Perception, 2003, volume 32, pages 947 – 962

DOI:10.1068/p5079

The relationship between change detection and recognition of centrally attended objects in motion pictures

Bonnie L Angelone, Daniel T Levin
Department of Psychology, Kent State University, PO Box 5190, Kent, OH 44242-0001, USA;
e-mail: bbruton@kent.edu

Daniel J Simons¶
Department of Psychology, Harvard University, Cambridge, MA 02138, USA
Received 10 September 2002, in revised form 28 April 2003

Abstract. Observers typically detect changes to central objects more readily than changes to marginal objects, but they sometimes miss changes to central, attended objects as well. However, even if observers do not report such changes, they may be able to recognize the changed object. In three experiments we explored change detection and recognition memory for several types of changes to central objects in motion pictures. Observers who failed to detect a change still performed at above chance levels on a recognition task in almost all conditions. In addition, observers who detected the change were no more accurate in their recognition than those who did not detect the change. Despite large differences in the detectability of changes across conditions, those observers who missed the change did not vary in their ability to recognize the changing object.

1 Introduction

Although we experience a detailed visual world, we often fail to detect changes in natural and artificial scenes, a phenomenon known as change blindness (Irwin 1996; Levin and Simons 1997; Pashler 1988; Phillips 1974; Rensink et al 1997; Simons 1996; Simons and Levin 1997; Simons and Levin 1998). Since early studies demonstrating change blindness, researchers have used people's inability to detect changes to gain insights into the nature of the mental representations underlying change detection. The most intuitive explanation for change blindness is that it reveals a failure to represent the visual world (O'Regan 1992; Simons 2000). The current experiments investigate this hypothesis by examining change detection for, and recognition of, centrally attended objects in motion pictures. We made a change to the central actor in a motion picture of a conversation. After viewing the video, observers were asked whether they had seen the unexpected change. They then tried to recognize the actor in a lineup. If change blindness were caused by a representation failure, then we would expect a close relationship between change detection and recognition performance.

1.1 *Change blindness and sparse representations*

Detecting a change implies the existence of an internal representation of the changed object or feature. The most intuitive explanation for change blindness is the absence of such a representation. According to this hypothesis, visual representations of either the pre-change or the post-change stimulus are insufficient (Simons 2000). If we assume that once the image or scene disappears, only abstract information remains (Horowitz and Wolfe 1998; O'Regan 1992; O'Regan and Noe 2001; Rensink 2000a), successful change detection would require an abstract representation of the changing feature from both the original and the modified scenes once they have disappeared.

¶ Now at the Department of Psychology, University of Illinois at Urbana Champain, 603 East Daniel Street, Champain, IL 61820, USA.

The view that change blindness results from a failure to represent a scene is consistent in some ways with object-file theory. Object files are thought to be detailed working-memory representations of real-world objects (Kahneman and Treisman 1984; Kahneman et al 1992). In addition, the theory suggests that people use these representations to detect changes. Kahneman and Treisman's original theory proposes that "focusing of attention on a target object not only enhances the salience of all its current properties—it also selectively reactivates the recent history of that object" (Kahneman et al 1992, page 176) and, "when the sensory situation changes the information in the files is updated, yielding the perceptual experience of changing or moving objects" (Kahneman et al 1992, page 178). This idea is bold because it gives the impression that, if we direct our attention to an object, then we perceive and retain detailed information about the properties of that object. In short, this theory implies that we have *detailed* representations of attended objects and that these representations help us to track objects over time.

Given its emphasis on detailed representations, the original version of object-file theory suggests that change blindness should not occur for attended objects. However, more recent research supports a modified version of the theory in which object files may contain abstract semantic information in addition to (or in place of) more concrete visual features (Gordon and Irwin 2000; Henderson 1994; Rensink 2000a; Rensink et al 2000; Wolfe and Bennett 1997). Thus, change blindness for attended objects is possible if the object file lacks information about the specific features that change, or if that information is abstract enough to encompass both versions of the changing feature. More important for present purposes, in both the original and the revised object-file theories, whatever information is in the object file is used to track the object over time. Therefore, a reasonable inference based on this theory is that, if a change to a central object goes undetected, information sufficient to distinguish the two versions of the attended object was not preserved in the object file.

1.2 *Change blindness and preserved representations*

Although representational sparseness is a possible cause for change blindness, recent research suggests that representations may be preserved in the face of change blindness. For example, in one study, an experimenter carrying a basketball approached a pedestrian and started a conversation (Simons et al 2002). While the two were talking, a group of students passed between them and surreptitiously took the basketball away. When asked if they noticed anything unusual or if anything had changed, only three observers mentioned the disappearance of the basketball. However, when specifically asked if the experimenter was carrying a basketball before the interruption, more than half of the observers claimed to recall the basketball. Most of these participants could describe the odd coloration of the ball, and none described it as a typical orange ball. Clearly, these observers represented the initial presence of the basketball and they were still looking at the experimenter after the ball was removed. Yet, they did not overtly mention the change until questioned by the experimenter. Although this is suggestive of preserved representations, there is no opportunity to compare observers who noticed the change to those who did not. Therefore, it is impossible to know whether these reports were, or were not, associated with success in detecting the change.

Another recent experiment did investigate the relationship between recognition performance and change detection with real-world person-changes (Levin et al 2002). Subjects who missed the change performed at chance levels on a lineup recognition task for the actors, whereas subjects who saw the change were significantly more successful. Although this is evidence for a representation failure, it occurred in a situation where subjects were not expecting to make any kind of response about the stimuli.

Therefore, subjects who missed the change had no reason to encode information, and so may have been particularly prone to a representational failure. In contrast, our experiments tested the relationship between recognition performance and change detection when subjects knew that they were part of an experiment and were told that they would have to answer questions about the stimulus. If the previously observed link between change detection and recognition performance were the result of this kind of very minimal real-world encoding, then it might be possible to observe a dissociation between the two in a situation where subjects may intentionally represent information.

