Available online at www.sciencedirect.com
SCIENCE DIRECT
Consciousness and Cognition 15 (2006) 409–422
Consciousness and Cognition
www.elsevier.com/locate/concog

Change blindness and priming: When it does and does not occur

Michael E. Silverman [a,*], Arien Mack [b]

[a] *Mount Sinai School of Medicine, Department of Psychiatry, Box 1230, NY 10029, USA*
[b] *New School University, NY, USA*

Received 3 August 2004
Available online 11 January 2006

Abstract

In a series of three experiments, we explored the nature of implicit representations in change blindness (CB). Using 3 × 3 letter arrays, we asked subjects (Ss) to locate changes in paired arrays separated by 80 ms ISIs, in which one, two or three letters of a row in the second array changed. In one testing version, a tone followed the second array, signaling a row for partial report (PR). In the other version, no PR was required. After Ss reported whether a change had been detected and the PR had been completed (if required), they were asked to identify a degraded letter trigram that was either novel, or from a previously shown row (repetition priming). Our findings indicate that when CB occurs, *both* the pre-change and post-change stimulus information primes despite its unavailability to consciousness. Surprisingly, findings also indicate that when change detection occurs only the post-change information primes.

Keywords: Change blindness; Change detection; Attention; Priming; Implicit memory

1. Introduction

Over the years, research on the phenomenon of change blindness has led to many important insights into visual perception, awareness and memory. More recent research, both in the laboratory (e.g., Rensink, O'Regan, & Clark, 1997) and in the real world (e.g., Simons & Levin, 1998), has more closely examined the link between attention and awareness, demonstrating that conscious perception of change requires focused attention and that, without attention, our representations of the visual world are extremely volatile. Over the past decade, change blindness research has also led to the striking finding that not only are we blind, to a great extent, to large and otherwise obvious changes in the visual world around us, but that we often remain completely unaware of this fact (Levin, Momen, Drivdahl, & Simons, 2000; Simons & Rensink, 2005).

That most observers are blind to significant visual changes might suggest that the original features of the object or objects undergoing change are not sufficiently represented or retained (Dennett, 1991). However, others have argued that this visual information is represented and that change blindness, instead, occurs

* Corresponding author.
E-mail address: michael.silverman@mssm.edu (M.E. Silverman).

1053-8100/$ - see front matter
doi:10.1016/j.concog.2005.08.003

because the memory trace of the original information is fleeting (Wolfe, 1999), overwritten (Simons, 2000), not compared (Simons, Chabris, Schnur, & Levin, 2002), or is retained in some non-accessible or non-reportable form (Mack & Rock, 1998). In any case, most researchers agree that changes are reportable only when observers attend to the change in a visual stimulus, encode that change in memory, and then compare the old (pre-change) information to the new (Becker, Pashler, & Anstis, 2000). In short, focused attention is needed for change detection, although attention alone does not necessarily guarantee it.

While a good deal of research on change blindness has primarily focused on the failure to detect visual changes (Grimes, 1996; Rensink, 2000; Rensink et al., 1997; Simons & Levin, 1998) and its implications, other research has attempted to determine what information, if any, is *preserved* in the face of change blindness. Some research has demonstrated that although visual perception may require stable representations, the failure to consciously perceive a visible stimulus, does not necessarily imply the absence of deep processing and encoding (Mack & Rock, 1998). In addition, accumulating evidence for unconscious or *implicit* perception (see, Bornstein & Pittman, 1992; Draine & Greenwald, 1998; Greenwald & Draine, 1998; MacLeod, 1998; Marcel, 1983; Moore & Egeth, 1997) now suggests that perceptual processing does occur outside of awareness without the allocation of attention (see Maki, Frigen, & Paulson, 1997 for a review of when priming does and does not occur in the attentional blink, 1997; Shapiro, Driver, Ward, & Sorenson, 1997).

This evidence, we believe, suggests that observers may, in fact, represent visual information even in the face of change blindness.[1] With this idea in mind, we set out to exploit (employing a similar approach to that used by Becker et al., 2000) the phenomenon of change blindness to determine the fate of visual representations in change blindness displays and how they may be affected by change detection.

In a series of three experiments, we use a perceptual identification task to assess what visual information is preserved despite its unavailability to conscious report. This first experiment determines the fate of the unreported information in an *iconic image* using a Sperling (1960) type 3×3 letter matrix. Experiments 2 and 3 modify this procedure to address the role of attention in visual representations in a change blindness task.

Together, these experiments address: whether incomplete subject reports are the result of the rapid decay of representations or limited conscious access to them (i.e., whether visual representations persist even after conscious access is no longer possible) and under what conditions the pre- and post-change information of a changing visual stimulus, which is not *selectively*[2] attended to, is nevertheless represented and encoded.

2. Experiment 1

Although our primary goal was to determine the fate of stimuli in change blindness displays, we first needed to determine whether information from a brief visual presentation was preserved without the benefit of selective attention and whether this information would prime a subsequent response despite its unavailability to consciousness. To do this, we chose to employ the now classic procedure developed by Sperling (1960) to determine whether the non-cued (unselected) rows in the partial report task truly decay like a fast-fading after-image, restrained from further processing by some kind of informational bottleneck as originally argued[3], or whether all the items in the stimulus displays, particularly those which had not been selectively attended to, are stored for a significantly longer duration despite their unavailability to consciousness. Our assumption is that the cued row is selectively attended to, whereas the other rows in the matrix are objects of distributed attention.

[1] Recently a few studies have attempted to demonstrate that the visual system maintains representations of *changes* (as opposed to *information*) that cannot be overtly reported (See Fernandez-Duque & Thornton, 2000, 2003; Smilek, Eastwood, & Merikle, 2000). Although, these findings can be better attributed to methodological artifacts (Mitroff, Simons, & Franconeri, 2002), they provide the grounds for devising a procedure to explore whether more information from a change blindness array is available then suggested by explicit reports.

