) difference was also significant [$t(14) = -5.964, p < .001$].

An analysis of accuracy combined across all array sizes (3–16) revealed similar results. A two-factor ANOVA revealed a significant main effect for recognition question (pre vs. post) [$F(1,31) = 49.526, p < .001$]. In particular, prechange accuracy ($M = 23\%$) was significantly lower than postchange accuracy ($M = 38\%$). There was no main effect for condition; pre-first accuracy ($M = 31\%$) was not significantly different from post-first accuracy ($M = 30\%$) [$F(1,31) = .112, p = .740$]. In addition, there was a significant interaction between condition and type of recognition question [$F(1,31) = 5.644, p = .024$]. Planned comparisons of accuracy for each condition revealed that in the pre-first condition the

difference between prechange accuracy ($M = 26\%$) and postchange accuracy ($M = 37\%$) was significant [$t(17) = -4.017, p = .001$], and in the post-first condition, the prechange ($M = 20\%$) and postchange ($M = 41\%$) difference was also significant [$t(14) = -5.558, p < .001$].

Further analysis revealed that the differences between pre- and postchange recognition performance was quite strong for certain array sizes even when performance was combined across conditions. There was a significant difference between pre- and postchange performance for Array Sizes 5, 6, and 8. The percentage of hits on target minus the percentage of lure false alarms for each subject was entered into a paired-samples t test. Prechange performance was significantly worse at Array Size 5 (pre = 38% and post = 66%) [$t(32) = -3.107, p = .004$]; at Array Size 6 (pre = 24% and post = 58%) [$t(32) = -3.552, p = .001$]; and at Array Size 8 (pre = 9% and post = 49%) [$t(32) = -4.635, p < .001$].

Hits on target. Similar results were found when hits on target were used as the measure of accuracy. The ANOVA of hits on target across small array sizes (3–9) revealed a significant main effect for recognition question (pre vs. post) [$F(1,31) = 51.309, p < .001$]. In particular, prechange accuracy ($M = 56\%$) was significantly lower than postchange accuracy ($M = 72\%$). There was no main effect for condition; pre-first accuracy ($M = 65\%$) was not significantly different from post-first accuracy ($M = 63\%$) [$F(1,31) = .208, p = .652$]. In addition, there was a significant interaction between condition and type of recognition question [$F(1,31) = 6.208, p = .018$]. Planned comparisons of accuracy at small array sizes for each condition revealed that in the pre-first condition the difference between prechange accuracy ($M = 60\%$) and postchange accuracy ($M = 70\%$) was significant [$t(17) =$

$-3.125, p = .006$], and in the post-first condition, the prechange ($M = 52\%$) and postchange ($M = 74\%$) difference was also significant [$t(14) = -7.69, p < .001$].

An analysis of accuracy combined across all array sizes (3–16) revealed similar results. A two-factor ANOVA revealed a significant main effect for recognition question (pre vs. post) [$F(1,31) = 63.318, p < .001$]. In particular, prechange accuracy ($M = 45\%$) was significantly lower than postchange accuracy ($M = 57\%$). There was no main effect for condition; pre-first accuracy ($M = 51\%$) was not significantly different from post-first accuracy ($M = 50\%$) [$F(1,31) = .081, p = .778$]. In addition, contrary to analyses of hits on target minus lure false alarms, the interaction between condition and type of recognition question was not significant [$F(1,31) = 3.336, p = .077$]. Planned comparisons of accuracy for each condition revealed that in the pre-first condition the difference between prechange accuracy ($M = 46\%$) and postchange accuracy ($M = 56\%$) was significant [$t(17) = -4.048, p = .001$], and in the post-first condition, the prechange ($M = 43\%$) and postchange ($M = 58\%$) difference was also significant [$t(14) = -8.018, p < .001$].

Analysis at each array size revealed that the differences between pre- and postchange recognition performance was quite strong for certain array sizes even when performance was combined across conditions. There was a significant difference between pre- and postchange performance for Array Sizes 5, 6, and 8. The percentage of hits on target minus the percentage of hits on lure false alarms for each subject was entered into a paired-samples t test. Prechange performance was significantly worse at Array Size 5 (pre = 60% and post = 77%) [$t(32) = -3.305, p = .002$]; at Array Size 6 (pre = 49% and post = 71%) [$t(32) = -4.432, p = .001$]; and at Array Size 8 (pre = 37% and post = 65%) [$t(32) = -4.723, p < .001$].