In the present report, we tested whether change detection is necessarily associated with recognition memory for the changing objects. In all three experiments, observers first completed an incidental change detection task in which they saw a short videotape of one actor asking another for directions. While the 'lost' actor was speaking, the camera cut to a different angle and the change was introduced. Before watching the tape, observers were not specifically told to look for a change. After viewing the tape, observers were first asked if they had seen the change and then they tried to identify the pre-change actor in a four-alternative forced-choice lineup (4AFC).

The design of our experiments allows us to test three distinct predictions about the relationship between change detection and recognition. First, a failure on the part of observers who missed the change to recognize the pre-change object would be consistent with a representation failure. Second, a representation failure should cause observers who miss the change to be less accurate on the lineup than observers who saw the change. Note that intermediate cases are possible here: observers who miss the change may perform above chance on the lineup while still being less accurate than observers who saw the change. Finally, if representation failures underlie change blindness, variations across conditions and stimuli that affect change detection should similarly affect recognition performance. Specifically, we examine representation failures as tested by our recognition test (see section 5 for explanation).

2 Experiment 1

In experiment 1, observers watched a short video showing a change to a centrally attended object. In the video, a female actor asked a pedestrian for directions and, upon a cut in camera angle, something she was wearing or carrying changed in color. Observers, who did not know that a change would occur, were asked if they had seen the change. They then performed a lineup recognition task for the pre-change actor. If a representation failure underlies change blindness, observers who failed to detect the change should perform at chance on the lineup-recognition task. In addition, observers who missed the change should be less accurate on the recognition lineup than observers who detected the change.

2.1 *Method*

2.1.1 *Participants.* Forty-five introductory psychology students from Kent State University participated in small groups ranging in size from one to six. Observers had completed a separate, unrelated experiment, and were asked if they would remain an extra 5 min for the current experiment. All participants who were asked agreed to participate. Thirty-five participants received course credit in their General Psychology class and ten received a candy bar for their participation.

2.1.2 *Videos.* Each participant viewed one of eight video clips depicting one student asking another student for directions on campus. The 'lost' student was wearing a basketball jersey and an art smock covered in paint. She was also carrying a basketball and an art portfolio (figure 1). She asked one of two questions: "Excuse me, I am late for class, can you tell me where the art studio is?" or "Excuse me, I am late for practice,

(a) (b)

Figure 1. Still images from the video used in experiment 1. (a) An example of the pre-change stimulus. Notice the art portfolio (white). (b) An example of the post-change stimulus. Notice the art portfolio (tan).

can you tell me where the gym is?"(1) Upon a cut in camera angle, one of four features changed in color (basketball, basketball jersey, art portfolio, or art smock). The change occurred mid-sentence, while the lost student was talking. All of the objects were plainly visible in both the pre-change and the post-change views.

2.1.3 *Lineups.* Each of the eight video clips had a corresponding lineup. Lineups were constructed with one target and three distractors, each of which differed from the post-change picture by one feature.(2) The target picture displayed the female actor with the correct feature and the distractors in the lineup were identical to the post-change actor's appearance with the exception of a single feature (see figure 2).

2.1.4 *Procedure.* Observers were asked if they would be willing to watch a 'test video' that lasted about 1 min. They were given no instructions about what to look for in the video. After the video ended, observers were asked if they would be willing to answer a few questions pertaining to the video (see appendix). The questions and the criterion used in this experiment were the same as those used by Simons and Levin (1998). The first question was: "Did you notice anything unusual in this video clip?". The second question was: "Did you notice anything unusual about the person asking for directions?". These first two questions were broad in scope to give observers a chance to report the change without telling them a change had occurred. The third question was: "Did you notice that some of the features of the person asking for directions unexpectedly change from the first shot to the second?". If observers answered "yes" to this direct question and reported any feature that could be a plausible correct answer for the feature that changed, they were recorded as having detected the change (hit). After answering these questions, observers were shown a picture of

(1) In the first two experiments, two different questions were used to test the hypothesis that change detection would increase if questions were relevant to the feature that changed. For example, if the lost student asked where the gym was and the feature that changed was the basketball jersey or the basketball, participants would be more likely to see the change than if the art portfolio or art smock changed. No such differences were observed.

(2) By accident, some of the lineups used in experiment 1 were slightly flawed. The target and the distractors were wearing basketball jerseys with a different number than was shown in the video. Results should still be considered helpful because experiment 2 replicates the findings of experiment 1. In addition, this flaw would decrease lineup accuracy, thereby working against the comparison failure hypothesis.

