Explicit Mechanisms Do Not Account for Implicit Localization and Identification of Change: An Empirical Reply to Mitroff et al. (2002)

Diego Fernandez-Duque
University of Toronto

Ian M. Thornton
Max Planck Institute for Biological Cybernetics

Several recent findings support the notion that changes in the environment can be implicitly represented by the visual system. S. R. Mitroff, D. J. Simons, and S. L. Franconeri (2002) challenged this view and proposed alternative interpretations based on explicit strategies. Across 4 experiments, the current study finds no empirical support for such alternative proposals. Experiment 1 shows that subjects do not rely on unchanged items when locating an unaware change. Experiments 2 and 3 show that unaware changes affect performance even when they occur at an unpredictable location. Experiment 4 shows that the unaware congruency effect does not depend simply on the pattern of the final display. The authors point to converging evidence from other methodologies and highlight several weaknesses in Mitroff et al.'s theoretical arguments. It is concluded here that implicit representation of change provides the most parsimonious explanation for both past and present findings.

Psychologists have long been interested in exploring the differences between implicit and explicit processing. This interest has crossed many domains, including memory (e.g., Milner, Corkin, & Teuber, 1968; Schacter, 1995), perception (e.g., Cowey & Stoerig, 1991; Marcel, 1983) and action (Castiello, Paulignan, & Jeannerod, 1991; Goodale, 1996). It reflects the general belief that understanding such differences could lead to important insights into the functioning of the brain.

Recently, we raised the possibility that the processing of change could occur implicitly (Fernandez-Duque & Thornton, 2000; Thornton & Fernandez-Duque, 2000). We suggested that the well-documented phenomenon of *change blindness*—in which observers often fail to explicitly report large changes to a visual environment (Rensink, 2002; Rensink, O'Regan, & Clark, 1997; Simons & Levin, 1997)—could be underestimating the ability of the visual system to represent and process change. One interpretation, typical of change blindness studies, is that large portions of the visible world go unrepresented (see Noë, Pessoa, & Thompson, 2000; O'Regan, 1992, for similar views). However, such studies have always relied on explicit reports; thus, they have measured the limits of perceptual awareness rather than the richness and stability of perceptual representations.

More specifically, being unaware of a change does not entail that a change has no impact on the visual system. Indeed, in our previous work, we have shown that changes can still affect performance even when they are not explicitly detected (Fernandez-Duque & Thornton, 2000; Thornton & Fernandez-Duque, 2000). Other lines of evidence also support the existence of implicit processing of change (Beck, Rees, Frith, & Lavie, 2001; Chun & Nakayama, 2000; Fernandez-Duque, Grossi, Thornton, & Neville, 2003; Hayhoe, Bensinger, & Ballard, 1998; Henderson & Hollingworth, 1999; Hollingworth & Henderson, 2002; Rensink, in press; Ryan, Althoff, Whitlow, & Cohen, 2000; Smilek, Eastwood, & Merikle, 2000; Watanabe, 2003; for a literature review, see Thornton & Fernandez-Duque, 2002; for a contrasting perspective, see Simons & Silverman, in press).

Despite this growing body of evidence, there is general agreement among researchers on the importance (and difficulty) of ruling out the contribution of explicit processes when making claims about potential implicit effects (Reingold & Merikle, 1990). In line with this concern, Mitroff, Simons, and Franconeri (2002) reexamined a number of the key behavioral findings relating to implicit change detection. In their article titled "The Siren Song of Implicit Change Detection," Mitroff et al. suggested that "although implicit change detection remains a theoretical possibility, the evidence to date does not provide compelling support for claims of implicit registration, localization, or identification" (p. 813). They concluded that all previously reported findings "could be explained by explicit mechanisms" (p. 813).

In the current article, we directly explore two of the central claims of Mitroff et al. (2002)—those relating to implicit localization and implicit identification—and find that in both cases the authors fail to provide viable alternatives to implicit mechanisms. In the first three experiments, we provide direct evidence that contradicts both the predictions (Experiment 1) and the findings (Experiments 2 and 3) of Mitroff et al.'s work. In Experiment 4, we rule out one other mechanism that might have explained the

findings without having to appeal to an implicit processing of change. We conclude that our initial proposal for implicit mechanisms remains the most parsimonious explanation for the data.

Experiment 1: A Direct Assessment of the Exclusion Strategy

Fernandez-Duque and Thornton (2000) demonstrated that, at least with simple displays, observers can report the location of an undetected change better than chance in a two-alternative forced-choice task (for data relating to complex scenes, see Mitroff & Simons, 2002). In our studies, observers were shown an array containing 8, 12, or 16 rectangles (half horizontal, half vertical) for 250 ms. The computer screen then went blank for 250 ms, after which the array reappeared with one of the items in a new orientation. The changed item and the item diametrically opposite it were then cued, and the observers were asked to choose which one had changed by clicking on the item with the mouse. Subsequently, observers were asked to report whether they had seen a change by making a keypress response. The finding of interest was that even when observers reported no conscious awareness of the change, they still selected the correct item at above-chance levels. We interpreted this result as an indication that the change had been detected at some level in the visual system and that this "nonconscious" trace was sufficient to drive localization even in the absence of awareness. Additional experiments using a similar paradigm revealed that the change was effective as an attentional cue when observers were aware of the change but not when they were unaware of the change. This dissociation between aware and unaware trials suggested that the localization of unaware changes was not due to a contamination from trials in which observers were aware of the change but failed to report it.

We did consider the possibility that better-than-chance localization might be due in part to an exclusion strategy by which observers chose the item opposite to the one they knew had not changed (see Fernandez-Duque & Thornton, 2000, pp. 339–340 and Footnote 5). That is, if observers were sure that one of the cued items had not changed, they should select the opposite item as the best candidate for the change. We produced a model based on this exclusion hypothesis that could be made to fit our experimental data quite closely. A similar model applied by Mitroff et al. (2002, Experiment 2) fits their data closely. An important limitation of these models is that they are highly dependent on assumptions about how many items can be *held* by the visual system from one display to the next.

