
The role of iconic memory in change-detection tasks

Mark W Becker, Harold Pashler, Stuart M Anstis

Department of Psychology, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0109, USA; e-mail: mbecker@ucsd.edu

Received 12 August 1999, in revised form 20 December 1999

Abstract. In three experiments, subjects attempted to detect the change of a single item in a visually presented array of items. Subjects' ability to detect a change was greatly reduced if a blank interstimulus interval (ISI) was inserted between the original array and an array in which one item had changed ('change blindness'). However, change detection improved when the location of the change was cued during the blank ISI. This suggests that people represent more information of a scene than change blindness might suggest. We test two possible hypotheses why, in the absence of a cue, this representation fails to produce good change detection. The first claims that the intervening events employed to create change blindness result in multiple neural transients which co-occur with the to-be-detected change. Poor detection rates occur because a serial search of all the transient locations is required to detect the change, during which time the representation of the original scene fades. The second claims that the occurrence of the second frame overwrites the representation of the first frame, unless that information is insulated against overwriting by attention. The results support the second hypothesis. We conclude that people may have a fairly rich visual representation of a scene *while the scene is present*, but fail to detect changes because they lack the ability to simultaneously represent two complete visual representations.

1 Introduction

Many recent experiments have shown that people are unable to detect large changes in visual scenes. This 'change blindness' is dramatic and robust and has been demonstrated in a number of ways (for review see Simons and Levin 1997). It occurs when two successive scenes are separated by a blank interval (Pashler 1988; Phillips and Singer 1974; Rensink et al 1997; Stelmach et al 1984), saccades (Grimes 1996; McConkie and Currie 1996; McConkie and Zola 1979), a camera cut (Levin and Simons 1997), or an occluding object (Simons and Levin 1998).

The techniques used to induce change blindness have varied, yet all the methods include an event separating the successive pictures. When the two pictures are not separated by such an event, an observer's ability to detect change is often nearly perfect. What do all of these intervening events have in common that could explain the resulting inability to notice a visual change?

One possible answer is that each manipulation generates multiple transients which occur at the same instant as the to-be-detected change. O'Regan et al (1999) showed that introducing multiple transients was sufficient to produce change blindness, even when these transients did not overlap the location of the change. In their experiment, the original and altered picture were presented in succession with no interstimulus interval (ISI). Observers failed to notice an alteration of the picture when it occurred simultaneously with the appearance of six ink blotches which were scattered around the scene. The addition of multiple transients (the blotches) which occur with the to-be-detected change induced change blindness.

The importance of visual transients for the detection of visual change has long been known. Phillips and Singer (1974) first reported that changes become difficult to detect when a blank ISI is inserted between the two scenes. In their experiment, they found that the longer the blank ISI, the worse people did at change detection. To explain this,

they proposed a model of change detection in which successful detections relied on the presence of a distinctive neural transient at the location of the change (Phillips and Singer 1974; Stelmach et al 1984). With zero ISI, changes were easily detected because the transient caused by the change occurred in isolation. However, with a long enough blank ISI, each item produced a transient when the second scene appeared. The transient caused by the change no longer occurred in isolation and thus was difficult to detect. In this early work, the initial scene was a display of fifty dots, and the change was an addition or deletion of a single dot. Since all of the stimuli were identical, subjects could not rely on information about the identity of particular stimuli to detect the change.

Recent experiments with character arrays (Pashler 1988) and pictures of scenes (Rensink et al 1997) used stimuli that have unique identities. Here an observer might potentially use information about the identity of items in a scene to detect a change. For instance, if a subject identified one of the items in the first picture as a dog, and then could not identify a dog in the second picture, he/she could infer that something had changed, solely on the basis of this identity information. However, even with these more complex stimuli people have difficulty detecting changes. It seems people can only use identity information to detect a change when they are attending to the item as it changes. Further, people seem to be able to monitor only about four items for change (Pashler 1988; Rensink 1997). The inability to use identity information to detect changes in unattended items has been interpreted as evidence that we do not represent items to which we are not currently attending (Rensink 1997; Rensink et al 1997). In short, these findings suggest that, unless we attend to an item, we retain insufficient information about the item to perform a change detection across extremely brief time periods such as an 80 ms blank frame inserted between two pictures.

However, many studies on iconic memory suggest that we have access to a great deal of visual information which lasts for at least 150 ms after the offset of stimuli. During this time, if researchers focus the subject's attention on an item, the subject can extract the item's identity (Averbach and Coriell 1961; Coltheart et al 1974; DiLollo and Dixon 1988; Gegenfurtner and Sperling 1993; Irwin and Yeomans 1986; Loftus et al 1985, 1992; Sperling 1960). These discrepant conclusions create an interesting conundrum. If it is true that 80 ms after a scene disappears we represent insufficient visual information to perform a change detection based on identity information, how does one explain the iconic-memory finding that, at similar delays, a person can still extract the identity from the iconic representation? Alternatively, if an iconic memory containing sufficient information for item identification exists throughout the blank ISI, why are people so poor at using this information to identify changes?

We begin to address this question by asking whether, during the blank ISI of a change-detection task, there is an iconic image which contains the information needed for accurate change detection. In experiment 1 we investigate this question by cueing the location of the change during this blank ISI.

2 General methods

2.1 Apparatus

All experiments were programmed in Macromedia Director and run on a Power Macintosh 400 with a 15-inch display running at 66 Hz.

2.2 Stimuli

Subjects saw a single array of six items equally spaced around a clock face with a diameter of 8 deg. All stimuli were presented on a white background. For a single trial, the six stimuli were either all black letters, all black symbols from a traditional typewriter keyboard, or all disks of different colors. Each stimulus subtended approximately 1.5 deg. After some ISI, the initial display was replaced by a test display.