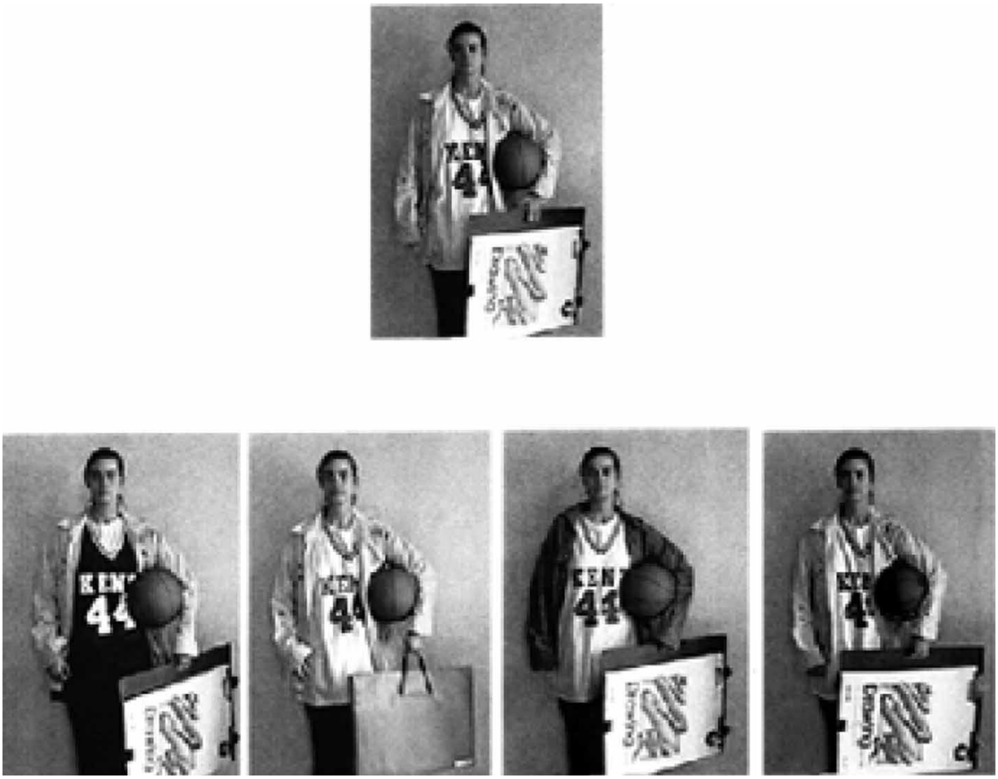

(a) (b) (c) (d)

Figure 2. Example lineup from experiment 1. This particular lineup corresponds to the video clip in which the art smock changed in color from blue to tan. The post-change picture is shown at the top. Participants were asked to select the pre-change picture from the four alternatives.

the lost student as she appeared at the end of the video (post-change) and were asked to choose what she looked like at the beginning of the video (pre-change) from among four possible pictures. All video clips were viewed on a standard 19-inch television at a distance of approximately 7.5 feet (2.25 m), with the display occupying approximately 12 deg of visual angle. Finally, each participant was debriefed and asked to sign a consent form allowing the use of his or her data.

2.2 *Results and discussion*

Only 6.7% of participants noticed the change. Among observers who missed the change, lineup accuracy was significantly above chance (45.2%, $\chi^2_{1,42} = 9.18$, $p < 0.01$). These data suggest that at least some observers who missed the change represented the pre-change stimulus.

This experiment shows that observers represented some detail from the video in the face of change blindness. Although the lineups were constructed to emphasize the features that changed, it is possible that a relative representation failure would distinguish observers who missed the change from those who saw it. Accordingly, both groups would have represented some features, but the factor that caused change blindness would be that observers who saw the change simply represented more features. In the present experiment, there is no way of comparing the hitters and missers because only three observers noticed the change. We sought to eliminate this problem in experiment 2 by increasing the detectability of the change.

3 Experiment 2

Experiment 2 was conducted to replicate the findings from experiment 1 and to determine if lineup accuracy varies as a function of change detection. The changes in method were meant to increase change detection, which would allow the possibility of comparing lineup performance between hitters and missers. Another benefit of increasing change detection was that it allowed us to test the relationship between the overall detectability of the change and the overall lineup performance. In experiment 2, we changed two features rather than one and we made each change more visually salient. We also added an identity change condition in which a new actor replaced the first (see Levin and Simons 1997). The representation-failure hypothesis predicts that with better change-detection performance, recognition accuracy should also increase, but only for observers who detected the change.

3.1 *Method*

3.1.1 *Participants.* A total of one-hundred-and-twenty-six Kent State University students participated in small groups ranging in size from one to seven. Alternate groups of observers were run in the feature-change and identity-change conditions for the first eighty-two observers: fifty-seven observers in the two-feature-change condition and twenty-five in the identity-change condition. The remaining forty-four observers all participated in the identity-change condition, for a total of sixty-nine observers in the identity-change condition. Participants either fulfilled a course requirement for General Psychology or received candy for their participation.

3.1.2 *Videos.* The general method for experiment 2 was similar to experiment 1 with a few exceptions. First, in the feature-change condition, we changed two features of the actor's appearance rather than just one (eg both the basketball jersey and the art smock changed color rather than just the basketball jersey). Second, the changes were made subjectively more salient; the basketball changed from orange to blue and white instead of orange to red and black; the jersey changed from blue to gold instead of black to white; and more paint was added to the art smocks. In the identity condition, a different actor replaced the original actor, but all of the clothing and objects were unchanged. There were a total of six video clips. Four depicted two-feature changes (in two clips the basketball and basketball jersey changed, and in the other two clips the art portfolio and art smock changed). The last two depicted a change in the identity of the central actor, counterbalanced for which actor appeared first. Given that our experiment is focused on visual change detection, we eliminated any auditory change in the identity videos by using a single audio track such that the voice remained constant throughout each video. The audio tracks were adjusted to allow sufficient synchronization with the video such that no observers detected anything unusual about the actors' voices.

3.1.3 *Lineups.* Each video clip was associated with a corresponding lineup. Observers were shown the post-change picture and they were asked to choose the correct pre-change picture from a set of four alternatives: one target and three distractors. The alternatives for the identity-change lineup were headshots of three same-age female foils and one headshot of the correct female actor. The alternatives for the feature-change lineup consisted of one target and three distractors that each differed from the post-change picture by two features. For example, in the target, both features had actually changed in the video; in one foil, the two features had not changed in the video; and in the other two foils, one feature had changed in the video and the other had not (see figure 3).

3.1.4 *Procedure.* The procedure was identical to that of experiment 1; observers answered the same set of questions and then completed the lineup task.

(a) Target
When compared to the post-change picture, two correct features are changed (the jersey and the basketball)

(b) Distractor 1
When compared to the post-change picture, two incorrect features are changed (the smock and the portfolio)

(c) Distractor 2
When compared to the post-change picture, one correct feature is changed (the jersey) and one incorrect feature is changed (the smock).