[2] It is important to note that we are assuming that all the rows of the letter matrices are the object of at least distributed attention whereas the letters in the cued row are likely to be the only ones that are the object of selective or focused attention.

[3] As noted by Haber (1983), theories of iconic storage by definition suggest a visual system in which a rapidly decaying retinal image is treated as a picture to be looked at and presumably read. These "select-then-identify" theories imply a perceptual system limited by an early attentional "bottle-neck" in which *only* those items selectively attended to from the rapidly fading image will proceed to some deeper level of processing and presumably to some longer-term store.

As in Sperling's original partial report procedure (1960), subjects were briefly shown a 3 × 3 consonant letter matrix. After the matrix was removed, a cue was sounded (a high, medium, or low tone) signaling which row to report. Once this partial report was made, subjects then tried to identify a degraded three letter trigram that was either novel or was one of the rows of the previously shown matrix. Decoding of these degraded trigrams served as our measure of repetition priming and implicit perceptual learning (for a review, see Bornstein & Pittman, 1992).

In the Sperling iconic memory paradigm, increasing the delay between the presentation of the matrix and the partial report cue results in a decline in the subjects' ability to report the cued row (Sperling, 1960). However, if priming were to occur and turns out to be unaffected by the delay, this would provide further proof of the encoding and storage of all the matrix letter rows regardless of whether or not they had been the recipient of selective attention and despite their increased unavailability to consciousness.

3. Method

3.1. Subjects

Seven subjects (3 male, 4 female) recruited from the New School University, received $25 for participating. All were native English speakers with normal vision. Subjects were not informed of the experimental hypothesis.

3.2. Apparatus

The experiment was run using PsyScope (v.1.2.5; Cohen, MacWhinney, Flatt, & Provost, 1993) on an Apple Computer with a 14-inch monitor (640 × 480, 75 Hz) positioned 76 cm from the subject.

3.3. Stimuli

Displays consisted of nine randomly selected letters from the 21 consonants including Y. Letters were presented in black 24 pt. Chicago font against a white background, each subtending an angle of .6° by .45°. The letters were arranged in a 3 × 3 grid subtending an angle of 2.6° by 2.2° in the center of the viewing screen. All presentation luminances (including pre- and post-exposure fields), as well as stimuli dimensions, were determined by measuring a uniform rectangle with a Candella Photometer calibrated against a standard light source. The measured luminance of the pre- and post-exposure fields was 21.0 ft-c. Degraded trigrams, used as the measure of priming, were created by randomly removing 75% of the pixels from each of the 3 stimulus letters. This level of degradation was chosen after performing earlier pilot studies to determine chance levels of decoding degraded trigrams. A degradation of 0.75 yielded an error rate of approximately 70%. This severe letter degradation rate was chosen to avoid the possibility that a subject's ability to decode the degraded trigrams might improve with practice and, thus, potentially mask any priming effect.

3.4. Procedure

The initial part of each trial involved a standard partial report task (Sperling, 1960). The experiment had two conditions: *immediate* cueing between the offset of the matrix and the cueing tone; and *delayed* cueing in which the cueing tone sounded 500 ms after the offset of the matrix. The order of the conditions was counterbalanced across subjects. Each subject received a total of 320 trials with 40 trials during each of eight separate daily sessions. Each trial began with the presentation of a fixation mark centered on the display. After subjects pressed a key to start the trial, a pre-exposure field appeared for 500 ms, followed by the stimulus matrix for 50 ms, and then a post-exposure field. The cueing tones were highly distinguishable (1000, 500, and 200 Hz) and signaled the observer to report the respective row of the matrix: a high tone signaled the report of the top row; a middle tone signaled report of the middle row; and a low tone signaled report of the bottom row. After hearing the cueing tone, subjects first reported the letters verbally and then typed them using the keypad. When these reports conflicted, the initial verbal report was used.

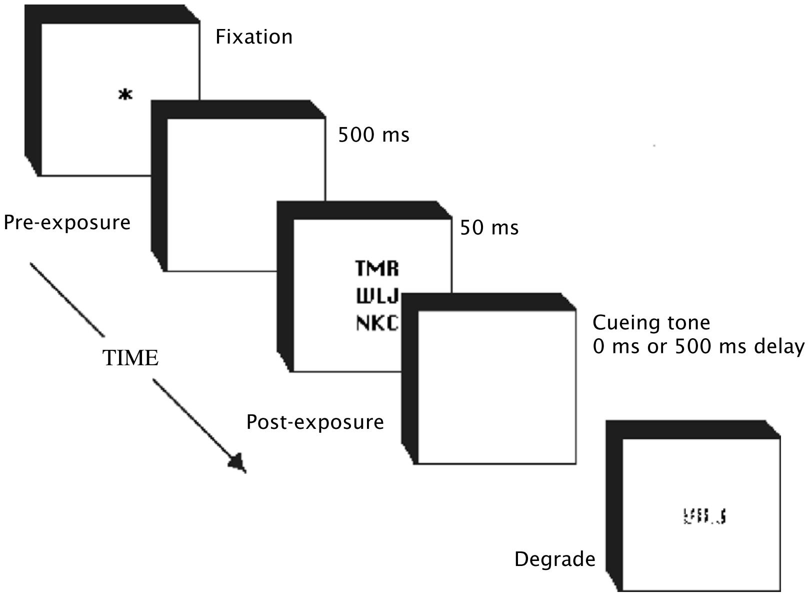

Fig. 1. Stimulus presentation for Experiment 1.

After the partial report component of each trial, subjects were asked to identify a degraded trigram chosen either from the cued or non-cued rows of the matrix or composed of novel letters (see Fig. 1). On average, ten seconds elapsed between the matrix offset and the onset of this degrade.