The number of items held by the visual system refers to those items that are thought at a given time to be processed to a level at which they can be accessed by conscious awareness. If a held item changes orientation, observers should be aware of that change. Conversely, if no change is detected, and observers somehow know which items are being held, they will be able to exclude these unchanged items as candidates for the changed target. The assumption about how many items are being held at any one time is crucial to the goodness of fit of these models. Assume that two items are held in mind across a flickering display, and the models closely fit the data reported by Fernandez-Duque and Thornton (2000) and by Mitroff et al. (2002). Alternatively, assume that only one item is being held, as is the case in Experiment 2 of Mitroff et al., and these models predict an effect half the size of the one

revealed by the empirical data in the eight-item display (14% vs. 7%). Assume that three items are being held, as shown in some studies by Rensink (2000), and the models predict an effect 50% larger than the one empirically obtained (21% vs. 14%).

Despite the shortcomings of these models, the exclusion strategy itself remains a possible alternative to implicit localization of change, and thus it merits further investigation. As mentioned above, the exclusion hypothesis suggests that observers will use their knowledge about a no-change item to infer, by exclusion, the location of a changed item. Therefore, the hypothesis depends on the reasonable assumption that observers will be able to locate no-change items better than chance. Experiment 1 directly tested this assumption.

In a two-alternative forced-choice task, observers were instructed to select the no-change item (i.e., the item opposite to the one that had changed). If observers in the unaware trials of Fernandez-Duque and Thornton (2000) were first identifying the no-change item and then locating the change by exclusion, selection of the no-change item in Experiment 1 should be better than chance. Chance or worse-than-chance performance in the unaware trials would argue against the exclusion strategy.

More specifically, the design of Experiment 1 made it possible to look for qualitative differences between aware and unaware changes. As in previous experiments (Fernandez-Duque & Thornton, 2000), observers were expected to perform better than chance when they were aware of the change. More importantly, observers were predicted to perform below chance in the unaware trials. This hypothesis was based on the reasoning that although the goal of the task was to choose the no-change item, the unaware change might act as a cue and thus bias observers' selection toward the location of change. This bias would lead observers to choose the wrong item (i.e., the changed item) more often than the target, driving their performance below chance level.

Method

Participants. Nine students from the University of Toronto participated in this experiment for course credit. All were right-handed, had normal or corrected-to-normal vision, and were naive as to the purpose of the experiment.

Equipment. Stimulus display and response collection were carried out on a Macintosh computer using custom-written software. Many of the routines were based on work by Steinman and Nawrot (1992), Pelli and Zhang (1991), and Rensink (1990).

Stimuli. Each display consisted of a ring of eight rectangles in a clockface design. Each rectangle was drawn in black on a uniform gray background. The same gray level was also used for a blank field that separated the first and second change displays. Each rectangle measured 10×30 pixels, which subtended $0.46° \times 1.38°$ visual angle. The ring arrangement ensured that all items were equidistant from each other and from the central fixation cross, which was approximately 4.6° from each item.

In the first display of each trial, half of the rectangles were horizontal and half were vertical (the distribution of orientation by position was randomized for each trial). In the second display, following the blank screen, the orientation of one rectangle was switched by having it rotate 90° about its center point. The duration of each display and of the blank interstimulus interval was fixed at 250 ms.

Each observer completed a total of 288 trials. In two thirds of the trials, a rectangle changed orientation, and in the other third (i.e., catch trials) no change occurred. The orientation change could occur with equal probabil-

ity at any one of the locations, with each object playing host to the change for 24 trials. Half of these changes at each location were horizontal to vertical, and half were vertical to horizontal. When a change occurred at a given location (the target), the orientation of the item diametrically opposite it (the distractor) was constrained. That is, after the change, half of the displays had *congruent* targets and distractors (i.e., both were either horizontal or vertical), and half had *incongruent* distractors (i.e., the target was vertical and the distractor horizontal, or vice versa). In the final display, two of the items, the changed target and the opposite distractor, were identified by their color turning from black to light gray.

Procedure. For each trial, observers were asked to fixate on a central cross until the target and distractor were highlighted. They were informed that on some trials, a change of orientation would occur and they should try to notice such a change. During the third display, when the target and distractor rectangles were highlighted, the observers' task was to move the central fixation point with the mouse and click on the rectangle opposite to the one they thought was most likely to have changed orientation. The mouse was always located in front of the observer, centered relative to the middle of the screen. Although being asked to "select the item opposite to the one that has changed" is functionally identical to being asked to "select the item that did not change," we chose the former wording to stress the importance of the change itself. This was done to bring the task relevance of the change and the general "mindset" of observers closer to those in our previous experiments on implicit localization of change (Fernandez-Duque & Thornton, 2000).

Having indicated their choice of target, observers where then asked to indicate whether they had actually seen an item change by pressing one of two marked keys. Observers were told to choose the key marked *aware* if they had noticed or sensed any change at all between the first two displays. The key marked *unaware* was to be chosen only if observers sensed nothing to indicate that a change had occurred. The instruction to adopt such a liberal criterion for awareness was aimed at minimizing the possible contamination of unaware responses by mislabeled aware responses. Observers received no feedback on whether their selection of the changed item was correct or incorrect or whether their awareness response was appropriate. After illustration of the general procedure, observers were given a number of practice trials and then completed a total of 288 trials.

Exclusion criteria. Observers who showed either extremely high or extremely low levels of awareness were taken to be not performing the task as instructed. No observers were excluded on this basis.