(c) Distractor 3
When compared to the post-change picture, one correct feature is changed (the basketball) and one incorrect feature is changed (the portfolio).

Figure 3. Example lineup from experiment 2. This particular lineup corresponds to the video clip in which the jersey and basketball changed color. The jersey changed from blue to gold and the basketball changed from blue and white to orange. The post-change picture is shown at the top. Participants were asked to select the pre-change picture from the four alternatives.

3.2 *Results and discussion*

Although we attempted to make the feature change more salient, across the four videos in the two-feature-change condition, only 12.3% of observers noticed the change. Consequently, it was again not possible to compare recognition performance of those observers who did and did not detect the change in the feature-change condition. Similar to experiment 1, 48% of participants who missed the change selected the correct choice in the lineup (see figure 4) significantly more than would be expected by chance ($\chi^2_{1,50} = 14.10$, $p < 0.001$).

Although the feature-change condition did not allow comparisons of those who did and did not detect the change, in the identity-change condition 56.5% of observers noticed the change. Inconsistent with the representation-failure hypothesis, those who noticed the change and those who missed the change were equally good on the recognition task (46.2% and 53.3%, respectively, $\chi^2_{1,69} = 0.35$, ns; see figure 4). Performance on the lineup-recognition task for both groups was significantly above what would be expected by chance ($\chi^2_{1,39} = 9.31$, $p < 0.01$ and $\chi^2_{1,30} = 12.84$, $p < 0.001$). The difference in change detection between the feature and identity conditions was significant, while the difference in lineup accuracy was not.

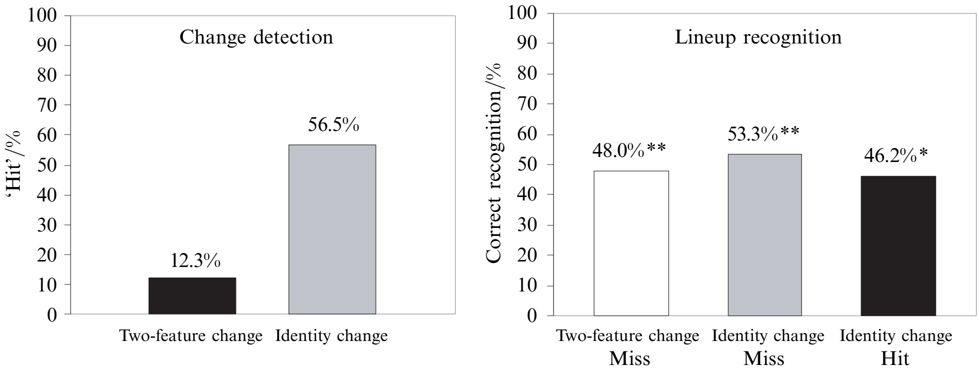

Figure 4. Change detection ('hit') and recognition results for the two conditions in experiment 2. These graphs show the percentage of participants who detected the change and the percentage of correct recognition on the lineup task (*$p < 0.01$, **$p < 0.001$). Note there are no data for participants who detected the two-feature change because there were only seven of them.

In addition, we conducted a lineup control experiment to ensure equality among our lineups. For the identity-change lineups we had participants watch the same videos as in the original experiment except the identity of the person remained the same throughout the video (it was always the post-change person). In addition, participants answered all of the same questions as in the original experiment (see appendix), including making a choice on a lineup. Once again, performance was not significantly different from chance of 25% correct ($M = 29.5\%$, $\chi^2_{1,78} = 0.838$, ns). For the feature-change lineups participants were told to imagine watching a video where a change occurred and that they did not see this change. They were then asked to choose or guess the correct pre-change actor from the lineups used in the original experiment. In this condition we were unable to create no-change videos as we did for the identity change lineups because we did not have video stimuli to do this. In the feature-change lineups performance was not significantly different from chance of 25% correct ($M = 14.2\%$, $\chi^2_{1,41} = 2.35$, ns).

3.3 *Summary*

As in experiment 1, we found support for preserved representations despite change blindness. Subjects who missed the change recognized the pre-change actor better than chance, in all conditions. Moreover, lineup accuracy was fairly constant and better than chance despite large differences in change detection performance. Change detection was 44% greater in the identity change than the two-feature-change condition, yet lineup accuracy was comparable (less than a 5% difference between the two conditions). Finally, in the identity-change condition there was no difference in lineup performance between the hitters and missers. This pattern of results suggests that observers represented the pre-change scenes, but it is possible that they did not compare them with the post-change scenes to determine whether a change had occurred.

4 Experiment 3

The identity-change condition from experiment 2 provides converging evidence against a broad representation failure hypothesis—not only did missers perform above chance on the lineup, but hitters and missers also did not differ on lineup accuracy. In both experiments 1 and 2, detection of the feature changes was too low to allow a comparison of hitters and missers. Although experiment 2 suggests that increased change detection is not associated with better recognition performance, the comparison was only possible by contrasting performance in the identity-change and feature-change conditions. These two conditions used different lineups which could account for the poorer than expected recognition performance in the identity-change condition. For example, the lineups in the identity-change condition might have been too difficult relative to the visibility of the change itself.

To alleviate these concerns, in experiment 3 we attempted to increase change detection in the feature-change condition by directing observers' attention to the changing features just prior to the change (Intraub 1981; Levin and Simons 1997; O'Regan et al 2000; Potter 1976; Rensink et al 1997). Observers viewed the videos while searching for a cue stimulus, which they were told might or might not appear on the body of the female actor or objects she was holding. For half of the observers, the cue appeared on one of the changing features just before the change occurred, and for the other half, no cue appeared. Finally, we asked observers to rate their confidence in their lineup choice to determine if a more sensitive measure would reveal differences between the hitters and missers. Again, the representation-failure hypothesis predicts that missers should be at chance on the lineup-recognition task, and that there should be a difference in recognition accuracy between hitters and missers.