3.5. Partial report coding

To estimate the number of items available to conscious report from the briefly presented letter matrix, the number of correctly reported letters was multiplied by 3. Given that subjects could not predict which row would be cued, a perfect score of 3 would indicate that all nine items in the matrix had been perceived and could have been reported if cued (Sperling, 1960). *Degraded Trigram Identification:* Degrades always appeared in the center of the screen. A response was considered correct only if *all* the letters were correctly identified and reported in their correct serial positions.

4. Results

In the immediate cuing condition, subjects correctly reported an average of 1.99 of the three letters from the cued row suggesting that they potentially could have reported, on average, 5.97 of the nine matrix letters ($SD = 1.28$). Consistent with the earlier Sperling (1960) findings, subjects performed less well in the delay condition, reporting on average only 1.53 letters from the cued row and thereby potentially recalling only 4.60 letters on average ($SD = 1.06$; $t(6) = 6.09$, $p < .01$). Thus as expected, delaying the cueing tone had a negative effect on the ability to correctly report letters in the partial report task.

There was a significant difference in decoding of degraded letters by condition (cued, non-cued and novel trials; $F(2,12) = 9.36$, $p < .004$) but this was entirely due to the difference between the decoding of the novel trigrams and the decoding of the cued and uncued trigrams. There was *no* significant difference in the decoding of degraded letters from the cued and non-cued rows ($t(6) = 1.09$, $p > .31$), while novel trigrams were correctly decoded significantly less often than either cued ($t(6) = 3.08$, $p < .03$) or non-cued rows ($t(6) = 3.55$, $p < .02$). This was also the case in the delayed cueing condition. Again there was no difference in the decoding of the non-cued and the cued rows ($t(6) = .203$, $p > .84$), and again both were successfully decoded significantly more often than the novel trigrams (cued, $t(6) = 4.61$, $p < .01$; non-cued, $t(6) = 2.83$, $p < .03$). Interestingly, there was *no* main effect of delay on decoding ($F(2,12) = .266$, $p > .77$). In other words, even though subjects could

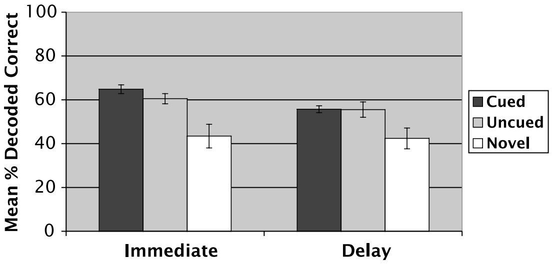

Fig. 2. Mean percent correct decoding from Experiment 1. Error bars represent standard deviation.

recall fewer letters from the matrix when the cue to recall occurred 500 ms after the matrix had disappeared, the delay had no impact on the ability to decode the degraded matrix items. That is, there was no difference between the immediate and delay conditions in the decoding of the cued and uncued trigrams (see Fig. 2).

5. Discussion

Iconic images are generally considered ephemeral. Only those parts of the stimulus to which selective attention is directed seem to survive in some longer-term memory (Sperling, 1960). Experiment 1 probed whether this information in fact decays, prevented from further processing by some kind of "informational bottleneck" or, as is suggested by an increasing body of literature, persists despite its unavailability to consciousness. We found that it does and based on these results, we now have a way of studying the fate of information in a change blindness paradigm—that is, the fate of the visual information from the pre- and post-change displays, and what effect change blindness and change detection may have on this information, and more specifically on the capacity of this information to prime subsequent responses.

6. Experiment 2a

Experiment 1 demonstrated *implicit* encoding of matrix items even though they could not be consciously recalled. Experiment 2a adopts this procedure to explore the nature of change blindness and the role of selective attention in change detection. Specifically, this experiment assesses whether or not *pre-change* information from an undetected change is preserved in implicit memory, as revealed though the same priming procedure now preceded by a change blindness task. If this information is preserved, a previously viewed, now degraded trigram should be decoded better than a novel trigram, suggesting that the failure of change detection does not necessarily imply the absence of representation of the pre-change stimulus information. This procedure also allows us to explore the fate of letter matrix information when changes are correctly reported.

7. Method

7.1. Subjects

Ten subjects (6 male, 4 female) recruited from the New School University received $25 for participating. All were native English speakers with normal vision. They were not informed of the experimental hypothesis. *Procedure:* The procedure from Experiment 1 was modified so that changes could be introduced into the display. Each trial began with the presentation of a centered fixation mark. To initiate each trial, subjects depressed the keyboard spacebar. The pre-exposure field appeared for 500 ms, followed by the pre-change stimulus matrix for 100 ms, then by a gray post-exposure field for 80 ms, and finally by the changed matrix for 100 ms. Changes, involving the replacement of *all three letters* of a *single row,* either top, middle or bottom, were counterbalanced and occurred on every trial. Immediately following the offset of the second (changed) matrix, the partial report cue was presented indicating the row to report. (see Fig. 3)

After hearing the cueing tone, subjects reported whether any change was perceived and, if so, at which location (i.e., top, middle, or bottom row). They then reported the letters from the cued row (as in Experiment 1)