Results

Observers reported being aware of the change in 56% of trials with a change and in 21% of catch trials. This percentage of aware trials is comparable to those found in our previous experiments (Fernandez-Duque & Thornton, 2000). The high percentage of false alarms is consistent with the instructions to adopt a liberal criterion for reporting a change. The discrimination index d_L revealed that, although they missed many changes and made some false alarms, observers were able to discriminate between the change trials and the catch trials at better-than-chance levels, d_L 1.850, $t(8) = 5.5$, $p < .01$, confidence interval (CI) = 1.07–2.62. When observers were aware of the change, they correctly chose the item opposite to the one that had changed in 81% of trials. Such performance was better than chance, $t(8) = 12.0$, $p < .01$, CI = 75%–86%, but by no means perfect, consistent with a liberal criterion for reporting awareness of change. More important, when observers were unaware of the change, they chose the target (i.e., the item opposite to the one that had changed) in only 46% of trials, a level of performance that is not better than chance, $t(8) = -1.2$, *ns,* CI = 40%–53%.

Discussion

The main goal of Experiment 1 was to directly assess the exclusion strategy as a mechanism for better-than-chance performance in studies of implicit localization of change. The findings clearly refute that possibility. When observers were unaware of the change, they performed at chance level at selecting the item that had not changed.

The second goal of the experiment was to look for a dissociation between aware and unaware performance. Although performance in unaware trials was 4% below chance, this effect failed to reach statistical significance. Thus, we found no evidence for a dissociation between aware and unaware trials. However, such a test sets a very stringent criterion, requiring not only that the change be implicitly registered but also that such registration not be available to influence goal-directed behavior. That is, implicitly registered changes must fail to directly influence the decision-making process by which observers choose the item opposite to the one that changed. Such a strict criterion is likely to underestimate the influence of unaware processes. In Experiments 2 and 3, we looked for evidence of dissociation between aware and unaware trials using different measures.

Experiment 2: Implicit Identification of Change and the Congruency Effect

One of the most compelling lines of evidence for the existence of implicit change detection comes from a series of experiments by Thornton and Fernandez-Duque (2000) in which responses to a simple orientation judgment—whether a rectangle was horizontally or vertically oriented—were shown to be biased by prior, unreported changes. This finding, deriving from a well-established, indirect measure—the congruency effect—is important because it suggests that in the absence of awareness, the visual system may maintain not only the fact the some change has taken place, or even just the possible location of that change, but also something about the identity of the change.

In a typical congruency task, the presence of distracting information can be shown to hurt performance when it conflicts with the target information. For example, the presence of a horizontal distractor leads to increased errors and longer response times (RTs) in reporting a vertical target. Examples of congruency effects abound in the literature, including the Simon task (Lu & Proctor, 1995; Simon & Small, 1969), the Flanker task (B. A. Ericksen & Ericksen, 1974), and the Stroop task (MacLeod, 1991; Stroop, 1935/1992). There are important differences among these paradigms (Lu & Proctor, 1995), but they all share the basic principle that incongruent information, when processed, can be detrimental to performance.

In some paradigms, the distracting information temporally precedes the target, and in those cases, the *prime* item highlights certain stimulus properties (e.g., orientation, location in space, meaning) that are also present in the *probed* target. When the prime is compatible (congruent) with the probe, it leads to better performance (higher accuracy and faster responses) than found in neutral trials, in which no priming takes place. When the prime is incompatible (incongruent) with the probe, it leads to impaired performance (increased errors and slower responses) relative to neutral trials.

In our experiment, observers were presented with a simple ring of eight rectangles that appeared briefly, was replaced by a blank screen, and then reappeared. On change trials (66%), one of the items changed orientation between the first and second presentations. This change constituted the prime phase of the congruency task. A change of orientation, say from horizontal to vertical, would increase the salience of *vertical,* the ending state of the change. Following the change sequence, one of the rectangles in the ring was highlighted, and observers were instructed to make a speeded response based on this probe item's orientation. For trials in which observers were aware of the change, probes with orientations incongruent to the changed item were reported more slowly and less accurately than were congruent probes. When observers were unaware of the change, a congruency effect was still found, but only in the error rates. Our explanation for the unaware error congruency effect was that even though the change was undetected, it was still registered, and the appropriate final orientation was primed. This priming then influenced the response to subsequent probes.

The finding of an unaware congruency effect by Thornton and Fernandez-Duque (2000) was replicated by Mitroff et al. (2002) in their Experiment 4A. However, Mitroff et al. favored an explanation based on explicit mechanisms, focusing on another aspect of our displays. They noted that probe items always occurred at a location diametrically opposite to the change location. They suggested that observers could learn this spatial relationship between the probe and the change and then use this association to shift attention back across the display, from the probe to the opposite item. This shifting of attention to the opposite item, not the change, would then raise the salience of the changed object, indirectly leading to a congruency effect. Thus, the change itself would only play a role in setting up the explicit association.

In support of their claim, Mitroff et al. (2002, Experiment 4B) reported results from a follow-up experiment in which the spatial link between the probe and the change was severed, and the congruency effect disappeared. However, it is important to note that the congruency effect in this study disappeared even when observers were fully aware of the change. The absence of a congruency effect in the aware trials seriously complicates the interpretation of these data. That is, the findings are uninformative about implicit detection of change because the modifications introduced by Mitroff et al. depleted their paradigm of the ability to reveal any congruency effect, unconscious or otherwise.

There are at least two reasons why the congruency effect might have been completely absent in Experiment 4B of Mitroff et al. (2002). First, accuracy rates for congruent trials in that experiment were at close to chance level (60.6%). In contrast, the other experiments that did show a congruency effect had higher accuracy rates in the congruent condition (77% in Experiment 4A of Mitroff et al., 2002; 82% in Experiment 2 of Thornton & Fernandez-Duque, 2000). Given that chance level in these two-alternative forced-choice tasks was 50%, accuracy rates of 60% in the control condition (congruent trials) might have been too low for an effect to reveal itself, although it is conceivable that incongruent trials could achieve below-chance levels.[1]

Another possible explanation for the lack of congruency effect in Experiment 4B of Mitroff et al. (2002) is the overall slowing of responses in that task. Congruency effects are known to be modulated by the speed of response, and in many cases the effect is abolished when RTs are too long (DeJong, Liang, & Lauber, 1994; Lu & Proctor, 1994). Unlike observers in the original experiment, those in Experiment 4B of Mitroff et al. did not receive speed feedback and therefore were not encouraged to respond in a speeded manner. The average RT for Experiment 4B, in which no congruent effect was found, was significantly longer than that in Experiment 4A, and it was 300 ms longer than those found in similar experiments by Thornton and Fernandez-Duque (2000).