Discussion

Experiment 1 shows that recognition accuracy (hits on target minus lure false alarms) for the postchange object was greater than accuracy for the prechange object. For some array sizes, this effect was quite large, ranging from 28% to 40% in the five-, six-, and eight-object arrays. Therefore, whatever the information signaling change (and therefore leading to accurate recognition for the postchange object), it is sometimes insufficient to identify the prechange object. This implies that a change can be detected without having a durable representation of the prechange object. It is important to emphasize that in an intentional task such as this one, subjects know they will have to report the identity of the pre- and postchange objects and therefore as soon as they detect the change, they should poll their VSTM for the prechange object. The fact that prechange recognition is worse than postchange recognition suggests that something about the prechange representation is unavailable the moment after it is used to detect the change.

The results from Experiment 1 indicate that our dependent measure (pre- and postchange object recognition) yields change detection capacities that are similar

to those of other dependent measures. Most change detection research employs a same–different paradigm. In this paradigm the subject is shown the prechange array and then the postchange array and asked to indicate if the two arrays were the same or if they were different. Simons (1996) used this paradigm with arrays of five objects very similar to those used in the present experiment. Using Pashler's (1988) formula to calculate change detection capacity in Simon's (1996) experiment, we found a capacity of 3.6 objects. This is very similar to the capacities we found for postchange object recognition performance across all conditions for the five object arrays (3.3 objects) using Sperling's (1960) formula for capacity. Therefore, it can be argued that recognition tests and same–different tests result in similar change detection capacities.

Given that prechange identification is poor, it is important to understand why. One possibility is that the postchange array had an advantage over the prechange array because the objects in the postchange were viewed twice (first in the prechange array, then again in the postchange array), making the memory for the objects stronger. Subjects could identify the new (postchange) object in the postchange array by eliminating objects that were in the prechange array. This elimination strategy should lead to very low lure false alarms (at least lower than those for the other distractors). However, this is not what we found; the lure was chosen 16% of the time and each distractor was chosen 10% of the time. Another possibility is that the representation of the prechange object was stable but not complete. The representation could be incomplete because it is a coarse representation that is barely sufficient to distinguish it from the postchange object. This would allow change detection and would focus attention on the postchange object but in some cases would not be sufficient to distinguish the prechange object from the three incorrect alternatives on the 4AFC test. For example, a basic contour statistic such as curvilinearity might be sufficient to contrast most pre- and postchange objects, but insufficient to distinguish prechange objects from the three distractors in the recognition test.

The coarse coding explanation assumes that the result of attention was not necessarily to conjoin features or to make the object coherent, but rather to make permanent a very primitive coding of the prechange object. We consider it unlikely that representations of natural objects would remain in such a primitive state in a situation where attention is distributed over a comparatively small number of objects. Even in preattentive vision, it appears that object representations are organized to some degree. Depth relationships are available (He & Nakayama, 1992), and it appears that familiar contour shapes are identified simultaneously with figure–ground relationships (Peterson, 1994). In addition, it appears that some simple form elements may be available (Enns & Rensink, 1990). Thus, it is likely that whatever the coding used to detect the change, it was probably specific enough to distinguish the target from the distractors in the 4AFC task. However,

the possibility does remain that some simple contour statistic is available, and that subjects can selectively attend to it in an effort to more broadly search the display.

Another possible explanation for poor prechange recognition is that the prechange object is part of a set of events that get segmented at the change. If this event is segmented similarly to sentences or motion pictures, then memory might be poor because the prechange object gets associated with an earlier part of the event, just as words that are part of early sentence segments are more poorly remembered (Carroll & Bever, 1976; Fodor & Bever, 1965). A similar possibility is that the postchange object has more time available to be consolidated into long-term memory. Finally, the postchange object could overwrite the prechange object. On this hypothesis, once the change is detected, the new object replaces at least some of the information represented about the old object (Brawn et al., 1999; Kahneman, 1968; Simons, 2000).

In Experiment 2 we tested the hypothesis that the postchange object overwrites the prechange object. We did this by having the prechange object disappear, rather than being replaced. Thus there was no postchange object, just a blank where the prechange object used to be. If the postchange object caused poor prechange performance in Experiment 1, removing it should result in increased accuracy for recognizing the prechange object.

EXPERIMENT 2

Method

Subjects. Eighteen introductory psychology students at Kent State University participated in exchange for class credit.

Procedure. The procedure was identical to that in Experiment 1 except the prechange object was not replaced by a postchange object. It simply disappeared. Therefore, the postchange array was

smaller than the prechange array by one object. Consequently, there was only a prechange 4AFC recognition test and only one condition (pre-only).