4.1 *Method*

4.1.1 *Participants.* One hundred and fifteen students were recruited from the Kent State University Student Center. Four observers were excluded because they said there was a cue present when there was no cue. Five observers were excluded because they reported that they had heard about the experiment or they had participated in a similar experiment. Finally, one observer was excluded because her child was present while she was participating and the child told her that the change had happened. The remaining one hundred and five observers completed the experiment one at a time alternating between each of the two conditions. A total of fifty observers were in the cue condition and fifty-five were in the no-cue condition. Each volunteer was entered into a draw for a US $50 prize, which was held at the end of the semester.

4.1.2 *Videos.* Stimuli were similar to the two-feature-change videos used in experiment 2, in which two features changed (see figure 5). Four sets of video clips were used. Each set contained three variants of the same video in which two features changed. Each of the two changed features was cued in one clip, and in the third clip no cue appeared. For example, in one set, the basketball and jersey changed. This set included one clip with no cue, a second clip that cued the basketball change, and the third that cued the jersey change. Thus, altogether, there were 12 videos in the experiment

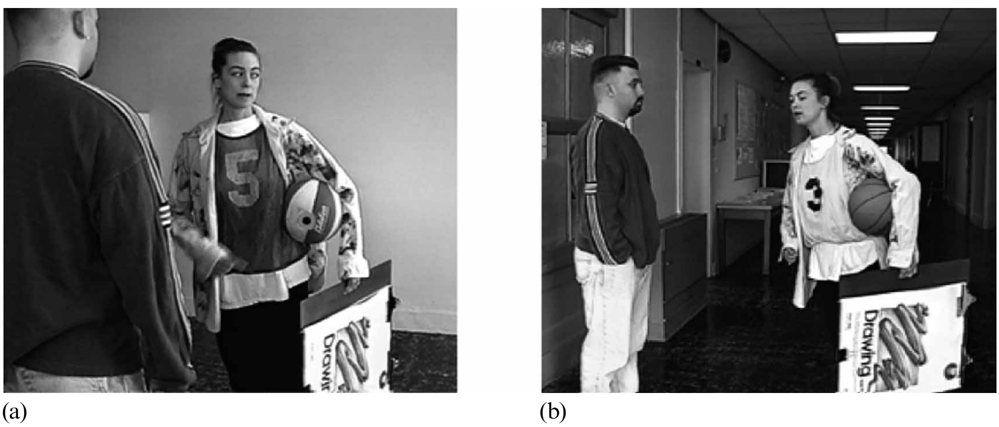

(a) (b)

Figure 5. Still images from the video used in experiment 3. (a) An example of the pre-change stimulus. Notice the basketball jersey (blue) and the basketball (blue and white). The cue appears 1 s before the cut in camera angle and signals the basketball change. (b) An example of the post-change stimulus. Notice the basketball jersey (yellow) and the basketball (orange).

(4 sets × 3 variants). When a cue appeared, it occurred 1 s before the feature changed and lasted for approximately 666 ms, offsetting 333 ms before the cut in camera angle. The cue (a red dot with a black center) was presented in different positions on the screen, depending on the feature that it cued, and it measured approximately 0.95 deg in diameter at a viewing distance of approximately 2.5 feet (0.75 m).

4.1.3 *Lineups.* Each set of three video clips was associated with a corresponding lineup. Observers were shown the post-change picture and they were asked to select the pre-change picture from a set of four alternatives. The alternatives in each lineup consisted of one target and three distractors, each of which differed from the post-change picture by one or two features. For example, the target differed by two features (both of which had actually changed in the video); one foil differed by two features (both of which had not changed in the video); and two foils differed by only one feature (which had actually changed in the video).

4.1.4 *Procedure.* Participants were told they would be watching a short video and they were asked to look for a red dot with a black center that may or may not appear on the female actor. They were also told that it could appear anywhere on the female actor, including her clothes or the objects she was holding. They were shown an example of the dot on paper and told to remember if they did or did not see it. Finally, they were told to watch the video until it ended because they would have to answer further questions about the video. All observers viewed one video clip and were randomly assigned either to the cue or the no-cue condition (participants in the no-cue condition were looking for a cue that did not appear). All video clips were presented on a standard 13-inch television at an approximate distance of 2.5 feet (0.75 m). The displays therefore subtended approximately 26 deg. After viewing the video clip, observers reported whether or not the cue had appeared and then answered the same set of questions from experiment 1. In addition, observers completed a lineup-recognition task and rated their confidence in their selection on a 1 (not very confident) to 7 (very confident) scale.

4.2 *Results*

Change-detection results were similar for the cue and no-cue conditions (46% and 47.3% of observers detected the change, respectively; see figure 6). Lineup accuracy results are also presented in figure 6 for all observers who did and did not notice the change. In the cue condition, all observers performed above chance (25%) on the recognition task regardless of whether they missed the change (59.3%, $\chi^2_{1,27} = 16.9$, $p < 0.001$) or detected it (56.5%, $\chi^2_{1,23} = 12.19$, $p < 0.001$), with no difference in accuracy between the two groups ($\chi^2_{1,50} = 0.038$, ns; see figure 6). In contrast, in the no-cue condition,