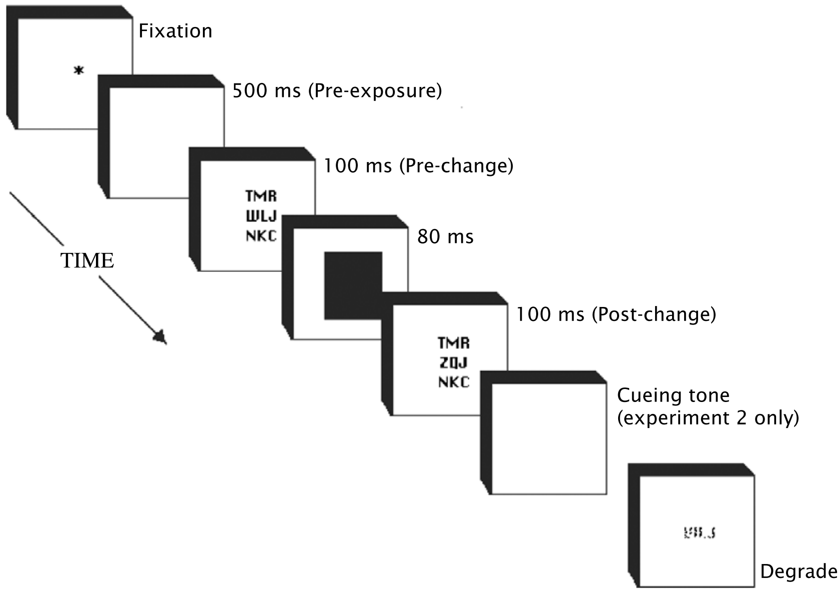

Fig. 3. Stimulus presentation for Experiments 2 and 3.

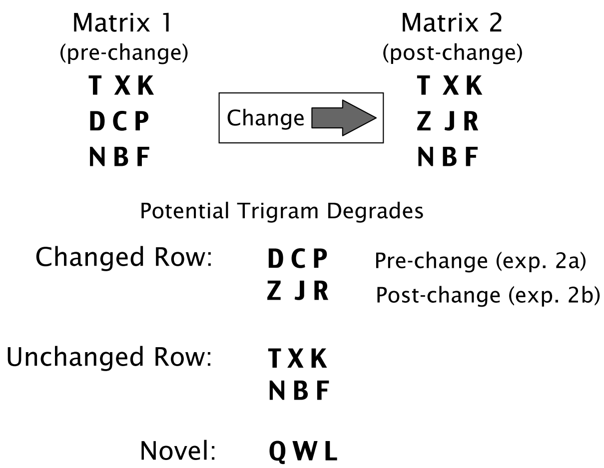

Fig. 4. Potential degrades for Experiments 2a and 2b.

of the second (changed) matrix. After these responses were made, a degrade of the changed, unchanged or novel trigram appeared on the screen, which subjects then attempted to decode. Trials involving changed row degrades used only the *pre-change* row of letters from the first matrix. (see Fig. 4)

Subjects completed 10 practice trials and a total of 324 counterbalanced trials during one testing session. Although changes occurred on every trial, subjects were told that changes might or might not occur on any given trial. They were allowed to rest as often as necessary.

7.2. Coding

The procedures for scoring the partial report and degrade decoding were identical to those in Experiment 1. Accuracy of change detection was based on a correct or incorrect response. If the subject was unable to report the change or reported a change in a row that had not changed, it was recorded as a miss.

8. Results

Subjects recalled an average of 3.99 of the nine letters ($SD = 1.38$) from the second (changed) matrix based on their partial report score, which was comparable to the delayed-cue, partial report score in Experiment 1 (This was probably because the majority of the subjects responded to the change task first, thus introducing a delay roughly equivalent to the one in the delayed cueing condition of the first Experiment). Mean change detection performance was 47.6% ($SD = 12.15$). When the changed row was also the row cued for a partial report and the change was detected, the mean calculated partial report was 6.30 compared to 3.20 letters correct when the change went undetected. When changes were not detected, evidence of significant differences among decoding of the trial-type degrades (changed, unchanged and novel trigrams) was observed ($F(2,18) = 42.31$, $p \leqslant .001$). As in Experiment 1, further analysis revealed that there was no statistical difference between the decoding of changed and unchanged trigrams ($t(9) = .517$, $p \geqslant .617$). But again the degraded novel trigrams were decoded significantly less often than trigrams that had previously appeared in the matrix (changed, $t(9) = 7.31$, $p \leqslant .001$. and unchanged, $t(9) = 7.51$, $p \leqslant .001$). Unexpectedly when the change was detected, the differences between the decoding of the changed, unchanged and novel trigrams disappeared ($F(2,18) = 2.01$, $p = .196$). (see Fig. 5)

9. Discussion

Experiment 2a explored the fate of the information in the pre-change array by contrasting the differences in response accuracy (i.e., decoding) between the changed, unchanged and novel degraded trigrams *when a change is undetected.* The obtained results are consistent with the hypothesis that the stimulus information in arrays in which a change goes undetected persists in a form adequate for lexical, repetition priming—without the benefit of selective attention which seems necessary for change detection.

While these results are no longer completely surprising, given our results from Experiment 1 and the extensive literature on implicit priming (see Bornstein & Pittman, 1992; for review), what was unforeseen was what happened when observers did detect the change. When the change was detected, none of the previously displayed trigrams from the first arrays seemed capable of priming. In other words when the change was detected the difference between the decoding of the changed, unchanged and novel degrades disappeared. Previously seen rows that had appeared in the first matrix were decoded *no better* than the novel trigrams, even though two of these trigrams (the ones from the two unchanged rows) had appeared twice, once in the first and once in the second matrix. This main effect of change detection evidenced in the failure of priming may be directly linked to task-related interference: i.e., the perceptual system might be actively inhibiting the trigrams that were irrelevant to change detection.

10. Experiment 2b

Experiment 2a demonstrated a main effect for change detection, with subjects failing to decode previously shown trigrams from the first matrix any better than novel trigrams—in direct contrast to the results when the change was not detected. Since subjects' partial reports had been more accurate (Experiment 2a) for the changed trigram from the second matrix (or the post-change row) when the changed row was cued *and* the change

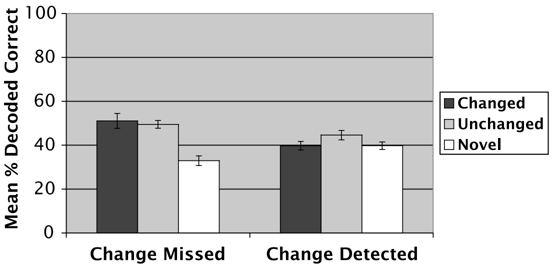

Fig. 5. Mean percent correct decoding results from Experiment 2a.

was detected, they should, we posited, be able to decode the post-change row degrade better than all other degrades. If so, the lack of priming results in Experiment 2a could, we reasoned, be attributed to interference caused by change detection. We explored this in Experiment 2b by asking subjects to decode degraded trigrams from the *post-change* rather than the *pre-change* matrix.