In no-change trials the test display was identical to the initial display. In change trials the test display was identical to the initial display except that a single item from the initial display was replaced by a different item from the same class of stimuli (ie a letter was replaced by a different letter). The three classes of stimuli (letters, symbols, and colors) cycled so that every fourth trial contained the same class of stimuli. This rotation was designed to reduce possible proactive interference between adjacent trials.

On some trials, a spatial-location cue was used to focus the subject's attention on a specific item in the display. This spatial-location cue was a red radial line that pointed to one of the items as the hour hand of a clock points to the time. The red line did not spatially overlap with the items of the search array. Subjects were told that when a cue appeared they needed to attend only to the location to which the cue pointed ("if a change occurs it will occur at the cued location; however, the presence of a cue does not necessarily mean that there will be a change"). In change trials this cue always pointed to the location of the changed item. In no-change trials its location was random.

2.3 Procedure

After the experimenter explained the task to the subject, the subject saw all thirty possible stimuli and named each one, so that the experimenter knew what the subject was calling each color and symbol.

Subjects then participated in a practice block of thirty trials, after which the experiment began. A single trial consisted of a central red fixation spot that was displayed for 2 s. Then six stimuli appeared around the fixation point for 256 ms. The six items disappeared, leaving just the fixation point during the ISI. If a cue appeared during the ISI, it remained on for the rest of the trial. After the ISI, six items reappeared around the fixation point and remained on the screen until the subject responded.

Subjects responded by saying "no change" or pointing with their finger to the location of the change. If the subject detected a change he/she also verbally identified the original item in the change location. For example if '7' became '4' the subject responded by pointing to the '4' and saying '7'. This identity response allowed an assessment of whether the subject detected the change by comparing the identity of the two items ('7' and '4').

The experimenter input the subjects' responses via the keyboard. All stimuli and the cue, if there was one, were erased leaving just the fixation point. The next trial began with 2 s with just the fixation point.

3 Experiment 1

3.1 Method

The first stimulus array was presented for 256 ms. The ISI between the first array and second array was 281 ms. One third of the trials (no-change) contained second arrays identical to the first array. The remaining two thirds of the trials (change) contained a second array with a replacement of a single item from the first array.

Half of the trials had a red radial line cue pointing to one of the item locations. This cue appeared either 16 ms, 82 ms, 149 ms, 215 ms, or 281 ms (simultaneous with frame 2 onset) after the offset of frame 1. The cue remained on until the subject responded. A schematic diagram of the method used in experiment 1 is shown in figure 1.

3.2 Results

In figure 2, the percentage of correct change detections and correct original item identifications at each cue delay is plotted against the performance without a cue. An analysis of variance of the detection data showed that there was a main effect of cue condition ($F_{5,45} = 4.4$, $p < 0.01$). Adding a cue significantly ($p < 0.05$) improved detection when the ISI between the offset of frame 1 and the onset of the cue was 16 ms ($F_{1,45} = 17.804$), 82 ms ($F_{1,45} = 12.15$), 149 ms ($F_{1,45} = 4.846$), and 215 ms ($F_{1,45} = 8.616$), but not when the cue appeared 281 ms after the offset of the first frame ($F_{1,45} = 3.037$, $p > 0.05$).

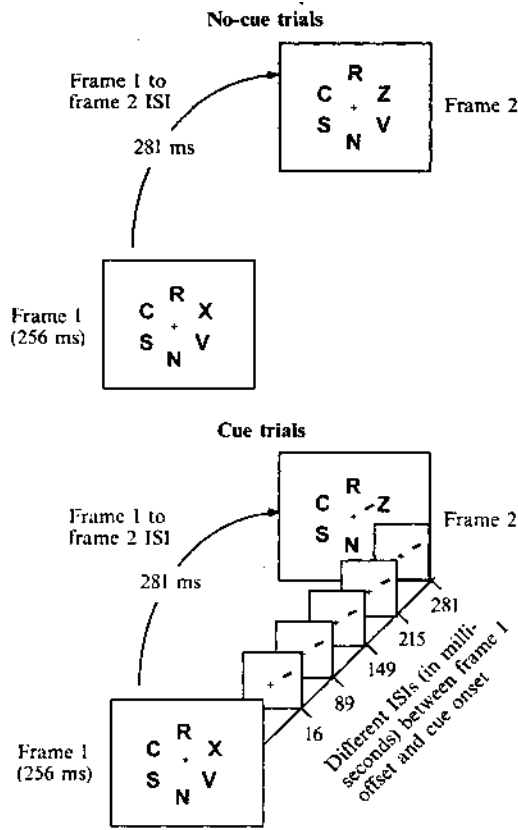


Figure 1. Method for experiment 1. The ISI between frame 1 and frame 2 is held constant at 281 ms. On cued trials, the timing of the onset of the cue varies from 16 ms to 281 ms depending on the condition of the trial.

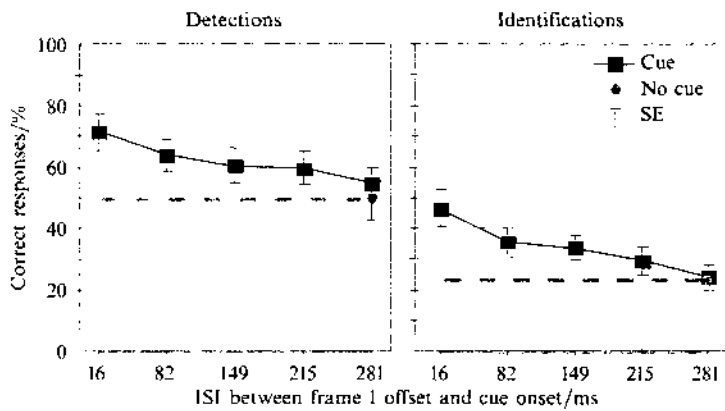


Figure 2. Results from experiment 1. Percentage of correct responses for the change-detection task and the original item-identification task for each cue delay. The performance without a cue for each task is also plotted as a line to make the comparison between cued performance and performance without a cue more clear. The bars show standard error.