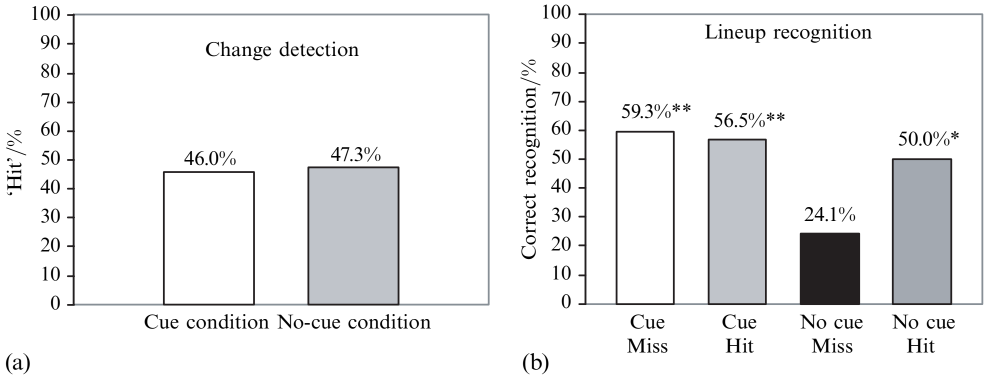

Figure 6. (a) The percentage of participants in each condition who detected the change for experiment 3. (b) The recognition accuracy on the lineup task for participants who did and did not detect the change in both the cue and the no-cue conditions (* $p < 0.01$; ** $p < 0.001$).

only observers who detected the change, performed above chance (50%, $\chi^2_{1,26} = 8.67$, $p < 0.01$). Those who missed the change performed with chance accuracy on the recognition task (24.1%, $\chi^2_{1,29} = 0.011$, ns). In addition, there was only a difference in lineup accuracy between the hitters and missers in the no-cue condition ($\chi^2_{1,55} = 3.96$, $p < 0.05$). Finally, there was no difference in lineup confidence ratings between the hitters (mean = 5.30) and missers (mean = 4.85), in the cue condition ($t_{54} = -1.03$, $p = 0.31$) or the no-cue condition (hitters mean = 4.85, missers mean = 4.86, $t_{48} = 0.04$, ns).

4.3 *Discussion*

Searching for a cue dramatically increased two-feature-change detection from that in experiment 2. Although the changes depicted in experiment 3 were the same as those in experiment 2, only 12.3% of observers noticed the change in experiment 2, whereas nearly half of the observers noticed the change in experiment 3 (46% in the cue condition and 47.3% in the no-cue condition). Most important, increased change-detection performance was not associated with improved lineup-recognition accuracy. Furthermore, as in the identity-change condition of experiment 2, recognition performance did not differ as a function of change detection when the change was cued. Once again, these findings provide evidence for preserved representations despite change blindness.

Interestingly, the results of the no-cue condition were consistent with the representation-failure hypothesis; observers who missed the change performed at chance on the lineup, whereas those who saw the change were correct approximately 50% of the time. As in an earlier series of experiments looking at lineup performance following a real-world identity change (Levin et al 2002), observers who missed the change showed chance recognition performance but those who detected the change did considerably better.

One possible explanation for the different results in the no-cue condition is that the cue task requires a working memory representation of the features of the cue itself. If observers in the no-cue condition must maintain those features throughout the film (because the cue never appears), then the capacity to represent other stimulus features would be reduced. In contrast, in the cue condition, the necessity of representing the stimulus is arguably terminated when the cue appears 1 s before the pre-change shot ends. Therefore, once observers have registered the cue they may have an opportunity to represent a few extra pre-change details even when they do not compare them across views.

5 General discussion

Three basic findings converge to support the hypothesis that change blindness is not caused by a failure to represent pre-change and post-change information. First, observers who missed a change were able to select the correct pre-change object from a lineup at above-chance levels in all conditions of all experiments except the no-cue condition of experiment 3. Second, in the identity condition of experiment 2 and the cue condition of experiment 3, observers who detected the change were no more accurate in the lineup task than observers who did not see the change. Finally, large between-experiment and between-condition differences in change-detection performance were associated with almost no difference in recognition accuracy. Change detection varied from 6% in experiment 1, to 56.5% in the identity-change condition of experiment 2, to 46%–47% in experiment 3; but recognition performance was roughly comparable across all of these conditions. Although most of the findings are inconsistent with the representation-failure explanation, the no-cue condition in experiment 3 was more consistent with the representation-failure hypothesis; observers who missed the change were less accurate in selecting the changed person from the lineup (performing at chance levels) than were those who saw the change.

Although representation failures are an intuitive explanation for change blindness, our results suggest that observers may also fail to detect changes if they fail to compare representations of the pre-change and post-change stimuli (Scott-Brown et al 2000; Simons 2000; Simons et al 2002). Evidence from other change-blindness research, as well as eyewitness recognition, suggest that observers have a representation of an initial set of features or objects that coexists with another version of that same event (McCloskey and Zaragoza 1985; Shapiro et al 1997). Viewers may represent information from the two presented scenes but assume the two views are consistent unless the meaning of the scene has changed or the experimenter triggers a comparison. Consequently, even if observers represent their visual world, they may still fail to detect changes if they do not examine the relationship between the two conflicting representations. This kind of failure is consistent with a number of findings in the concept and decisionmaking literatures. For example, people might simultaneously hold two beliefs that directly conflict with each other (Brewer and Samarapungavan 1991). The 'Moses illusion' also illustrates something akin to a comparison failure. In this case, people have difficulty detecting the inconsistencies in questions such as "How many animals of each kind did Moses take on the ark?", even though they know that Noah built the ark (Erickson and Mattson 1981; Reder and Kusbit 1991). Reder and Kusbit (1991) argue that this illusion occurs because people do not exhaustively match the features in the question with their existing knowledge of the answer.

As mentioned in the introduction, our conclusions about preserved representations must be limited to those representations that are consciously accessible (and to some degree durable) and that are tested by our recognition lineups. This constraint leaves open the possibility that change detection may be more closely correlated with other, perhaps implicit, representations. Although such implicit representations might underlie change detection (although see Mitroff et al 2002), we argue that the contribution of these data rests on the ecological validity of our 4AFC test, regardless of the particular mix of implicit and explicit representations that it reflects. Accordingly, the success or failure of a given representation must be closely indexed to the function that representation is presumed to serve. Otherwise, we will become caught in an endless argument about the degree to which some other test might reveal a difficult-to-specify representation existing somewhere in the cognitive system.