To test the idea that the null effect reported by Mitroff et al. (2002, Experiment 4B) was not due to the lack of a spatial link between the probe and the changing item but instead was due to some other artifact, such as a floor effect or long RTs, we tested a modified version of our original experiments (Thornton & Fernandez-Duque, 2000). We removed the spatial contingency between probe and change in the invalid trials (i.e., trials in which the probed item was not the changed item). We also removed the feedback relating to accuracy but kept the feedback relating to speed of response, to encourage observers to respond as quickly as possible. Finally, we increased the duration of the probe to 40 ms to make the task less challenging, in the hope of speeding responses and increasing overall accuracy rates.

If change, as suggested by Mitroff et al. (2002), was not the cause of the congruency effect, our emphasis on speeded responses and lower error rates should be inconsequential. If, in contrast, a congruency effect were to emerge under these conditions, it would rule out the associative learning hypothesis and favor instead a more direct role for the change itself.

Method

Participants. Ten students from the University of Toronto participated in this experiment for course credit. All were right-handed, had normal or corrected-to-normal vision, and were naive as to the purpose of the experiment.

Equipment and stimuli. These were identical to those used in Experiment 1, except that in the third display, only one item was highlighted. This item was the probe for the speeded orientation response. It changed color from black to light gray and remained visible for 40 ms. During this time, the remaining seven rectangles were also visible. After 40 ms, the whole screen went blank.

Design. The design of this experiment was closely modeled on Experiment 1 of Thornton and Fernandez-Duque (2000) rather than on the exact design of Mitroff et al. (2002, Experiment 4B). We chose to replicate our previous work because these are the findings that have been challenged.

The most important design feature of this experiment concerns the relationship between the changed item and the subsequent orientation probe. Each rectangle in the ring was probed on 56 trials. In 25% of these trials, no change had taken place. Such catch trials provide useful information for determining sensitivity to change.

In a further 25% of the trials, a change was present and the orientation probe was identical to the changed item, sharing both its location and orientation (*valid* trials). For the remaining 50% of the trials, a change was also present, but the location of this changed item had no systematic relation with the location of the probed item (*invalid* trials).

---

[1] Low accuracy rates in Experiment 4A of Mitroff et al. (2002) also complicate the interpretation of their null results for RT data. Even if average RTs are estimated from correct responses only, it is reasonable to assume, given the low overall accuracy rates, that many of the correct responses were guesses.

The lack of a spatial relation between the changed item and the probe on invalid trials was the main modification of the present experiment, relative to our previous work. The relation between the orientation of the probe and the orientation of the changed item on these trials was, however, systematically varied. In half of these trials (25% of total trials), the final orientation of the changed item was the same as that of the probed item (congruent trials). For the remaining half of the invalid trials (25% of total trials), the final orientation of the changed item had the opposite orientation to that of the changed item (incongruent trials). Differences in performance between invalid congruent and invalid incongruent trials are of most relevance for exploring the congruency effect.

Procedure. For each trial, observers were asked to fixate on a central cross until the probe item was highlighted in the final display. They were informed that on some trials, a change of orientation would occur and that they should try to notice such a change. They were also informed that during the third display a probe item would be highlighted, and they were instructed to report as rapidly as possible the orientation of the probe by pressing one of two indicated keys (letters *L* and *S* on the keyboard, counterbalanced across observers). Half of all trials, across all conditions, required a horizontal response, and half required a vertical response.

Following the speeded orientation response, observers indicated whether they had been aware of any change during the flicker part of the display. This awareness response was collected using a go/no-go protocol. After observers had responded to the probe's orientation, a question mark appeared on the center of the screen. Observers were instructed to press the spacebar if they had been aware of the change (the go response) or do nothing if they had been unaware (the no-go response). Regardless of the nature of the awareness response, the subsequent trial began after a delay of 2 s. We used this go/no-go method to reduce task demands. Previous research from our laboratories has shown similar findings, irrespective of the type of method used to report awareness of a change (Thornton & Fernandez-Duque, 2002). As in Experiment 1, observers were instructed to adopt a liberal criterion for awareness.

To familiarize observers with the display, each session began with 20 practice trials in which the probe item remained visible for 200 ms. This was followed by 20 practice trials with a 40-ms probe. During practice, feedback was provided on both the accuracy of orientation responses (single beep) and the speed (double beep). The main experimental session was divided into four blocks of 112 trials, for a total of 448. During these trials, the duration of the probe was fixed at 40 ms, and feedback was provided only on the speed of response (double *beep* to responses slower than 1,200 ms).

Exclusion criteria. One observer was replaced because he reported more false alarms than hits, suggesting either that he had not understood the instructions or had not cooperated with the task. A second observer was excluded because during debriefing at the end of the study it became clear that he had misunderstood the instructions and reported awareness of a change only when the probed item had been changed. Analysis was thus conducted on data from 8 observers.

Results

As found in other studies of change blindness, observers were generally quite poor at detecting changes. In trials with a change, observers indicated being aware of the change 61% of the time (hit rate), whereas changes were reported on 18% of catch trials (false alarm rate). These values are comparable to results from our previous studies (e.g., Thornton & Fernandez-Duque, 2000, Experiment 2). The high percentage of false alarms indicates that observers properly followed the instruction to adopt a liberal criterion for reporting awareness of change. The discrimination index d_L revealed that, although they missed many changes and made some false alarms, observers were able to discriminate between the change trials and the catch trials at better-than-chance levels, $d_L = 2.330$, $t(7) = 6.4$, $p < .01$, CI = 1.9–2.7. These findings also replicate previous results (Thornton & Fernandez-Duque, 2000).