An analysis of variance on the identity data also showed a significant main effect for cue condition ($F_{5,45} = 4.605, p < 0.01$). Adding a cue also improved identification when the ISI between the offset of frame 1 and the onset of the cue was 16 ms ($F_{1,45} = 18.14, p < 0.05$) and 82 ms ($F_{1,45} = 4.327, p < 0.05$), but not for the longer ISIs (149, 215, and 281 ms).

3.3 Discussion

Cueing the location of a change during the blank ISI enhances the ability both to detect changes and to identify the original item in a change location. The data are consistent with studies (Averbach and Coriell 1961; Gegenfurtner and Sperling 1993; Sperling 1960) which suggest that there is an iconic-memory trace of the first array that persists during the blank ISI. If attention is switched to the appropriate location by a cue, information from this iconic image can be used to detect changes. Additionally, the number of identifications of the original items increased in step with the number of overall change detections, suggesting that the cue aids change detection by allowing the observer to compare the identity of the item from the first scene with the item in the second scene. The number of changes detected and items identified decreased as the time between the offset of frame 1 and the onset of the cue increased, presumably as the icon of frame 1 faded.

We conclude that, during the blank ISI of a change-detection task, there is an internal representation which is sufficiently detailed for a subject to detect a change on the basis of identity information. Given this result, we must explain why this iconic representation fails to facilitate change detection when no cue is present. Two possible hypotheses which could explain this failure are outlined below.

3.4 Hypothesis 1: Change blindness due to a fading iconic image

Rensink et al (1997) concluded that the iconic memory of the first scene is masked by the second scene. They based this conclusion on the finding that an 80 ms blank ISI inserted between the original and altered picture produced change blindness. They reasoned that iconic memory should last for longer than this 80 ms blank frame, and interpreted people's failure to notice changes with such short ISIs as evidence of the masking of iconic memory. While this is one possible interpretation of their finding, both their theory of change blindness and their findings do not rule out the possibility that change blindness results from a fading iconic image which is not completely masked by the onset of a second picture.

The main premise of their theory is that "visual perception of change in an object occurs only when that object is given focused attention" (Rensink et al 1997, page 372). This is said to apply to both 'normal' (zero ISI separating the images) change situations and also to changes which occur when there are multiple transients which co-occur with the change. They claim that change-detection performance for normal conditions is excellent because the motion transient caused by the change captures attention. With attention quickly drawn to the change item, the change is detected. However, when the change co-occurs with multiple transients, there is not a single clear transient to attract attention to the site of the change; thus a slow item-by-item search begins and produces change blindness.

Within this theoretical framework, we believe that their findings are wholly consistent with a rapidly fading iconic image which is not overwritten by the occurrence of the second scene. The only difference between a normal change and one which occurs simultaneously with other transients is that the position of the change is signaled by the lone transient in the first case, whereas the multiple transients create positional uncertainty in the latter case. Thus changes which occur in isolation can be thought of as analogous to the partial-report paradigm used by Sperling (1960). The sole transient, instead of a tone, cues attention to the correct location and allows access to the iconic information of the first scene before it fades. Similarly, changes which co-occur with other transients are analogous to the whole-report paradigm. Without a cue to guide attention, the observer must search each item. Poor change detection results from the failure of this search to reach the location of the change before the iconic trace of the first frame has faded. This interpretation would be consistent with Rensink et al's theory of change blindness and fit their data, but does not invoke masking.

In addition, if the appearance of the second frame overwrites the iconic image of the first scene, it is not clear how Rensink et al's (1997) theory would predict excellent change detection with near-zero ISI changes. The original theory suggests that attention is necessary for change detections because it "lets [items] be entered into a relatively durable store, such as visual short-term memory, so that comparisons can be made" (Rensink et al 1997, page 372). Further, the postulated role of attention is the same for changes which occur in isolation and changes which occur with other transients; the only difference is how attention is allocated. Logically, during a normal change detection the transient which attracts attention to the site of the change only occurs once the second frame appears; thus attention cannot be directed to the site of the change until the arrival of the second scene. If this second scene masks the first scene's iconic image, there would be no information about the first scene remaining for attention to enter into a durable store. Thus one would predict poor detection rates.

Some may discount this hypothesis because it requires iconic information to be retained past the onset of the second frame. These critics may point out that many iconic-memory experiments have shown that iconic memory is masked by the onset of second-frame stimuli (Gegenfurtner and Sperling 1993; Irwin and Yeomans 1986; Loftus et al 1992). However, the task in typical iconic-memory experiments is to name the original item. This task requires that enough information persists about the original item to make an identification. It is possible that the appearance of a second frame interferes with the iconic image of the original frame to a degree which makes identification impossible, but still retains enough information about the first scene to support change detection. This argument relies on the assumption that identification of the original item requires more information than noticing that the new item is in some way different from its predecessor. The assumption is logically appealing and supported by the results of experiment 1 which show that people are better at detecting changes than identifying items.