11. Method

11.1. Subjects

Five subjects (1 male, 4 female) recruited from the New School University, received $25 for participating. All were native English speakers with normal vision. They were not informed of the experimental hypothesis.

11.2. Procedure

Procedures were identical to those in Experiment 2a except that the changed trigram degrades were now taken from the *post-change* matrix rather than the *pre-change* matrix. Since two of the trigrams were the same in both matrices, this, of course, meant that the only difference between Experiments 2a and 2b was in presenting the changed trigram as a degraded stimulus to be decoded.

12. Results

As in Experiment 2a, subjects recalled an average of 4.20 of the 9.00 letters from the matrix (SD;= 1.06) and detected 39.2% of the changes ($SD = 19.16$). When the changed row was cued *and* the change was detected, subjects, upon being asked to report the cued row, were again much more likely to report the cued row correctly (6.5), than when the change was not detected (3.7) (see Fig. 6).

Similar to Experiment 2a, there were significant differences between the decoding of changed rows, unchanged rows and novel degrades when the change was not detected ($F(2,8) = 23.74$, $p \leqslant .015$). However, again there were no differences between the decoding of changed and unchanged row degrades ($t(4) = .828$, $p \geqslant .454$), while degraded novel trigrams were decoded significantly less often than trigrams that had previously appeared in the matrix (changed, $t(4) = 6.23$, $p \leqslant .004$. and unchanged, $t(4) = 6.20$, $p \leqslant .004$). These results, consistent with our earlier findings from Experiment 2a, indicate that visual stimuli in change blindness displays are capable of priming in the absence of change detection.

Again as in Experiment 2a demonstrated that when the change was detected, no *pre-change* matrix row trigrams produced priming. As Fig. 6 shows, decoding performance of the changed row in the post-change matrix was significantly better than the unchanged ($t(4) = 2.78$, $p \leqslant .050$) and novel rows ($t(4) = 3.83$, $p \leqslant .018$). However, there was again no difference between the decoding of the novel trigram degrades and the unchanged row degrades ($t(4) = .092$, $p \leqslant .939$).

13. Discussion

Perhaps our most surprising result was that when the change was detected, neither of the unchanged trigram rows primed the degraded identification task, even though these rows had appeared twice and been

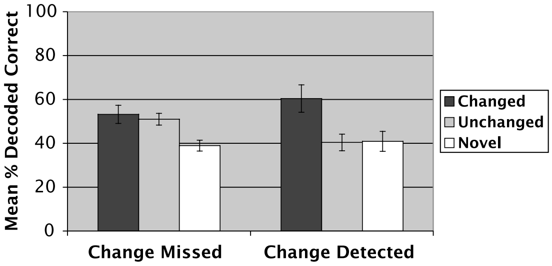

Fig. 6. Mean percent correct decoding results from Experiment 2b.

viewed for a total of 200 ms. Why does the detection of a change eliminate priming for unchanged stimuli? One possible explanation is that the simultaneous tasks of change detection *and* partial report together caused interference, thus decreasing the likelihood of perceptual priming. This would be consistent with the fact that manipulations of attention involving response selection can affect perceptual memory by disrupting encoding (for review see Lavie, 2005). However, an alternative explanation might be that when a change is detected, the awareness of the change may weaken or repress all the other stimulus information that otherwise would be evidenced by priming (Jiang & Chun, 2001). Were this the case, this could suggest that when change is detected the perceptual system attempts to actively insulate the attended to information from both proactive and retroactive interferences so as to better preserve it.

14. Experiment 3

Experiments 2a and b demonstrated a main effect of change detection, which, as we have already suggested, could be due either to a dual-task interference effect or due to inhibition that accompanies and promotes change detection. To test for dual-task interference, we removed the partial report task from Experiments 3a and b. However, we suspected that in its absence, if changes continued to involve all three letters in a matrix row as it had in all the previous Experiments, subjects would quickly learn to attend to only one letter in each row, thereby detecting changes with minimal effort. To eliminate this as a possible strategy, we therefore varied the number of letters (one, two or three) that changed in the given row from one trial to the next. In Experiment 3a, we examine the *pre-change* information and, in Experiment 3b, the *post-change* information.

15. Method

15.1. Subjects

Twenty subjects (7 male, 13 female) recruited from Harvard University, received $10 for participating. All were English speakers with normal vision. 10 subjects participated in each Experiment. They were not informed of the experimental hypothesis.

15.2. Stimuli

Stimulus objects and dimensions were identical to those in all previous Experiments. However, the experiment and trigram degradation was now controlled using custom software written with Vision Shell C libraries. Instead of changing an entire row, changes were now to one, two or all three letters in a particular row, with all combinations and locations counterbalanced. (see Fig. 7)

15.3. Procedure

The procedures were the same as in Experiments 2a and b except that the partial report cue and task were eliminated, and subjects now entered their response to the degraded stimuli using the computer keypad.

15.4. Coding

Measures for change detection and degrade decoding were identical to those used in earlier Experiments.

16. Results (Experiment 3a)

Comparable to Experiments 2a and b, mean change detection was 38.9% ($SD = 11.5$). The mean percentage of change detection based on the number of letters changed in the changed row were $1 = 24.7\%$ ($SD = 12.8$); $2 = 39.9\%$ ($SD = 14.2$) and $3 = 47.7\ \%$ ($SD = 14.0$), demonstrating that the likelihood of detection increased with the number of changes to a row ($F(2,18) = 23.19,\ p \leqslant .001$).