As is evident in Table 1, responses were much faster, much less variable, and much more accurate than those reported by Mitroff et al. (2002, Experiment 4B), showing the expected influence of the RT feedback and longer probe duration. As in our previous experiments (Thornton & Fernandez-Duque, 2000), only correct responses were included when computing median RTs. Consistent with the idea that observers were engaged in a speeded response, observers were slower in responding after detecting a change than they were when the change went undetected, $F(1, 7) = 19.0$, $p < .01$.

The main question motivating this experiment was whether the absence of a spatial relation between probe and distractor would abolish the congruency effect for observers encouraged to make a speeded response. Data from invalid trials (i.e., trials in which the

Table 1

Accuracy Percentages and Median Response Times for Unaware and Aware Responses in Experiment 4B of Mitroff et al. (2002) and Experiment 2 of the Current Study

	% accuracy		Response time (ms)	
Trial type	Unaware	Aware	Unaware	Aware
Mitroff et al. (2002, Experiment 4B)				
No change	64.8 (11)	62.1 (22)	1,085 (377)	1,092 (397)
Valid	71.3 (20)	80.1 (14)	1,132 (579)	1,034 (404)
Invalid congruent	60.4 (12)	60.8 (8)	1,112 (386)	1,080 (390)
Invalid incongruent	61.3 (12)	58.3 (10)	1,076 (349)	1,072 (372)
Experiment 2				
No change	94.7 (3.5)	96.0 (4.5)	610 (47)	676 (43)
Valid	95.1 (4.4)	92.0 (9.5)	589 (40)	625 (71)
Invalid congruent	92.0 (5.5)	90.9 (4.2)	617 (60)	697 (66)
Invalid incongruent	87.8 (4.5)	83.9 (9.7)	649 (41)	727 (43)

Note. Aware responses to no-change trials are false alarms. Standard deviations are given in parentheses.

changed item was not the one probed) were submitted to a within-subject analysis of variance (ANOVA) with two variables: awareness of change (yes, no) and congruency (congruent, incongruent). When the probed item had an orientation incongruent to the changed item, observers made more errors, $F(1, 7) = 7.8, p < .03$, and were slower, $F(1, 7) = 7.0, p < .03$, than when the orientation was congruent. These effects were present both in aware and unaware trials, and there were no interactions between awareness and congruency in accuracy, $F(1, 7) = 0.6, ns$, or in RT, $F(1, 7) = 0.1, ns$.

Next, we looked at whether the congruency effect was modulated by the spatial relation between the probed location and the location of change and whether this modulation was different for aware and unaware trials. These analyses were designed to reveal a possible dissociation between aware and unaware trials. The data were split into *close* and *far* trials. In close trials, the changed item appeared next to the probe item ($45°$); in far trials, the changed item and the probe were more distant from each other ($90°$, $135°$, and $180°$). We pursued this analysis, instead of including each bin separately, because the error rate was rather small and there were not enough errors to achieve stable estimates. A preliminary analysis revealed no statistical differences in the congruency effects of these three bins. We explored both RTs and accuracy rates (see Table 2).

For the RT data, there was a strong trend toward a three-way Awareness × Congruency × Location interaction, $F(1, 7) = 4.9$, $p < .06$. The influence of probe location on the congruency effect was conditional on whether the observer was aware of the change. Follow-up analyses for each level of probe location were conducted. These revealed that aware trials led to larger congruency effects than did unaware trials, but only at the close location: congruency by awareness interaction at the close location, $F(1, 7) = 12.0, p < .01$; congruency by awareness interaction at the far location, $F(1, 7) = 0.1, ns$ (see Table 2). One possible interpretation of this result is that the close spatial proximity helped observers to slow down and resolve conflict, but only when they were aware of the change. By being aware of an incongruent change

close to the probe, observers might also have been able to minimize their errors by slowing down. This interpretation is consistent with findings from previous experiments showing the congruency RT effect to be most evident during aware trials (Thornton & Fernandez-Duque, 2000). This interpretation also makes some predictions regarding the pattern of errors. For incongruent trials, it predicts that close proximity to the change should lead to an increased number of errors, but only for unaware trials, which are the ones in which observers are least likely to slow down.

For the accuracy data, the three-way Awareness × Congruency × Location interaction failed to reach statistical significance, $F(1, 7) = 2.5, p < .15$. However, further inspection of the data did reveal the presence of the predicted pattern. Specifically, close proximity of change and probe led to an increased congruency effect, but only in the unaware trials, $t(7) = 3.4, p < .01$. For aware trials, the cost of incongruent trials at the close location was not significant, suggesting that by slowing down in incongruent aware trials, observers were able to minimize errors, $t(7) = 0.4, ns$.

Discussion

The results of this experiment show that a spatial link between the changed item and the probed item was not responsible for the congruency effect observed in our previous experiments (Thornton & Fernandez-Duque, 2000). Thus, we found no evidence for the explicit mechanism that Mitroff et al. (2002) offered as an alternative to the implicit detection of change. Specifically, when the link between the changed and probed items was removed, healthy aware and unaware effects were still observed as long as observers were encouraged to make speeded responses and a longer probe duration prevented floor accuracy effects.

The presence of an awareness main effect in the RTs in the current experiment further suggests that observers were making speeded responses and were slowed down by the detection of a change, possibly in an attempt to reduce errors. This is a typical finding in speeded dual-response tasks such as this. The fact that this awareness effect was absent in Mitroff et al.'s Experiment 4B (2002) further supports the view that the observers in their experiment were not making speeded responses.

Although the current results are consistent with implicit representation of change, they reveal an unaware RT congruency effect that was not present in our previous studies (e.g., Thornton & Fernandez-Duque, 2000). In the past, we have interpreted the absence of such an effect as evidence that unaware responses were not being contaminated by unreported aware trials. Despite the lack of such a dissociation in the current study, there are other reasons to believe that the unaware congruency effects are genuine. For example, responses to aware trials were generally much slower than responses to unaware ones. This suggests some real difference between these two conditions and indicates that contamination, if present, would have to be the exception rather than the rule.