In short, an account which explains change blindness in terms of a fading iconic image that is either cued, by a lone transient in the zero-ISI case, or uncued, when multiple transients co-occur with the change, seems plausible and would provide a parsimonious explanation of the change-blindness findings. This account is tested in experiment 2.

4 Experiment 2

If the inability to detect a change which co-occurs with other transients is a result of the iconic trace decaying while a serial or limited capacity search is conducted, a spatial-location cue during the second frame (instead of during the ISI) should improve performance by eliminating the need for a search.

4.1 Method

The method was identical to experiment 1 except that the blank ISI was reduced to 82 ms, and the cue, when present, appeared simultaneously with the second array. In addition, there was a zero-ISI control condition. Experiment 1 contained a condition in which the cue was simultaneous with the second frame; however, that condition's ISI of 281 ms may have been so long that the iconic image of the first array faded prior to the onset of frame 2. The 82 ms ISI should be short enough for the iconic image of the first array to be still strong when the second array appears. A schematic diagram of the method used in experiment 2 is shown in figure 3.

4.2 Results

The cue did not improve performance. The percentage of correct change detections and original object identifications is plotted in figure 4 for the three conditions (zero-ISI, ISI, and ISI with cue). Detection shows a significant effect of condition

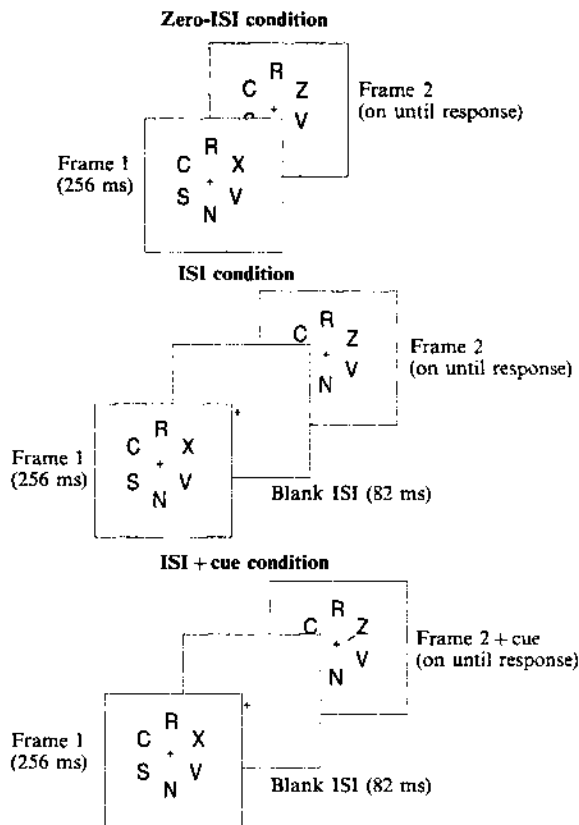


Figure 3. Method for experiment 2. The control condition has a 0 ms ISI. There are two conditions with an 82 ms ISI, one with and one without a cue that appears with the onset of frame 2.

($F_{2,32} = 64.319, p < 0.01$). Detection in the zero-ISI condition was significantly better than in the ISI condition without a cue ($F_{1,32} = 103.2, p < 0.01$), and the ISI condition with a cue ($F_{1,32} = 89.243, p < 0.01$). The ISI condition with a cue did not differ from the ISI condition without a cue ($F_{1,32} = 0.507, p > 0.05$).

Similarly planned comparisons of the identification data show no significant difference between the two ISI conditions ($F_{1,32} = 0.0, p > 0.05$), but show both the ISI conditions differing significantly from the zero-ISI condition ($F_{1,32} = 9.693, p < 0.01$).

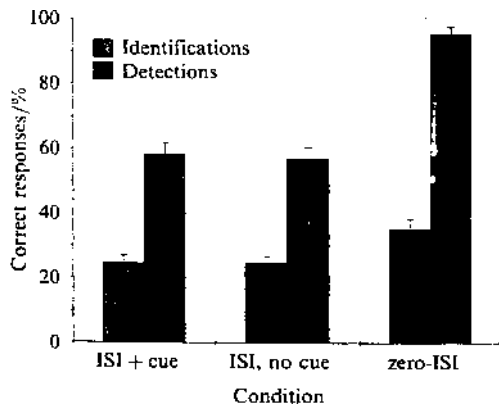


Figure 4. Results from experiment 2. Percentage of correct change detections and original item identifications for a condition with an 82 ms ISI and a cue in frame 2, a condition with an 82 ms ISI without a cue in frame 2, and a condition with a 0 ms ISI. The bars show standard error.

4.3 Discussion

The fading-iconic-image hypothesis of change blindness claimed that, with a 82 ms ISI, people fail to detect changes because the iconic representation of the first frame decays while a limited-capacity attentional mechanism is searching the location of every transient. If this hypothesis were true, then a cue in the second frame should have substantially improved detection by eliminating the need for a serial search. However, neither detection nor identification was improved by cueing the location of the change when the cue appeared simultaneously with the second frame. This finding suggests that the detection failures observed in traditional change-blindness studies are not a result of the positional uncertainty introduced by the multiple transients which occur with the to-be-detected change.

In addition, the finding that adding a cue to the second frame of the ISI condition fails to improve performance suggests that the cue is not accessing enough additional information to produce an effect. The finding that this cue is ineffective supports Rensink et al's (1997) claim that the appearance of the second frame overwrites the iconic-image representation of the first array. The iconic system seems competitive, with new information replacing or masking old information. But if this is true, why does the zero-ISI condition produce more identifications of the original item than either of the ISI conditions? It is possible that the lack of an interstimulus interval increases identification by providing motion cues which can improve guessing.