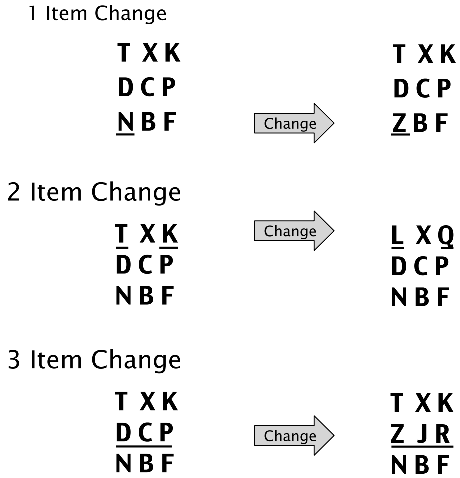

Fig. 7. Example of the potential change types (1, 2 or all stimuli of a given row) in Experiments 3a and 3b.

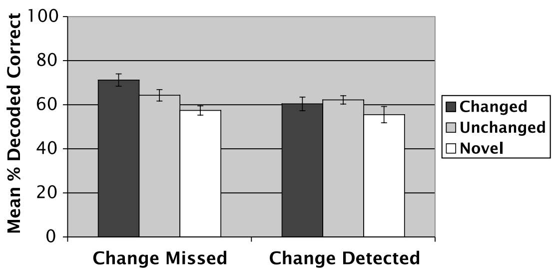

Fig. 8. Mean percent correct decoding results from Experiment 3a.

Significant differences were found in decoding accuracy between the changed row, unchanged row and novel trigrams when the change was not detected ($F(2,18) = 8.15$, $p \leqslant .003$). Once again, although there was no statistical difference between the decoding of the degraded trigrams from the rows that had changed and those that had not ($t(9) = 1.56$, $p \geqslant .152$), the degraded novel trigrams were decoded significantly less often than those that had appeared in the pre-change matrix (changed, $t(9) = 4.17$, $p \leqslant .002$ and unchanged $t(9) = 3.25$, $p \leqslant .010$) (see Fig. 8).

Analysis of the results indicated that when change was detected, there again were *no* differences in decoding between the changed row, unchanged row and novel trigram degrades ($F(2,18) = 2.13$, $p = .148$).

17. Discussion

These results are consistent with our earlier finding that, when change goes undetected, the pre-change information remains sufficiently represented by the visual system to prime, despite its unavailability to consciousness. However, when the change is detected, degraded trigrams from the first matrix are *not* decoded any better than degraded novel trigrams. This finding argues against an "informational bottleneck" and, instead, supports an inhibition account. Thus, in the instance when attention is correctly directed to the changed row *and* the change is detected, post-change information seems to suppress or inhibit all other information that, in its absence, would be otherwise represented sufficiently to prime.

18. Results (Experiment 3b)

Mean change detection across all trial types was 44% ($SD = 9.29$), comparable to Experiments 2a, 2b, and 3a. The mean percent of change detection based on the number of letters in the changed row were $1 = 30.6\%$

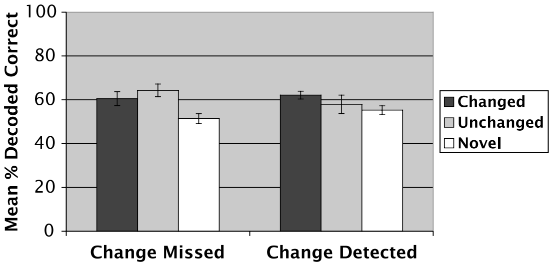

Fig. 9. Mean percent correct decoding results from Experiment 3b.

($SD = 7.4$); $2 = 47.1\%$ ($SD = 11.3$) and $3 = 54.9\%$ ($SD = 11.8$), again demonstrating that change detection was more likely as the number of changed letters in the row increased ($F(2, 18) = 49.52$, $p \leqslant .001$).

A review of the differences in decoding accuracy between the changed row, unchanged row and novel trigrams, when the change was not detected, again provided evidence that there were differences among the three trial types ($F(2, 18) = 7.36$, $p \leqslant .005$). Once again, no statistical difference between the decoding of the changed and unchanged trigram degrades was found ($t(9) = .927$, $p \geqslant .378$). However, degraded novel trigrams were decoded significantly less often than those rows that had previously appeared in the matrix (changed; $t(9) = 2.76$, $p \leqslant .022$. and unchanged; $t(9) = 4.56$, $p \leqslant .001$). The review of those trials in which the change was detected again indicated no significant differences among decoding of the three trigram types ($F(2,18) = 2.39$, $p \leqslant .120$). However, as a result of our findings from Experiment 2b, we hypothesized that, *when the change is detected*: (1) only the changed trigram from the post-change matrix should prime; and (2) change detection would result in an inhibition of all of the other *unattended* matrix information. In Experiment 3b, this hypothesis was confirmed - when change was detected, degraded novel trigrams were decoded significantly less frequently than those trigrams that had appeared *and* changed ($t(9) = 3.57$, $p \leqslant .006$) (see Fig. 9).

19. Discussion

These results are consistent with our earlier finding that when change is not detected, all the stimuli in the trigram arrays are capable of priming even though they have not benefited from selective attention.

However, consistent with our results from Experiment 2b, when the change was detected only the changed trigram produced priming. These results, therefore, indicate that: (1) when the change is detected, the only trigram that primes is the changed trigram from the post-change matrix which means that when change is detected, it is accompanied by an inhibition of all other *irrelevant* matrix information. This suggests that the unattended visual information may be volatile and that without focused attention it may be effectively inhibited by the stimulus information that has been attended to (because it is critical to the task at hand), and is, therefore, more stable.