Furthermore, for a minority of contaminated trials to account for the current data, the magnitude of the congruency effect on those trials would need to be much larger than that found on the aware trials. This follows from the fact that the size of the RT congruency effect for the unaware condition was as large as that for the aware condition. A contamination by aware trials should therefore have increased the variability in the unaware condition. In contrast to

Table 2

Accuracy Percentages and Median Response Times on Invalid Trials at Close and Far Locations in Experiment 2

Probe/change relation	% accuracy		Response time (ms)	
	Unaware	Aware	Unaware	Aware
Invalid trial data				
Close				
Congruent	96 (4)	89 (16)	644 (63)	698 (47)
Incongruent	83 (10)	86 (8)	657 (50)	754 (38)
Far				
Congruent	89 (7)	90 (5)	626 (58)	712 (72)
Incongruent	89 (7)	83 (11)	665 (56)	750 (45)
Congruency effect				
Close	13	3	13	56
Far	0	7	29	38

Note. Close = $45°$; Far = $90°$, $135°$ or $180°$. Standard deviations are given in parentheses.

this prediction, an inspection of the standard deviations for unaware conditions revealed no evidence of such increased variability. Finally, a contamination by aware trials could not explain why a lack of awareness decreased the congruency effect for RTs but increased it for error rates, as was the case when the change occurred close to the probe.

The unaware congruency effect for RTs in the current experiment may stem from the speeding up of the overall RTs. Observers in the current experiment were given instructions that emphasized speed. Consequently, RTs were consistently shorter than they were in our original experiments (Thornton & Fernandez-Duque, 2000). It is also possible that the unaware congruency effect became evident due to the higher accuracy rates. With increasing numbers of correct trials, the interpretation of RT patterns becomes more reliable.

In summary, the findings from Experiment 2 reveal that the congruency effect was not driven by the artifact proposed by Mitroff et al. (2002). When that potential artifact was removed by severing the spatial link between change and probe, an implicit effect was still observed.

Experiment 3: Congruency Effect in Fast and Slow Responders

In Experiment 2, we purposefully speeded responses by adding feedback when observers were too slow. We also heightened accuracy rates by displaying the probe for a longer duration. These manipulations were effective in revealing a congruency effect in the absence of a spatial link between the change and the probe, even when observers were unaware of the change. In Experiment 3, we asked whether these findings would generalize to a briefer probe and lower accuracy rates. We also explored more directly whether slow responses might account for the absence of a congruency effect in Experiment 4B of Mitroff et al. (2002) by analyzing observers with fast and slow responses separately.

Another motivation for Experiment 3 was to assess whether the unaware congruency effect was present for accuracy rates only (as in Thornton & Fernandez-Duque, 2000) or was present also for RT effects (as in Experiment 2). A lack of RT congruency effect in unaware trials under these conditions would serve to rule out contamination by aware trials.

Method

Participants. Twelve students from the University of Toronto participated in this experiment for course credit. All participants were right-handed, had normal or corrected-to-normal vision, and were naive as to the purpose of the experiment.

Equipment, stimuli, and procedure. These were identical to those used in Experiment 2, except for the following modifications. We did not include speed feedback, we made the probe visible for only 20 ms (rather than 40 ms), and we instructed subjects to respond as accurately as possible, emphasizing that accuracy was more important than speed.

Exclusion criteria. Observers who showed either extremely high or extremely low levels of awareness were taken to be not performing the task as instructed. One participant reported being aware in only 2% of trials. Thus, the following analyses were carried out on 11 of the 12 sets of data collected.

Results

As expected, given the lack of speed feedback and the shorter probe duration, responses were slower, more variable, and less accurate than they were in Experiment 2 (see Table 3). In this sense, the data resemble those reported in Experiment 4B of Mitroff et al. (2002).

We first analyzed the data from all subjects as a single group, and then we split the data into *fast* and *slow* groups. The slow group comprised observers with median RTs longer than 1,000 ms ($n = 6$).

Awareness. In trials with a change, observers indicated being aware of the change 66% of the time (hit rate), whereas changes were reported on 18% of catch trials (false alarm rate). These values are comparable to those from Experiment 2 and our previous studies (e.g., Thornton & Fernandez-Duque, 2000, Experiment 2). The high percentage of false alarms indicates that observers properly followed the instruction to adopt a liberal criterion for reporting awareness of change. The discrimination index d_L revealed that, although they missed many changes and made some false alarms, observers were able to discriminate between change trials and catch trials at better-than-chance levels, $d_L = 2.34$, $t(9) = 7$, $p < .01$. These findings also replicate results from our previous experiments (Thornton & Fernandez-Duque, 2000).

There were no significant differences in change discriminability between observers in the fast ($d_L = 2.2$) and slow ($d_L = 2.4$) groups, $t(9) = 0.7$, *ns*. When there was a change, observers in the slow group were more likely to report it than were observers in the fast group (73% vs. 57%), $t(9) = 2.8$, $p < .02$, but they also had a tendency to report more false alarms (20% vs. 14%), $t(9) = 1.7$, $p < .11$.

Accuracy. Data from invalid trials (i.e., trials in which the changed item was not the one probed) were submitted to a within-

Table 3
Accuracy Percentages and Median Response Times for Invalid Trials

Group and condition	% accuracy		Response time (ms)	
	Unaware	Aware	Unaware	Aware
Slow				
Congruent	78.6 (7)	82.9 (8)	1,263 (201)	1,303 (443)
Incongruent	66.5 (11)	56.2 (13)	1,272 (270)	1,300 (267)
Fast				
Congruent	80.0 (7)	88.7 (6)	752 (138)	715 (184)
Incongruent	72.1 (7)	45.2 (24)	765 (122)	836 (164)

Note. The slow group comprises observers with median response times longer than 1,000 ms. Standard deviations are given in parentheses.

subject analysis of variance with two variables: awareness of change (yes, no) and congruency (congruent, incongruent). When the probed item had an orientation incongruent to the changed item, observers made more errors than when the orientation was congruent, $F(1, 10) = 26$, $p < .01$. Thus, unlike in Experiment 4 of Mitroff et al. (2002), we found a significant congruency effect. Although this congruency effect was larger when observers were aware of the change [Congruency $\times$ Awareness interaction: $F(1, 10) = 11$, $p < .01$], a follow-up analysis revealed that the congruency effect was also significant for unaware trials, $F(1, 10) = 17$, $p < .01$. The effect was present both for aware and unaware trials in both the fast and slow groups [fast group, aware: $t(4) = 3$, $p < .04$; fast group, unaware: $t(4) = 3$, $p < .04$; slow group, aware: $t(5) = 4$, $p < .01$; slow group, unaware: $t(5) = 3$, $p < .03$].