A second possibility is that the 82 ms ISI is too long of a delay and the iconic representation has faded to the degree it cannot support change detection. In experiment 3 we further examine whether the appearance of the second scene overwrites the representation of the first frame.

4.4 Hypothesis 2: Change blindness due to task demands

A second explanation why change blindness depends on the occurrence of multiple transients rests on the notion that the addition of multiple transients changes the nature of the task. When the only transient in the display is caused by the to-be-detected change, as is the case with zero ISI, the detection of any transient is enough information to make the change detection. As such, this type of change detection requires very little information about the previous display and can be made on the basis of very-low-level motion detectors. However, if there are multiple transients co-occurring the nature of the task changes from a pure detection of a transient to a discrimination of whether one of the transients is caused by a change of form—a task which requires a comparison to previous information.

Note that this hypothesis is generally in agreement with Rensink et al's (1997) proposal that attention is needed to detect changes which occur when an ISI separates the two images. However, it suggests that under zero-ISI conditions, the change-detection task becomes a much easier task which does not require that attention move information from the first display into a more durable store. Instead, for these changes, the simple detection of a transient is sufficient for the correct detection of change.

Under this task-demand explanation of change blindness it is possible that there is an iconic representation persisting through the interstimulus interval, but that the system which contains this image is subject to the constraint that new information at a given location overwrites old information (Averbach and Coriell 1961; Sperling 1960). On this view, the iconic-memory representation would be useless for performing a task which requires a comparison across time, such as an 80 ms ISI change-detection task. This hypothesis assumes that there is a representation of unattended parts of a scene while the scene is viewed, and for a time thereafter (if nothing else follows it). However, since this representation is held in a single buffer, it is incapable of supporting change detection, unless the necessary information is transferred via attention to a more durable memory

system, which is not subject to erasure by new visual information (Gegenfurtner and Sperling 1993; Rensink et al 1997).

This hypothesis would also predict that 'mud splashes' hitting the screen would disrupt change detection (O'Regan et al 1999). First, the co-occurrence of multiple transients would shift the task from a pure detection of a transient to the discrimination of what type of transient is involved, increasing the informational demands of the task. Second, the new item is in the same location as the item in the first frame and would mask the iconic image of the first frame.

If this hypothesis is true, we may have a great deal of visual information on-line, but may lack the ability to store this information. In experiment 3, we used a spatial-location cue to test this hypothesis.

5 Experiment 3 5.1

Method

In experiment 3, a spatial-location cue was once again used. However, the cue always appeared 16 ms after the offset of the first array, at which time the iconic trace of the first array should have been very strong. While the latency between the first-array offset and the cue was held constant at 16 ms, the latency between the cue onset and the onset of the second array varied. The second array either appeared simultaneously with the cue or followed the cue onset by 66 ms, 133 ms, 199 ms, or 256 ms (figure 5). Varying this interval also necessarily affected the amount of time between the offset of the first array and the onset of the second array, so that the ISI between frames was either 16 ms, 82 ms, 149 ms, 215 ms, or 281 ms.

On the basis of previous change-detection studies with short ISIs, one should expect detection to fall off as the ISI becomes longer (Phillips and Singer 1974; Stelmach et al 1984). However, another possibility exists. Suppose there is an iconic image of the initial scene and, by attending to a specific item in the icon, an item can be moved to a more stable representation (short-term memory). Further, suppose that the attentional mechanism takes time to move information from the visual buffer into short-term memory. Then it follows that a cue, which directs attention to the change location, will be more effective if the time between the onset of the cue and the appearance of frame 2 is longer. The additional time will allow the attentional mechanism to transfer information from the icon into short-term memory before the icon is overwritten by the appearance of the second display. Thus, if the cue always appears 16 ms after the offset of frame 1, the ability to detect changes should increase, not decrease, as the ISI between frame 1 and frame 2 increases.

5.2 Results

Both detection and identification performance improved as the time between the cue and the onset of the second frame increased (figure 6). An analysis of variance of the detection data shows that the cue condition is significantly better than the no-cue condition ($F_{1,9} = 60.307, p < 0.01$). In addition, there is a significant interaction between the ISI condition and whether the cue was present or not ($F_{4,36} = 4.559, p < 0.01$). Planned comparisons between the cued versus uncued trials at each ISI show that the cue condition was not significantly better than the no-cue condition when the second frame appeared simultaneously with the cue (16 ms after the offset of the first frame) ($F_{1,36} = 0.0, p > 0.05$). However, the cued condition produced significantly ($p < 0.05$) better detection rates than the uncued condition when the cue onset preceded the second-frame onset by 66 ms ($F_{1,36} = 13.352$), 133 ms ($F_{1,36} = 12.476$), 199 ms ($F_{1,36} = 12.476$), and 256 ms ($F_{1,36} = 35.617$).

Similar results were obtained for the identification data. There was a significant main effect of cueing ($F_{1,9} = 58.522, p < 0.01$), and the interaction between cueing