20. General discussion/Conclusions

Despite the common folk belief in, and the "subjective impression" (Dennett, 1991) of, a coherent and richly detailed world, recent empirical research has led many to believe that we actually perceive much less than we think we do. Much of this research points to the fact that, in the absence of focused attention, our visual representations of the world are quite ephemeral. However, not everyone agrees that change blindness absolutely necessitates sparse or absent representations. That is, change blindness, much like inattentional blindness, may occur with relatively detailed and possibly even complete representations (see Simons & Rensink, 2005 for commentary).

In general, we believe that our results are consistent with, and support the hypothesis that, stimulus information from brief visual presentations exist, and *persist,* even in the face of change blindness and the absence of selective attention. The results from Experiment 1, we believe, strongly suggest that *briefly presented* visual stimuli that do not benefit from *selective* attention are nevertheless encoded and that, despite their volatility, visual representations persist even though they are not consciously available. This finding might not be so surprising in 2005, but certainly would have been in 1960 when Sperling had done the original work. Oddly his conclusion that the uncued portion of the iconic image quickly decayed to our knowledge has never been questioned despite the accumulating evidence since the 1980s (e.g., Marcel, 1983; Schacter, 1987) demonstrating that unattended stimuli are capable of priming. In addition, we believe that the clear evidence of priming in the face of change blindness (Experiments 2 and 3), demonstrates that the failure of change detection does not necessarily imply the absence of *multiple* visual representations (pre- and post-change information[4]).

The accumulating evidence, our own and that of other researchers (e.g., Greenwald & Draine, 1998; MacLeod, 1998; Mack & Rock, 1998; Marcel, 1983; Schacter, 1987), indicating the encoding and enduring representations of stimulus input of which we are not consciously aware or cannot consciously recall, led us first to question the received view that iconic images and iconic memories were transient and only endured if they benefited from selective attention. Having found evidence of the encoding of this information (Experiment 1) we were then able to explore the fate of the stimulus information in change blindness displays, both when change was and was not detected. While it is clear that even if all the information in change blindness displays were encoded and remained available, change blindness might nevertheless occur due to the failure of the necessary comparison process (Angelone, Levin, & Simons, 2003; Mitroff, Simons, & Levin, 2004), it seemed important to find out whether or not this input actually does or does not get encoded (e.g., Beck & Levin, 2003; Noë, Pessoa, & Thompson, 2000; O'Regan & Noë, 2002). If it does, it would provide additional evidence that the primary cause of change blindness resides in the failure of the comparison process. The subsequent experiments on change blindness (Experiments 2 and 3) provide important evidence about the fate of the stimuli in change blindness arrays.

At the very least, the results from Experiments 2 and 3 confirm earlier results indicating that stimulus information in change blindness arrays may be preserved despite the failure of change detection (Mitroff et al., 2004). In addition these results may also indicate that: (1) encoded representations which are not accessible to consciousness are also not available to the comparison process necessary for change detection (Rensink, 2002) or, (2) that the comparison process fails for some yet unknown other reason. The second alternative seems more likely to us than the first since Mitroff et al. (2004) have shown that even when the information from change blindness arrays is preserved, and accessible to recall, change detection nevertheless may fail. This being the case, it seems likely that change detection failed to occur in our Experiments because the comparison process did not operate successfully for reasons yet to be determined. It also means, as has been stated by others, that the absence of change detection does not necessarily imply an absence of encoded representations of the input.

Perhaps the more interesting question however (because not the subject of prior speculation), relates to the consistent and novel findings from Experiments 2 and 3 when change is detected. Whenever attention was successfully directed to the changed row *and* the change was, therefore, explicitly detected, we found that the *post-change* information seemingly interfered with the representations of all other extraneous information. In other words, the only information robust enough to prime a degraded stimulus was the *post-change* information.

What do these findings tell us about change detection? A possible, but we think incorrect answer, is that when change detection occurs the stimuli from the pre- change display are overwritten by the post-change display and are thereby erased when attention is captured by the change. Overwriting has been offered as one possible explanation of change *blindness* (e.g., Beck & Levin, 2003; Simons, 2000) and of metacontrast masking (see Enns & Di Lollo, 2000) but it does not seem to be a viable explanation for the failure of priming we observed when change is detected. First, if the changed trigram were to overwrite the comparable trigram from the initial array, change ought not to be perceived since the basis for detecting the change, namely the detecting of the difference between the pre- and post trigram in a particular row of the matrix ought not to be possible—in fact this is why overwriting has been offered as an explanation of change blindness. Moreover, even if

[4] Subjects were interviewed at the close of each of the Experiments. Not one of them reported any knowledge of a relationship between the letters appearing in any of the matrices and the subsequently presented degraded trigrams (Experiments 1, 2, and 3). Notably, when subjects did comment upon a possible relationship between the two tasks, they overwhelmingly indicated that the decoding of the degraded stimuli was, instead, a precursor, or in some was connected with, the upcoming matrix or matrices.

the changed trigram were to overwrite the comparable pre- change trigram, the other two unchanged trigrams which appear in both the pre- and post-change arrays are the same, so these trigrams ought still be represented and therefore capable of priming which they appear not to be.

If not overwriting, what then might account for this failure of any but the changed trigram from the post-change matrix to prime when change is detected? While we do not know the answer, it seems possible that when change is detected, the visual system, in an effort to protect the changed input and maximize change detection, inhibits all the other stimulus information in the arrays or attenuates it so effectively that it is unable to compete with the changed stimulus representation and consequently incapable of priming. This possible explanation has the virtue of being consistent with other accounts of selective attention, particularly that proposed by Treisman (1960) to account for the results of dichotic listening experiments. While the notion that an attentional filter attenuates input irrelevant to the subject's task in a dichotic listening experiment might or might not be the correct account of this body of data,[5] this type of explanation is also consistent with what has been referred to as a *coherence theory*, which presumes that attention acts by locking on to certain representations, while disrupting or replacing others (Rensink, 2000). However, further work is needed, for example, to determine whether visual information *must* in some way be ''protected'' or whether current testing methods are simply not yet sensitive enough to probe/elicit such memories. Future studies should therefore be directed towards attempting to determine the neural activity underlying these behaviors and how selective attentional and implicit information interact (for a review of the current theoretical issues see Hannula, Simons, & Cohen, 2005).