A related issue is that some versions of a representation-failure hypothesis might account for our data. As many authors have noted, visual representations do not all reside in a single neural subsystem. Accordingly, representations in one system might underlie change detection and representations in another system might underlie recognition performance (Chun and Nakayama 2000; Rensink 2000a, 2000b). For example, the ventral visual system is often presumed to underlie conscious object recognition while the dorsal system is presumed to direct object tracking and action-relevant perception (Goodale and Humphrey 2001; Mishkin et al 1983). If our recognition task relies on one system while the change-detection task relies on another, then it is possible that participants could have representations in one system that allow for accurate recognition performance despite insufficient representations in the change system. In contrast, participants could also have sufficient information in both systems, but then the relevant comparison failure would occur because the outputs from the two systems were never coordinated. According to this reasoning, we can rule out a complete representation failure, but our data cannot exclude the possibility that representations may exist elsewhere.

A more general version of the above argument is that the recognition test is inherently not a valid test of representations that might underlie change blindness. However, this option runs counter both to previous findings and to the no-cue findings in experiment 3, that 4AFC performance can be very closely tied to change detection.

In Levin et al (2002), observers in real-world tasks who failed to notice a change performed at chance levels on post-change lineups, while observers who saw the change were quite accurate (ranging from 63% to 81% correct on 4AFC lineups), suggesting a representation failure. In the real-world, observers may encode fewer features, focusing only on those that they believe to be central to the meaning of the scene. In contrast, observers in the laboratory know they are in an experiment and are specifically told to pay close attention to the stimuli and to remember them (they are told they will be asked questions about the videos in experiment 3). Consequently, they might employ a general strategy of attending to and remembering as much information as possible. An important implication of this argument is that feature representations can be manipulated independently of change detection, and that intentionally remembering visual features may not be sufficient for change detection.

A final concern about any experiment relying on explicit recognition from lineups is that the specific lineups might influence level of performance. In all cases, observers were shown the post-change stimulus and were asked to select the pre-change stimulus from the four alternatives. This choice inherently involves four contrasts between the post-change item and each lineup alternative. Hence, choosing the target involves a decision about what changed, more than a decision about the appearance of the pre-change feature. In a sense, this lack of specificity simplifies the representational demands of the task, because observers need not have represented the specific characteristics of the changed feature to succeed. Therefore, the amount of detail represented by participants who missed the change might have been overestimated by the lineup task. Similarly, the lineup task might underestimate the amount of detail represented by those participants who detected the change because it is too gross a measure to assess the completeness of their representation. For example, participants who missed the change might have only a vague impression of what might have changed, but that might have been sufficient to produce comparable performance to those who noticed the change. This concern, however, applies to any recognition task requiring selection from a set of distractors: all such designs test a specific target–distractor contrast, and it is always possible to argue that some other contrast goes untested. In addition, the concern that participants could rely on knowledge of what had changed applies more to the feature changes than to the identity changes. In the identity-change condition, the target and distractor items in the lineup were all faces. Consequently, observers might need a more detailed representation in order to detect changes. Yet, the identity-change lineups also yielded equivalent performance regardless of whether or not the change was detected.

An interesting implication of the finding that lineup performance is unaffected by change detection is that seeing a change does not guarantee that observers retain a representation of the pre-change object. One might reasonably expect the detection of a change to induce an orienting response, thereby leading to an activation of pre-change and post-change feature representations. Yet, if observers formed such a representation upon detecting a change, then those participants detecting a change should show superior recognition performance. This prediction is inconsistent with the results of the identity condition of experiment 2 and the cue condition of experiment 3. Accordingly, change detection does not guarantee a representation of the changed features; observers need a representation to detect a change, but the change-detection process does not guarantee subsequent access to that representation. This conclusion is supported by recent data showing that even when observers are intentionally searching for a change and know they will have to complete a recognition test for the pre-change and post-change items, they are often considerably less accurate in identifying the pre-change object than the post-change object (Beck and Levin 2003). At least under these conditions, observers cannot always create a durable representation of the pre-change object even when they are trying to do so.

6 Summary and conclusions

The pattern of our results suggests that change blindness is sometimes unrelated to recognition performance. It is not always the case that observers fail to detect a change because of a broad failure to represent any visual information. Even observers who missed a change had represented the pre-change stimulus sufficiently that they could often recognize the pre-change actor. Furthermore, differences in change detection were not related to differences in lineup performance. Therefore, representation, like attention, is necessary but not sufficient for change detection. To successfully detect a change, observers must both represent the pre-change features and compare them to the post-change features. These findings give the sense that online representations of scenes are similar to representations of beliefs and concepts in that we often fail to check our representations for internal consistency.

It is important to note that we do not argue that change detection and recognition will always be dissociated. Other findings suggest that change detection is, in some cases, closely related to recognition (Levin et al 2002). So, an effort to find *the* cause of change blindness would be misguided. Instead, a more useful approach would be to specify each element of the change-detection task, and to determine how and under what circumstances each subprocess might fail. This is a difficult task, but completing it effectively would lead to a well-articulated understanding of visual representation(s).

Acknowledgments. The contents of this paper were part of a master's thesis completed by Bonnie Angelone. The authors gratefully acknowledge Tina Smith and David Angelone for serving as actors in all the videos. Also, thank you to Bill Merriman, Maria Zaragoza, Jan Crowther, and Ron Rensink for reading an earlier version of this manuscript. Daniel J Simons was supported by NSF BCS-9905578 and by an Alfred P Sloan Research Fellowship.