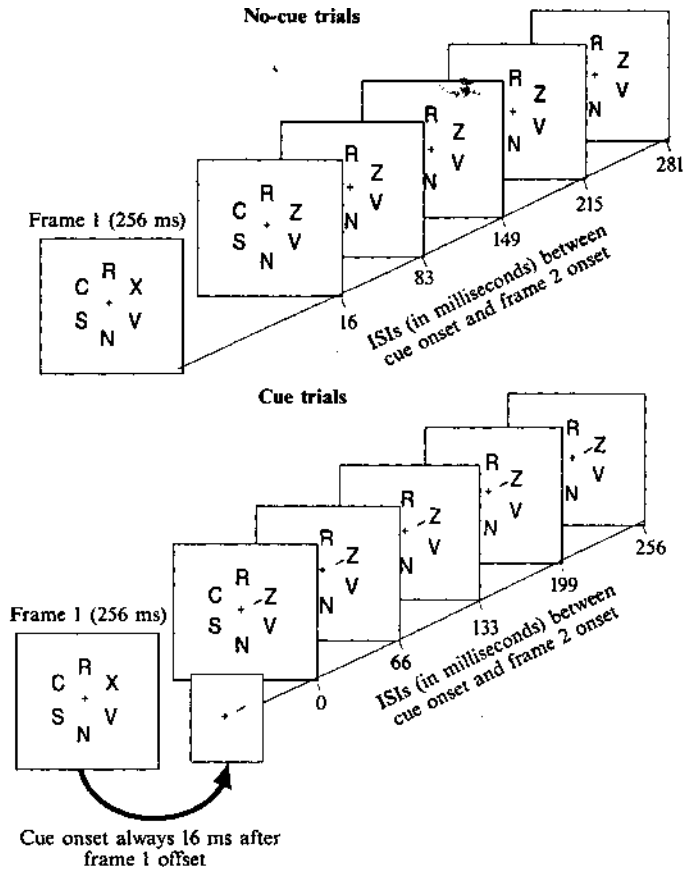


Figure 5. Method for experiment 3. In cue trials the cue always follows the offset of frame 1 by 16 ms. The ISI between frame 1 offset and frame 2 onset varies from 16 ms (simultaneous with the cue) to 281 ms depending on condition. There are also no-cue trials at each frame 1 to frame 2 ISI (not shown in the figure).

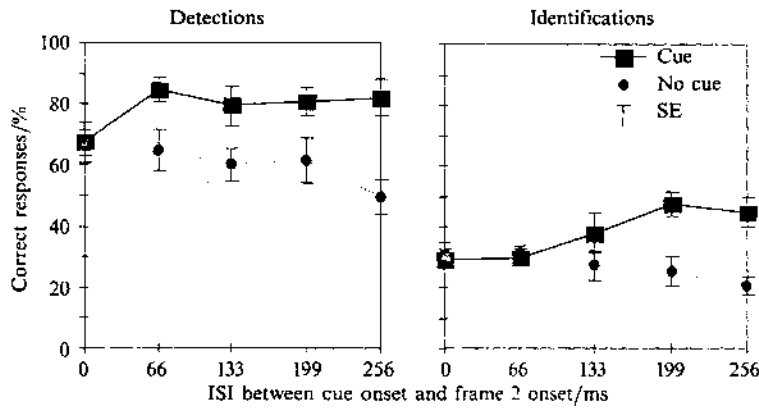


Figure 6. Results from experiment 3. The percentage of correct detections and original item identifications for each ISI between cue onset and frame 2 onset. These data are plotted against a no-cue condition with the same ISI between frame 1 and frame 2. The bars show standard error.

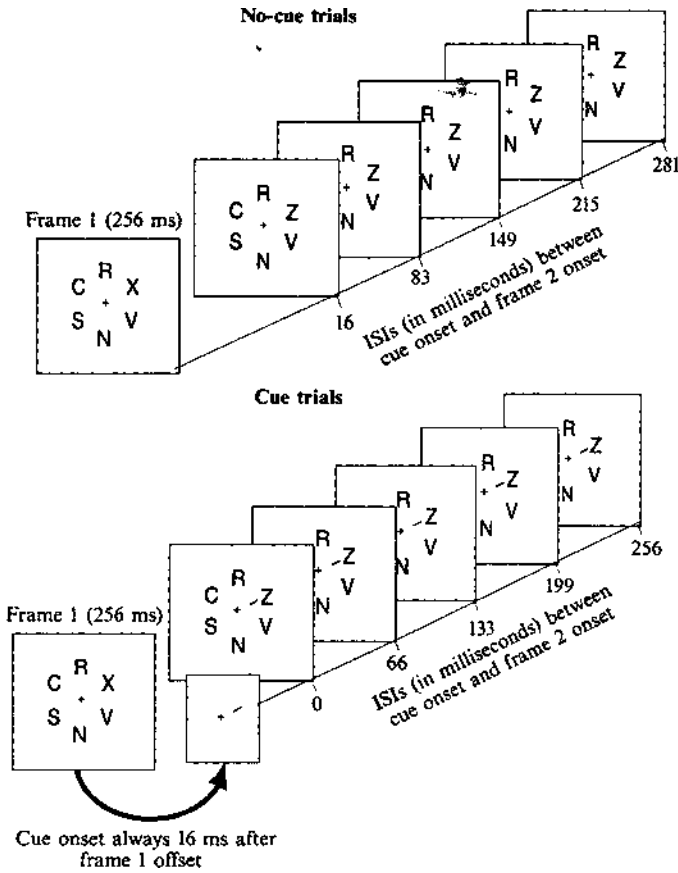


Figure 5. Method for experiment 3. In cue trials the cue always follows the offset of frame 1 by 16 ms. The ISI between frame 1 offset and frame 2 onset varies from 16 ms (simultaneous with the cue) to 281 ms depending on condition. There are also no-cue trials at each frame 1 to frame 2 ISI (not shown in the figure).