In conclusion, we believe our findings provide further insight into the nature of *change perception* and the role of attention in it. Our study also raises new and important questions regarding the nature of multiple representations, though one, in particular, now intrigues us: namely, if the visual system does retain multiple representations, without the benefit of focused attention, what purpose might these consciously inaccessible representations serve?

Acknowledgments

Thanks to Stephen Mitroff for help in programming the presentation and analysis software. Thanks to Daniel J. Simons for comments on earlier versions of the manuscript and helpful discussion.

References

Angelone, B. L., Levin, D. T., & Simons, D. J. (2003). The roles of representation and comparison failures in change blindness. *Perception, 32*, 947–962.

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

Becker, M. W., Pashler, H., & Anstis, S. (2000). The role of iconic memory in change detection. *Perception, 29*, 273–286.

Bornstein, R. F., & Pittman, T. S. (1992). *Perception without awareness: Cognitive, clinical, and social perspectives*. New York: Guilford Press.

Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments and Computers, 25*(2), 257–271.

Dennett, D. C. (1991). *Consciousness explained*. Boston: Little Brown.

Draine, S. C., & Greenwald, A. G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology: General, 127*, 286–303.

Enns, J. T., & Di Lollo, V. (2000). What's new in visual masking? *Trends in Cognitive Sciences, 4*, 345–352.

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.

Fernandez-Duque, D., & Thornton, I. M. (2003). Explicit mechanisms do not account for implicit localization and identification of change. *Journal of Experimental Psychology: Human Perception and Performance, 29*(5), 846–858.

[5] The question of whether attentional selection occurs early or late in the processing stream or, depending on perceptual load (Yi, Woodman, Widders, Marios, & Chun, 2004), occurs either early or late is still a debated question. This is not the appropriate venue for rehearsing the early versus late account of selective attention but interested readers are referred to (Pashler, 1998) for discussion of this issue.

Greenwald, A. G., & Draine, S. C. (1998). Distinguishing unconscious from conscious cognition—Reasonable assumptions and replicable findings: Reply to Merikle and Reingold (1998) and Dosher (1998). *Journal of Experimental Psychology: General, 127*, 320–324.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.). *Perception: Vancouver studies in cognitive science* (vol. 2, pp. 89–110). New York: Oxford University Press.

Haber, R. N. (1983). The impending demise of the icon: A critique of the concept of iconic storage in visual information processing. *Behavioral and Brain Sciences, 6*, 1–54.

Hannula, D., Simons, D. J., & Cohen, N. (2005). Imaging implicit perception: Promise and pitfalls. *Nature Reviews Neuroscience, 6*, 247–255.

Jiang, Y. H., & Chun, M. M. (2001). Selective attention modulates implicit learning. *Quarterly Journal of Experimental Psychology, 54*, 1105–1124.

Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Science, 9*(2), 75–82.

Levin, D. T., Momen, N., Drivdahl, S. B., & Simons, D. J. (2000). Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition, 7*, 397–412.

MacLeod, C. (1998). Implicit perception: Perceptual processing without awareness. In K. Kirsner & C. Speelman (Eds.), *Implicit and explicit mental processes* (pp. 57–78). Mahwah, NJ, USA: Lawrence Erlbaum Associates Inc. Publishers.

Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: MIT Press.

Maki, W. S., Frigen, K., & Paulson, K. (1997). Associative priming by targets and distractors during rapid serial visual presentation: Does word meaning survive the attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 23*, 1014–1034.

Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15*, 197–237.

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*(4), 798–815.

Mitroff, S. R., Simons, D. J., & Levin, D. T. (2004). Nothing compares 2 views: Change blindness can occur despite preserved access to the changed information. *Perception and Psychophysics, 66*, 1268–1281.

Moore, C. M., & Egeth, H. (1997). Perception without attention: Evidence of grouping under conditions of inattention. *Journal of Experimental Psychology: Human Perception and Performance, 23*, 339–352.

Noë, A., Pessoa, L., & Thompson, E. (2000). Beyond the grand illusion: What change blindness really teaches us about vision. *Visual Cognition, 7*, 93–106.

O'Regan, J. K., & Noë, A. (2002). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24*, 939–1031.

Pashler, H. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.

Rensink, R. A. (2000). The dynamic representation of scenes. *Visual Cognition, 7*, 17–42.

Rensink, R. A. (2002). Change detection. *Annual Review of Psychology, 53*, 245–277.

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

Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 501–518.

Shapiro, K., Driver, J., Ward, R., & Sorenson, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science, 8*, 95–100.

Simons, D. J., Chabris, C. F., Schnur, T. T., & Levin, D. T. (2002). Evidence for preserved representations in change blindness. *Consciousness and Cognition, 11*, 78–97.

Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition, 7*, 1–15.

Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during real-world interaction. *Psychonomic Bulletin and Review, 4*(5), 644–649.

Simons, D. J., & Rensink, R. A. (2005). Change blindness: past, present, and Future. *Trends in Cognitive Sciences, 9*(1), 16–20.

Smilek, D., Eastwood, J. D., & Merikle, P. M. (2000). Does unattended information facilitate change detection. *Journal of Experimental Psychology; Human Perception and Performance, 26*(2), 480–487.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied, 74*, 1–29.

Treisman, A. (1960). Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology, 12*, 242–248.

Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting memories* (pp. 71–94). Cambridge, MA: MIT Press.

Yi, D.-J., Woodman, G. F., Widders, D., Marios, R., & Chun, M. M. (2004). Neural fate of ignored stimuli: Dissociable effects of perceptual and working memory load. *Nature Neuroscience, 7*(9), 992–996.