RT. Data from correct responses to invalid trials were submitted to a within-subjects analysis of variance with two variables: awareness of change (yes, no) and congruency (congruent, incongruent). This analysis showed no congruency effect, $F(1, 10) = 0.6$, *ns.* In this sense, the data replicate the results of Experiment 4B of Mitroff et al. (2002).

However, when the data were split into fast and slow groups, a more interesting pattern emerged. As expected, RTs for observers in the slow group were more variable than were those for observers in the fast group. More important, for slow-group observers, the congruency effect was absent both in the aware and in the unaware trials. An inspection of the data reveals a 3-ms effect in the reverse direction for aware changes, $t(5) = -0.1$, *ns,* and a 9-ms effect for unaware ones, $t(5) = 0.3$, *ns.* The pattern of data for observers with short RTs is markedly different. For these observers, there was a large (121 ms) congruency effect when they were aware of the change, $t(4) = 3.0$, $p < .04$, but only a small (13 ms) nonsignificant difference when they were unaware of the change, $t(4) = 0.3$, *ns.*

Discussion

The findings of Experiment 3 generalize those of Experiment 2 to a situation in which the task was made more difficult by reducing the duration of the probe, leading to an overall reduction in accuracy rates. Unlike the results reported by Mitroff et al. (2002) in their Experiment 4B, there was clear evidence that incongruency led to an increased number of errors, both in aware and unaware trials. One possible explanation for this discrepancy lies with the accuracy rates for congruent trials, which, despite an overall drop in performance, were still higher in this experiment than were those in Experiment 4B of Mitroff et al. Those trials serve as a control condition against which to compare performance in incongruent trials. Thus, congruent trials with accuracy rates close to chance may have obscured the cost of incongruency in Experiment 4B of Mitroff et al. Another difference between the two paradigms was the percentage of valid trials (i.e., trials in which the change item was probed). Mitroff et al. used a lower percentage of valid trials than we did. It is conceivable that this spatial relation between the change and probed item had, over the course of the experiment, an influence on the invalid trials (i.e., trials in which the probe occurred at a no-change location). For example, a large number of valid trials might have biased the attentional control settings toward the changed item, because changed items would then have been probed on a disproportionate

number of trials. When an invalid trial occurred, the bias would have enhanced the influence that the changed item had on performance. Consistent with this proposal, there is already evidence that observers can implicitly learn location contingencies between cue and probe and that such learned contingencies can automatically influence attention allocation (Lambert, Naikar, McLachlan, & Aitken, 1999).

The RT data failed to show a congruency effect when responses from the slow and fast groups were analyzed together. In this regard, the findings were in agreement with those of Mitroff et al. (2002). However, as predicted, the absence of a congruency effect in the RTs could be traced to observers with very long RTs. For those observers, there was no congruency effect whatsoever. In contrast, for observers who adopted a speeded response, there was a large congruency effect, but only when they were aware of the change. This dissociation between aware and unaware trials (i.e., RT effect for aware trials only but accuracy effect for both aware and unaware trials) argues against a contamination hypothesis. Given that aware trials affect both RTs and accuracy, a contamination hypothesis would predict the unaware congruency effect in both measures.

Experiment 4: Ruling Out Final Number of Items as a Cause for Congruency

In reporting Experiments 2 and 3, we argue that the presence of a congruency effect in the error rates of unaware trials was an indication of implicit representation of change. However, in both of these experiments, the final frame of the change display contains an unequal number of horizontal and vertical items. For example, a change from vertical to horizontal would result in a final frame with three vertical items and five horizontal items. On an invalid trial with this configuration, the correct response to the probe item would be horizontal for congruent trials and vertical for incongruent trials. If an observer were unaware of the change but responded on the basis of number of items in the final display— that is, selected the orientation with the most exemplars— then they would be accurate on congruent trials and inaccurate on incongruent trials. Such a pattern of responding would predict the observed unaware congruency effect without having to invoke an implicit representation of change.

It seems unlikely that observers used the number of items in a strategic way. During debriefing at the end of all our previous experiments with this design, no observer ever reported using the number of items to guide their responses. Furthermore, change arrays appear too rapidly for observers to explicitly estimate the number of items of each orientation. Of course, given the main message of this article, we cannot rule out an alternative explanation simply because observers are unlikely to have explicit knowledge. Therefore, in this final experiment, we presented an equal number of horizontal and vertical items in the final change array.

Method

Participants. Fifteen university students took part in this study for partial course credit. All had normal or corrected-to-normal vision, were right-handed and were naive as to the purpose of the experiment.

Equipment and stimuli. The initial array contained five rectangles of one orientation and three of the other. This was to ensure that the second array contained an equal number of horizontal and vertical items. The

design also included catch trials, which had four horizontal and four vertical rectangles in each frame, as well as trials which began with a 5/3 distribution but ended with 6/2. These additional trials were included to reduce the probability that observers would try to use the distribution of orientations in the first display as an explicit or implicit predictor of the change. These 6/2 trials were not included in the main analysis because the question of central interest was whether an unaware congruency effect would be present in the absence of an orientation bias. As in Experiment 3, the final probe item remained visible for 20 ms only. Unlike in Experiment 3, the probed item in invalid trials was always the item diametrically opposite to the change.