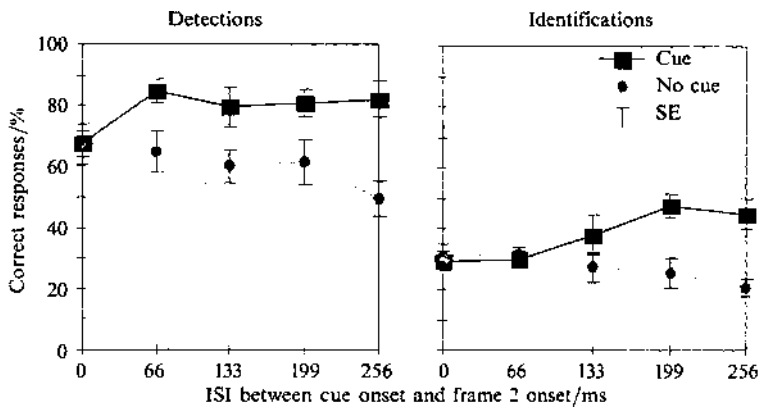


Figure 6. Results from experiment 3. The percentage of correct detections and original item identifications for each ISI between cue onset and frame 2 onset. These data are plotted against a no-cue condition with the same ISI between frame 1 and frame 2. The bars show standard error.

and ISI condition was significant ($F_{4,36} = 4.665, p < 0.01$). Planned comparisons between the cued and uncued trials at each ISI show that the cue condition is significantly better than the no-cue condition when the cue onset preceded the second-frame onset by 199 ms ($F_{1,36} = 15.2, p < 0.01$), or by 256 ms ($F_{1,36} = 18.089, p < 0.01$).

Planned comparisons within the cue condition show that the 66 ms cued condition has significantly more detections than the 16 ms (simultaneous with frame 2) condition ($F_{1,36} = 22.479, p < 0.01$). In fact, for detections, all delayed-cue ISIs were significantly different from the 16 ms cued ISI and none of them was significantly different from another. For the identifications, however, the 66 ms ISI cued condition was not significantly different from the 16 ms cued condition ($F_{1,36} = 0.033, p > 0.1$), but the 133 ms cued condition was greater than the 16 ms cued condition ($F_{1,36} = 5.62, p < 0.05$). All ISIs longer than 133 ms also produced significantly more identifications than the 16 ms ISI condition and none was significantly different from another.

Planned comparisons within the uncued condition show that performance decreased with greater ISI. The 16 ms and 66 ms uncued conditions had significantly more detections than the 256 ms condition ($F_{1,36} = 10.814, p < 0.01$; and $F_{1,36} = 7.847, p < 0.01$). For the identifications, the 16 ms and 66 ms uncued conditions were not different from one another. When they were pooled together, these conditions produced significantly better identification performance than the 256 ms ISI ($F_{1,36} = 4.471, p < 0.05$).

5.3 Discussion

Most studies looking at ISIs in the range we used have found that increasing the time between the presentation of two pictures decreases people's ability to detect a discrepancy between them (Phillips and Singer 1974; Stelmach et al 1984). Experiment 3 replicated this finding when there was no cue. However, when a cue was presented 16 ms into the blank ISI, the ability to detect a change actually increased with greater ISI. This pattern of results in the cued condition is very informative. First, a cue which occurs 16 ms after the offset of frame 1 increased change detections. This finding provides evidence that the nervous system contains a more detailed (but not necessarily consciously accessible) representation of the scene than the notion of change blindness might lead one to believe. Second, the finding that a cue occurring at the same 16 ms delay is not beneficial when it appears simultaneously with the onset of frame 2 suggests that the representation of the original scene is overwritten or severely degraded by the onset of the second scene. Finally, the finding that increasing the time between the cue and the second frame improves performance suggests that the attentional process can, when given enough time, insulate information against overwriting by the second frame.

In addition, the amount of time between the cue and the onset of the second frame necessary for the cue to improve performance was shorter for change detection (66 ms) than identification (133 ms). If the attentional process gradually moves information from the iconic trace to a more stable form, this temporal difference may be related to the informational demands of each task. Detecting that an item at a location has changed may require far less information than is required to actually identify the original stimulus.

These results suggest the existence of an iconic trace which is overwritten by new visual information. However, directing attention to a specific item in an iconic trace transfers that item to a short-term-memory buffer. Once an item has been moved to this separate buffer it may be compared with the later-occurring item in that location, leading to accurate change detection. These concepts are all in line with earlier proposals about the nature of iconic memory (Averbach and Coriell 1961; Gegenfurtner and Sperling 1993; Sperling 1960).

It should be pointed out that even in the most ideal situation, with a cue following first-frame offset by 16 ms and with 256 ms between the cue and second frame, the percentage of correct original item identifications is only 48%. This performance seems extremely low. In fact, throughout all three experiments the performance on the identification task was relatively poor. These poor identification rates might be caused by the new item in the change location interfering with the retrieval of the original item in that location. If this were the case, changing an item to a different kind of item (eg a letter changing to a symbol) might lower the amount of interference, much like a change of category results in a release from proactive interference (Wickens et al 1968). To assess this, we looked informally at identification of original items when letters changed to symbols and found that identification performance for these was higher (~ 75%). On the basis of this informal study, it seems reasonable to suspect that interference may be responsible for the extremely poor performance on the identification task; however, more research is needed to confirm that it is responsible for the low identification rates.

6 General discussion

Previous change-blindness experiments have demonstrated that, in the absence of a clearly defined transient, people are very poor at detecting changes in visual scenes. One possible interpretation of this finding is that people represent very little of their visual environment, only representing the few items to which they can attend. An extreme form of this interpretation would suggest that when we view a picture we are essentially blind to all but a few items in the scene. Naturally, this account seems to challenge our subjective experience of a detailed visual world.