Results

As in Experiment 1, accuracy was defined as hits on target minus lure false alarms for analysis and figures. Figure 3 shows the percentage of hits on the target minus the percentage of lure false alarms at each array size for the prechange object in this experiment and for the prechange object from the pre-first condition in Experiment 1. The graph shows that accuracy for the prechange object in Experiment 2 was very similar to that in Experiment 1.

Hits on target minus lure false alarms. Accuracy for the prechange objects across small array sizes (3–9) was 47%. This was very similar to the accuracy in the pre-first condition from Experiment 1 ($M = 44\%$) [$t(33) = -.452$, $p = .655$]. Similar results were found for prechange accuracy across all array sizes ($M = 31\%$) for the pre-only condition in this experiment and the pre-first condition in Experiment 1 ($M = 26\%$) [$t(33) = -.977$, $p = .336$].

Hits on target. Accuracy for the prechange objects across small array sizes (3–9) was 61%. This was very similar to the accuracy in the pre-first condition from Experiment 1 ($M = 60\%$) [$t(33) = -.273$, $p = .787$]. Similar results were found for prechange accuracy across all array sizes ($M = 49\%$) for the pre-only condition in this experiment and the pre-first condition in Experiment 1 ($M = 46\%$) [$t(33) = -.817$, $p = .420$].

Discussion

Prechange recognition performance did not improve in Experiment 2. Therefore, it remains possible that

Figure 3. Percentage of hits on target minus percentage of lure false alarms (accuracy) for the pre-first 4AFC in Experiment 1 and the pre-only 4AFC in Experiment 2.

prechange recognition was poor because there was only a coarse representation of the prechange object or because the prechange object was an early segment of a set of events. However, the presence of the postchange object was not responsible for poor recognition of the prechange object. These findings run counter to at least one version of the overwriting theory. That is, a new object must replace the old object, and further, the new object must occupy the same spot in the array configuration as the old object. However, it is also possible that any new stimulus in the same location as the prechange object is sufficient for overwriting. On this hypothesis, new visual information of any kind is sufficient to mask some part of the prechange object, leading to disruption of the prechange representation. Accordingly, the stimulus shift that occurs between the presentation of the pre- and postchange arrays is enough to overwrite the representation of the prechange object. In Experiment 3, we removed the stimulus shift, which should result in increased prechange accuracy if the representation is being disrupted by masking.

EXPERIMENT 3

Method

Subjects. Twenty-eight introductory psychology students at Kent State University participated in exchange for class credit.

Procedure. This procedure was the same as in Experiment 2 except for the following changes. Subjects completed 56 trials similar to those from the pre-first condition in Experiment 1 (in which subjects saw the prechange object replaced the postchange object, then completed a prechange 4AFC recognition test followed by a postchange 4AFC recognition test) and 56 trials similar to those in Experiment 2 (in which there was no postchange object and subjects completed only a prechange 4AFC). Therefore, the trials with

a postchange object consisted of two 4AFC recognition tests (prechange and postchange), and the trials without a postchange object consisted of only the prechange 4AFC recognition test. The 112 trials were randomly presented. Subjects knew they would get both types of trials, but they did not know before each trial whether the prechange object would be replaced or simply disappear. In all trials the stimulus shift was removed. Therefore, the postchange array was presented in the same location as the prechange array after the 350-msec ISI.

Results

As in Experiments 1 and 2, accuracy was calculated as the percentage of hits on the target minus the percentage of lure false alarms. Figure 4 shows the percentage of hits on the target minus the percentage of lure false alarms at each array size for the prechange object and the postchange object in the pre-first condition and for the prechange object in the pre-only condition. In the pre-first condition, accuracy for the prechange object was consistently lower than accuracy for the postchange object. However, in the pre-only condition, accuracy for the prechange object was greater than accuracy for both the pre- and postchange objects in the pre-first condition.

Hits on target minus lure false alarms. A repeated measures MANOVA with accuracy (at the small array sizes) for the three types of recognition question (pre-first, post, and pre-only) as a within-subjects factor was significant [$F(2,54) = 15.175, p < .001$]. An analysis of all pairwise comparisons at the small array sizes revealed that pre-only accuracy ($M = 67\%$) was significantly higher than post accuracy ($M = 56\%$) [$F(1,27) = 8.01, p = .009$] and that post accuracy was significantly higher than pre-first accuracy ($M = 45\%$) [$F(1,27) = 7.05, p = .013$].