References

Beck M R, Levin D T, 2003 "The role of representational volatility in recognizing pre- and post-change objects" *Perception & Psychophysics* **65** 458 – 468

Brewer W F, Samarapungavan A, 1991 "Children's theories vs. scientific theories: Differences in reasoning or differences in knowledge?", in *Cognition and the Symbolic Processes: Applied and Ecological Perspectives* Eds R R Hoffman, D S Palermo (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 209 – 232

Chun M M, Nakayama K, 2000 "On the functional role of implicit visual memory for the adaptive deployment of attention across scenes" *Visual Cognition* **7** 65 – 81

Erickson T D, Mattson M E, 1981 "From words to meaning: A semantic illusion" *Journal of Verbal Learning and Verbal Behavior* **20** 540 – 551

Goodale M A, Humphrey G K, 2001 "Separate visual systems for action and perception", in *Blackwell Handbook of Perception* Ed. E B Goldstein (Oxford: Blackwell) pp 311 – 343

Gordon R D, Irwin D E, 1996 "What's in an object file? Evidence from priming studies" *Perception & Psychophysics* **58** 1260 – 1277

Gordon R D, Irwin D E, 2000 "The role of physical and conceptual properties in preserving object continuity" *Journal of Experimental Psychology: Learning, Memory and Cognition* **26** 136 – 150

Henderson J M, 1994 "Two representational systems in dynamic visual identification" *Journal of Experimental Psychology: General* **123** 410 – 426

Horowitz T S, Wolfe J M, 1998 "Visual search has no memory" *Nature* **394** 575 – 577

Intraub H, 1981 "Rapid conceptual identification of sequentially presented pictures" *Journal of Experimental Psychology: Human Perception and Performance* **7** 604 – 610

Irwin D E, 1996 "Integrating information across saccadic eye movements" *Current Directions in Psychological Science* **5** 94 – 100

Kahneman D, Treisman A, 1984 "Changing views of attention and automaticity", in *Varieties of Attention* Eds R Parasuraman, D A Davies (New York: Academic Press) pp 29 – 61

Kahneman D, Treisman A, Gibbs B J, 1992 "The reviewing of object files: Object-specific integration of information" *Cognitive Psychology* **24** 175 – 219

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

Levin D T, Simons D J, Angelone B L, Chabris C F, 2002 "Memory for centrally attended changing objects in an incidental real-world change detection paradigm" *British Journal of Psychology* **93** 289 – 302

McCloskey M, Zaragoza M, 1985 "Postevent information and memory: Reply to Loftus, Schooler, and Wagenaar" *Journal of Experimental Psychology: General* **114** 381 – 387

Mishkin M, Ungerleider L G, Macko K A, 1983 "Object vision and spatial vision: Two cortical pathways" *Trends in Neurosciences* **6** 414 – 417

Mitroff S R, Simons D J, Franconeri S L, 2002 "The siren song of implicit change detection" *Journal of Experimental Psychology: Human Perception and Performance* **28** 798 – 815

O'Regan J K, 1992 "Solving the 'real' mysteries of visual perception: The world as an outside memory" *Canadian Journal of Psychology* **46** 461 – 488

O'Regan J K, Deubel H, Clark J J, Rensink R A, 2000 "Picture changes during blinks: Looking without seeing and seeing without looking" *Visual Cognition* **7** 191 – 211

O'Regan J K, Noe A, 2001 "A sensorimotor account of vision and visual consciousness" *Behavioral and Brain Sciences* **24** 939 – 1031

Pashler H, 1988 "Familiarity and visual change detection" *Perception & Psychophysics* **44** 369 – 378

Phillips W A, 1974 "On the distinction between sensory storage and short-term visual memory" *Perception & Psychophysics* **16** 283 – 290

Potter M C, 1976 "Short-term conceptual memory for pictures" *Journal of Experimental Psychology: Human Learning and Memory* **2** 509 – 522

Reder L M, Kusbit G W, 1991 "Locus of the Moses Illusion: Imperfect encoding, retrieval, or match?" *Journal of Memory and Language* **30** 385 – 406

Rensink R A, 2000a "The dynamic representation of scenes" *Visual Cognition* **7** 17 – 42

Rensink R A, 2000b "Seeing, sensing, and scrutinizing" *Vision Research* **40** 1469 – 1487

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

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

Scott-Brown K C, Baker M R, Orbach H S, 2000 "Comparison blindness" *Visual Cognition* **7** 253 – 267

Shapiro K L, Arnell K A, Raymond J E, 1997 "The attentional blink: A view on attention and a glimpse on consciousness" *Trends in Cognitive Sciences* **1** 291 – 296

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, Chabris C F, Schnur T, Levin D T, 2002 "Evidence for preserved representations in change blindness" *Consciousness and Cognition* **11** 78 – 97

Simons D J, Levin D T, 1997 "Change blindness" *Trends in Cognitive Sciences* **1** 261 – 267

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

Wolfe J M, Bennett S C, 1997 "Preattentive object files: Shapeless bundles of basic features" *Vision Research* **37** 25 – 43

Appendix

Questions used to determine if participants noticed the change. These questions start broad and become more focused so as to give participants maximum chances to answer that they had seen the change without telling them one had occurred.

Question 1: Did you notice anything unusual in this video clip? Y N

Question 2: Did you notice anything unusual about the person asking for directions? Y N

If so, what? ....

Question 3: Did you notice that some of the features of the person asking for directions unexpectedly change from the first shot to the second? Y N

If so, fill in what changed ....
Describe what it looked like before it changed ....
What color was it? ....

Question 4: (Lineup) SHOW LIST: There were two shots in this video; the second shot is at the top of the page. On the bottom of the page are four different versions of how she might have looked in the first shot. Which of these four best represents how she looked in the first shot?

Question 5: Have you ever heard about or participated in research like this before? Y N

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 32 2003

www.perceptionweb.com