Another possible interpretation is that we have a fairly rich representation of a scene while that scene is present, but lack the ability to store more than a tiny bit of this representation in a way that allows it to be compared to subsequent scenes. The change-detection task requires a comparison of two scenes across time. Therefore, even if people do represent much of their visual environment while that scene is present, they may still fail at a change-detection task because they are unable to simultaneously represent and compare two visual scenes. In fact, recent experiments using measures thought to be more sensitive than a verbal report of change, suggest that people encode more information than their verbal responses indicate (Hayhoe et al 1998; Fernandez-Duque and Thornton 2000).

The data from the current experiments are consistent with this latter interpretation. In experiments 1 and 3 change-detection rates improved when a subject's attention was drawn to the location of a change during the blank ISI. This finding implies that the subject does have an internal representation, probably in the form of iconic memory, of much of the scene. However, drawing the subject's attention to the location of the change simultaneously with the appearance of the second scene yielded no benefit (experiments 2 and 3). This finding suggests that the appearance of the new scene overwrites the representation of the first scene. Thus change blindness might reflect an inability to store much information from a scene, rather than a failure to represent much information about a scene as it is being viewed. In sum the inability revealed here seems wholly consistent with early proposals about the workings of iconic memory. In fact, many of the findings which have been described as counterintuitive in the change-detection literature are exactly in accord with the predictions one would make from the iconic-memory literature (Averbach and Coriell 1961; Gegenfurtner and Sperling 1993; Sperling 1960).

While our results verify that there is an internal representation of much of a visual scene and provide a possible explanation of why this representation does not result in good change detection, they say nothing about how much processing this

representation receives prior to attention. It is possible that prior to attentional processing the representation is simply the persistence of retinal neural firing. That is, the representation may not explicitly code the identity or form of the items in the environment. However, it is logically possible that, prior to attention, there is an explicit representation of the identity of many items in the scene (but see Pashler 1984).

Unfortunately the change-detection paradigm may not be able to distinguish between these two possibilities. Change detection always requires both an explicit representation of the information from the original scene and the storage of that information for comparison to a later-occurring scene; thus the method is not ideally suited for assessing the depth of processing of unattended items. Any failure to detect a change could result from either a failure to explicitly represent information from the first display or a failure to store that explicit information for comparison with the later scene. In conclusion, while the change-detection paradigm is inherently interesting, and provides a useful tool for investigating the amount of information people retain from one scene to the next, it is not well suited for investigating the detail and complexity of the representations formed while scenes are actually viewed. Inferring what an observer represents or is aware of while viewing a scene on the basis of explicit verbal change detection is, at the very least, questionable.

Acknowledgements. This work was supported by the National Science Foundation (SBR 9729778) and the National Institute of Mental Health (R01-MH45584-09).

References

- Averbach E, Coriell A S, 1961 "Short term memory in vision" *Bell Systems Technical Journal* **40** 309-328
- Coltheart M, Lea D C, Thompson K, 1974 "In defense of iconic memory" *Quarterly Journal of Experimental Psychology* **26** 633 – 641
- DiLollo V, Dixon P, 1988 "Two forms of persistence in visual information processing" *Journal of Experimental Psychology: Human Perception and Performance* **14** 671 -681
- Fernandez-Duque D, Thornton I M, 2000 "Change detection without awareness: Do explicit reports underestimate the representation of change in the visual system?" *Visual Cognition* **7** 323-344
- Gegenfurtner K R, Sperling G, 1993 "Information transfer in iconic memory experiments" *Journal of Experimental Psychology: Human Perception and Performance* **19** 845-866
- Grimes J, 1996 "On the failure to detect changes in scenes across saccades", in *Perception* Ed. K Akins (Oxford: Oxford University Press) pp 89-110
- Hayhoe M, Bensinger D G, Ballard D H, 1998 "Task constraints in visual working memory" *Vision Research* **38** 125-137
- Irwin D E, Yeomans J M, 1986 "Sensory registration and informational persistence" *Journal of Experimental Psychology: Human Perception and Performance* **12** 343-360
- 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 and Performance* **18** 530 - 549
- Loftus G R, Johnson C A, Shimamura A P, 1985 "How much is an icon worth?" *Journal of Experimental Psychology: Human Perception and Performance* **11** 1-13
- McConkie G W, Currie C B, 1996 "Visual stability across saccades while viewing complex pictures" *Journal of Experimental Psychology: Human Perception and Performance* **22** 563-581
- McConkie G W, Zola D, 1979 "Is visual information integrated across successive fixations in reading?" *Perception & Psychophysics* **25** 221-224
- O'Regan J K, Rensink R A, Clark J J, 1999 "Change-blindness as a result of 'mudsplashes'" *Nature (London)* **398** 34
- Pashler H, 1984 "Evidence against late selection: Stimulus quality effects in previewed displays" *Journal of Experimental Psychology: Human Perception and Performance* **10** 429-448
- Pashler H, 1988 "Familiarity and visual change detection" *Perception & Psychophysics* **44** 369-378
- Phillips W A, Singer W, 1974 "Function and interaction of on and off transients in vision. I. Psychophysics" *Experimental Brain Research* **19** 493-506

- Rensink R A, 1997 "How much of a scene is seen? The role of attention in scene perception" *Investigative Ophthalmology & Visual Science* **38**(4) S707
- 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, 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
- Sperling G, 1960 "The information available in brief visual presentations" *Psychological Monographs: General and Applied* **74** 1-29
- Stelmach L B, Bourassa M, DiLollo V, 1984 "Detection of stimulus change: The hypothetical roles of visual transient responses" *Perception & Psychophysics* **35** 245 - 255
- Wickens D D, Clark S E, Hill F A, Wittlinger R P, 1968 "Investigation of grammatical class as an encoding category in short-term memory" *Journal of Experimental Psychology* **78** 599 - 604