Figure 4. Percentage of hits on target minus percentage of lure false alarms (accuracy) for all conditions in Experiment 3.

In analyses conducted using accuracy across all array sizes (3–16), a repeated measures MANOVA for the three types of recognition question (pre-first, post, and pre-only) as a within-subjects factor was significant [$F(2,54) = 14.010, p < .001$]. Pairwise comparisons revealed that pre-only accuracy ($M = 41\%$) was not significantly higher than post accuracy ($M = 39\%$) [$F(1,27) = .427, p = .519$] and that post accuracy was significantly higher than pre-first accuracy ($M = 26\%$) [$F(1,27) = 16.566, p < .001$].

Hits on target. A repeated measures MANOVA with accuracy (at the small array sizes) for the three types of recognition question (pre-first, post, and pre-only) as a within-subjects factor was significant [$F(2,54) = 14.194, p < .001$]. An analysis of all pairwise comparisons at the small array sizes revealed that there was a nonsignificant trend for higher pre-only accuracy ($M = 76\%$) than for post accuracy ($M = 71\%$) [$F(1,27) = -2.922, p = .099$], and that post accuracy was significantly higher than pre-first accuracy ($M = 61\%$) [$F(1,27) = 11.013, p = .003$].

In analyses conducted using accuracy across all array sizes (3–16), a repeated measures MANOVA for the three types of recognition question (pre-first, post, and pre-only) as a within-subjects factor was significant [$F(2,54) = 17.253, p < .001$]. Pairwise comparisons revealed that pre-only accuracy ($M = 57\%$) was not different from post accuracy ($M = 57\%$) [$F(1,27) = .020, p = .890$], and that post accuracy was significantly higher than pre-first accuracy ($M = 47\%$) [$F(1,27) = 23.620, p < .001$].

Discussion

When the prechange object was replaced and there was no stimulus shift, accuracy for both the pre- and postchange objects was almost identical to accuracy for the comparable conditions in Experiment 1 (prechange accuracy, $M = 44\%$, and postchange accuracy, $M = 57\%$). However, when the prechange object was not replaced and there was no stimulus shift, accuracy for the prechange object was not lower than accuracy for the postchange object in the pre-first condition. These results indicate that the representation of the prechange object is overwritten if there is either a stimulus shift or a postchange object. However, if neither of these events occurs, the representation of the prechange object supports accurate recognition. Therefore, changes can be detected using volatile representations that support accurate recognition only if they are not disrupted by a stimulus shift or a postchange object.

In addition to revealing a general overwriting effect, Experiment 3 also provides evidence against the coarse coding hypothesis and the segmentation hypothesis mentioned in the discussion of Experiment 1. The coarse coding hypothesis assumes that only a coarse but durable representation of the prechange object is stored. However, if reduced pre-change accuracy were caused by a coarse-but-durable representation, then one would not expect that removal of the postchange array shift would

increase prechange performance. Inherent in this alternative is the notion that the coarse representation will not be overwritten. If invulnerability to overwriting were not part of this hypothesis, one would end up with a coarse-but-volatile hypothesis similar to a more general overwriting hypothesis. Such a hypothesis would suggest that changes are missed not only because the information is volatile, but also because it is coarse. Thus, we would draw the same conclusion that we would otherwise draw: that changes can be detected using volatile representations that do not support accurate recognition. The segmentation hypothesis assumes that poor prechange performance occurs because the prechange object is part of an early segment in the ongoing event represented by the complete trial. In Experiment 3, however, the prechange object was still part of this early segment, but was nonetheless recognized accurately. Similarly, these results cannot be due to less long-term memory consolidation for the prechange object, because consolidation time for the prechange object was not different in the pre-only and pre-first conditions.

Another possible explanation for these results is that a low-level iconic representation is preserved in the pre-only condition. That is, the empty spot in the postchange array may serve as a partial report cue for objects in the prechange array. Therefore, the subject is able to recall objects in that row and accurately identify the prechange object. This is somewhat plausible given the time constraints in Experiment 3. The ISI is 350 msec and an iconic representation could persist for this amount of time (Loftus, Duncan, & Gehrig, 1992; Sperling, 1960). Therefore, we ran 13 subjects in an experiment identical to Experiment 3 except the ISI was increased to 2,000 msec. This experiment yielded the same results as Experiment 3. Performance (hits on target minus lure false alarms for small display sizes) for the pre-first trials ($M = 43\%$) was significantly worse than performance for the post trials ($M = 54\%$) [$t(12) = -2.695, p = .020$]. However, pre-only trials ($M = 59\%$) were not significantly different from post trials [$t(12) = -1.086, p = .299$]. Therefore, it is unlikely that an iconic representation caused the results in Experiment 3.

GENERAL DISCUSSION

The experiments presented here show that postchange objects can be recognized without recognition of the prechange object. Therefore, the representation of the prechange object is sufficient to cue the change, but not always sufficient to be coded into a form useful for the recognition test. The results of Experiments 2 and 3 demonstrate that poor recognition of the prechange object occurs because it is overwritten by the postchange array if the array contains a postchange object or if the array shifts to the left or the right. These experiments also demonstrate that sometimes there is a durable representation of the prechange object representation that is not overwritten by the postchange object or the stimulus

shift. In fact, the prechange object was chosen more often than the lure and the distractors in all three experiments. Therefore, either durable or volatile representations can be used for change detection.

The explanation we have been proposing to account for these results is that the representation for the prechange object can be volatile and easily lost. We are making this inference about representations on the basis of a recognition test that occurs several seconds after the change. We therefore acknowledge the possibility that some other test of object representations would reveal equal performance on the pre- and postchange objects. It is, for example, possible that equal representations of the two objects were created initially, and that the prechange representation became inaccessible in the 3-sec interval between the change and the recognition test, or that the prechange representation was never consciously accessible. Therefore, a more immediate test of the prechange object or an implicit test might have resulted in more accurate prechange performance. However, just about any conceivable test of a given representation falls victim to the argument that it is not exhaustive. Therefore, we take a functional view of representations and emphasize that we are testing for a representation that is useful in consciously recognizing elements of a previously viewed scene. On this view, the putative representational failure in prechange objects is appropriately matched to our methodology and further delimits an ecologically important way that representations can fail.

Another possible explanation for the results presented here is that the representation of the prechange object was not volatile but was simply weaker because it received less focal attention than the postchange object. This possibility assumes that subjects scanned the objects in the prechange array one at a time and that in some cases they detected a change when a fixation on one or more nonchanging objects intervened between fixations on the pre- and postchange objects. Studies have shown that only a limited number of objects (three to four) can be held in VSTM while the eyes scan a scene (Irwin, 1996). Therefore the prechange object suffers the disadvantage of having been coded several spots earlier in the sequence of represented objects and therefore may have decayed prior to the change.

However, this "weak representation" hypothesis doesn't seem problematic for the representational volatility explanation. Almost any way of understanding poor prechange performance would assume that representations of the prechange object suffer some disadvantage before the change is detected. The point is that the representation was strong enough to cue the change, but not strong enough to drive accurate recognition performance or to recode the moment the change was visible. It is important to emphasize here that prechange performance was poor even in situations where subjects knew that the only information they would have to report was the prechange object (as in Experiment 2). Therefore, it appears that information suitable for abstraction was not available about

the prechange object, even in the moment after the change was detected.

In the introduction, Rensink's (2000) coherence theory was introduced as a possible explanation for CB. Coherence theory states that objects in the focus of attention have stable spatiotemporal representations of the structure of the objects but not necessarily the parts of the objects. Therefore, the present results suggest that although the representation of the prechange object may be stable enough to support accurate change detection, this stability of the parts can be disrupted by a shift in the postchange array or the presence of a postchange object. We can conclude that the coherent representations described in coherence theory are not stable in that they can be overwritten by the postchange object or other visual stimuli. According to coherence theory, if a change is detected, attention was directed to the changing object before and after the change. The present findings complement or extend coherence theory by showing that attention can result in representations that support accurate change detection but are too unstable to support accurate recognition. Therefore, attention may not lead to completely stable representations, and completely stable representations may not be necessary to detect changes.

REFERENCES

- BECKER, M. W., PASHLER, H., & ANSTIS, S. M. (2000). How quickly we forget: Change blindness in a highest digit task shows the volatility of visual representations. Poster presented at the annual meeting of the Association for Research in Vision and Ophthalmology, Fort Lauderdale.
- 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.
- BRAWN, P. T., SNOWDEN, R. J., & WOLFE, J. M. (1999). *The minimal conditions for "change blindness": What is replaces what was*. Poster presented at the annual meeting of the Association for Research in Vision and Ophthalmology, Fort Lauderdale.
- CARROLL, J. M., & BEVER, T. G. (1976). Segmentation in cinema perception. *Science*, **191**, 1053-1055.
- CHUN, M. M., & POTTER, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 109-127.
- DENNETT, D. C. (1991). *Consciousness explained*. Boston: Little, Brown.
- ENNS, J. T., & RENSINK, R. A. (1990). Sensitivity to three-dimensional orientation in visual search. *Psychological Science*, **1**, 323-326.
- FODOR, J. A., & BEVER, T. G. (1965). The psychological reality of linguistic structures. *Journal of Verbal Learning & Verbal Behavior*, **4**, 414-420.
- GRIMES, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.), *Perception* (Vancouver Studies in Cognitive Science, Vol. 2, pp. 89-110). New York: Oxford University Press.
- HE, Z. J., & NAKAYAMA, K. (1992). Surface versus features in visual search. *Nature*, **359**, 231-233.
- HENDERSON, J. M. (1997). Transsaccadic memory and integration during real-world object perception. *Psychological Science*, **8**, 51-55.
- HOLLINGSWORTH, A., & HENDERSON, J. M. (2000). Semantic informativeness mediates the detection of changes in natural scenes. *Visual Cognition*, **7**, 213-235.
- IRWIN, D. E. (1991). Information integration across saccadic eye movements. *Cognitive Psychology*, **18**, 420-456.
- IRWIN, D. E. (1996). Integrating information across saccadic eye movements. *Current Directions in Psychological Science*, **5**, 94-100.
- IRWIN, D. E., ZACKS, J. L., & BROWN, J. S. (1990). Visual memory and

- the perception of a stable visual environment. *Perception & Psychophysics*, **47**, 35-46.
- KAHNEMAN, D. (1968). Method, findings, and theory in studies of visual masking. *Psychological Bulletin*, **70**, 404-425.
- LEVIN, D. T., & SIMONS, D. J. (1997). Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin & Review*, **4**, 501-506.
- LOFTUS, G. R., DUNCAN, J., & GEHRIG, P. (1992). On the time course of perceptual information that results from a brief visual presentation. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 530-549.
- MCCONKIE, G. W., & CURRIE, C. B. (1996). Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 563-581.
- MONDY, S., & COLTHEART, V. (2000). Detection and identification of change in naturalistic scenes. *Visual Cognition*, **7**, 281-296.
- O'REGAN, J. K., DEUBEL, H., CLARK, J. J., & RENSINK, R. (2000). 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. (1988). Familiarity and visual change detection. *Perception & Psychophysics*, **44**, 369-378.
- PETERSON, M. A. (1994). Object recognition processes can and do operate before figure-ground organization. *Current Directions in Psychological Science*, **3**, 105-111.
- PHILLIPS, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics*, **16**, 283-290.
- POTTER, M. C. (1999). Understanding sentences and scenes: The role of conceptual short-term memory. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli* (pp. 13-46). Cambridge, MA: MIT Press.
- RAYNER, K., MCCONKIE, G. W., & EHRLICH, S. (1978). Eye movements and integrating information across fixations. *Journal of Experimental Psychology: Human Perception & Performance*, **4**, 529-544.
- RENSINK, R. A. (2000). The dynamic representation of scenes. *Visual Cognition*, **7**, 17-24.
- RENSINK, R. A. (2002). Change detection. *Annual Review of Psychology*, **53**, 245-277.
- 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**, 368-373.
- SIMONS, D. J. (1996). In sight, out of mind: When object representations fail. *Psychological Science*, **7**, 301-305.
- SIMONS, D. J. (2000). Current approaches to change blindness. *Visual Cognition*, **7**, 1-15.
- SIMONS, D. J., & LEVIN, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, **1**, 262-267.
- SIMONS, D. J., & LEVIN, D. T. (1998). Failure to detect changes to people in a real-world interaction. *Psychonomic Bulletin & 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 & Memory*, **6**, 174-215.
- SPELTING, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, **74**(11, Whole No. 498), 1-29.
- ZELINSKY, G. J., & MURPHY, G. L. (2000). Synchronizing visual and language processing: An effect of object name length on eye movements. *Psychological Science*, **11**, 125-131.

NOTES

1. Pashler's (1988) formula for capacity: $H = k/16 + (16 - k/16 * g)$, where H = hit rate, g = false alarm rate, 16 is the number of objects in the array, and k = capacity.

2. Sperling's (1960) formula for capacity: Percentage correct * array size = capacity. Here we used hits on target minus lure false alarms as the measure of accuracy.

(Manuscript received April 25, 2001;
revision accepted for publication September 25, 2002.)