Procedure. The procedure was altered from that of Experiment 3 in two respects. First, the pattern of orientations was modified so that on critical trials there were always four horizontal and four vertical items in the final change display. Second, instructions emphasized speed over accuracy in the orientation task.

Exclusion criteria. Observers who showed either extremely high or extremely low levels of awareness were taken to be not performing the task as instructed. Three observers were excluded on these grounds. Thus, the following analyses were carried out on data from 12 observers.

Results

Overall levels of awareness were quite high, with observers reporting a change on 69% of trials in which an item changed orientation. This level of awareness was probably due to a combination of a liberal awareness criterion and added visual uncertainty due to the briefly presented probe. Consistent with this explanation, observers also reported some form of change on 24% of catch trials. This was true for trials in which observers' responses in the orientation task were correct, $d_L = 2.069$, $t(11) = 11.4$, $p < .01$, CI = 1.67–2.47), and also for error trials, $d_L = 2.227$, $t(11) = 2.8$, $p < .02$, CI = 0.48–3.98).

Table 4 summarizes the data of Experiment 4. As for Experiment 3, there was a main effect of congruency in the error rates, $F(1, 11) = 28.3$, $p < .01$, which was modulated by awareness level, $F(1, 11) = 4.04$, $p = .07$. This marginal interaction resulted from a difference in magnitude between the congruency for aware trials (30%) and unaware trials (16%), but in both conditions the effect was significant, $t(11) = 5.23$, $p < .01$, CI = 17–42, and $t(11) = 2.97$, $p < .05$, CI = 4–27, respectively.

The RT data revealed a marginally significant congruency main effect, $F(1, 11) = 4.44$, $p = .06$, MSE = 4,899, which was driven exclusively by the aware trials, $t(11) = 2.02$, $p = .07$, CI = 5–121. There was no RT congruency effect in the unaware trials.

Discussion

The purpose of Experiment 4 was to examine whether the unequal number of horizontal and vertical items in the final display was the cause of the congruency effect in our previous experiments. Such an influence could be the consequence of an explicit strategy—for example, counting the number of items with each orientation and responding on the basis of the orientation of the smaller set—or could occur implicitly—for example, by the aggregate influence of each rectangle onto the orientation task. Either of those explanations would mean that a static feature of the postchange display, rather than the change itself, was the cause of the congruency effect. The results of Experiment 4 clearly argue against such a notion. Even when the number of vertical and horizontal items in the final frame was the same, a significant congruency effect was found.

General Discussion

In the four experiments reported in this article, we found no evidence to support Mitroff et al.'s (2002) claim that explicit mechanisms can account for the implicit localization and implicit identification of change (Fernandez-Duque & Thornton, 2000; Thornton & Fernandez-Duque, 2000). Experiment 1 demonstrated that the selection of unchanging items—an exclusion strategy— was not responsible for the ability to localize change in the absence of awareness. Specifically, we showed that subjects were at chance in localizing the nonchange item. Experiment 2 demonstrated that congruency effects do not depend on a spatial link between the changed item and a subsequent probe, directly contradicting Mitroff et al.'s associative learning hypothesis. Experiment 3 provided further evidence for such congruency effects in the absence of a spatial link and suggested that the lack of these effects in the aware trials of Mitroff et al.'s study was probably due to unusually long RTs or very large error rates. Experiment 4 ruled out another alternative interpretation, namely that the unequal number of vertical and horizontal items in the final display could account for the congruency effect.

In the current article, we have focused almost exclusively on our own behavioral evidence for the implicit localization and implicit identification of change (Fernandez-Duque & Thornton, 2000; Thornton & Fernandez-Duque, 2000), and we have addressed the specific concerns raised by Mitroff et al. (2002). However, other researchers have independently provided evidence for implicit

Table 4

Accuracy Percentages and Median Response Times for Unaware and Aware Responses in Experiment 4

Trial type	% accuracy		Response time (ms)	
	Unaware	Aware	Unaware	Aware
No change	77.4 (10.0)	70.0 (15.0)	649 (105)	663 (127)
Valid	82.9 (9.4)	86.2 (8.8)	640 (111)	596 (131)
Invalid congruent	80.5 (10.8)	82.3 (9.5)	646 (110)	633 (152)
Invalid incongruent	64.8 (17.8)	52.7 (21.0)	673 (163)	691 (129)

Note. Aware responses to no-change trials are false alarms. Standard deviations are given in parentheses.

change detection that is not subject to the criticisms of Mitroff et al., and elsewhere we have reviewed this literature in more detail (Thornton & Fernandez-Duque, 2002; see also Simons & Silverman, in press). Here, we briefly mention four other studies, the first of which was also discussed by Mitroff et al. (the other three appeared after publication of their critique).

Smilek et al. (2000) combined a flicker paradigm with a standard visual search task to show that large changes lead to shallower search slopes than do smaller changes. This finding was originally interpreted as suggesting that attention was guided to the location of large changes more effectively than to the location of the small changes. Mitroff et al. (2002) have challenged this original interpretation of the finding on the grounds that slope differences may have stemmed from an increased difficulty in detecting small changes. According to this alternative interpretation, the size of change is not a factor in guiding attention. Instead, attention is allocated to every item with equal probability, but attended items with a small change are more likely to be missed than attended items with a large change. Such misses lead to disproportionately long search times, and in this way they disproportionately increase the duration of search for small changes.[2]

Although this alternative explanation certainly identifies a factor that could affect performance in this task, experiments conducted after the publication of Mitroff et al.'s (2002) critique would seem to favor Smilek et al.'s (2000) original interpretation. One of these follow-up experiments revealed that the bias of large changes toward shallower search is dependent on the duration of the mask. When the interval duration is increased from 80 ms to 300 ms, the bias toward large changes disappears. An account based on small changes being missed more often than large changes would need to pose different target-detection mechanisms for short and long intervals. Indeed, some researchers have adopted this position, arguing that longer durations favor object-based processing whereas shorter durations favor feature-based processing (Richards, Jolicœur, Stolz, & Vogel-Sprott, 2001). Another experiment reveals that the modulation of search slopes by the size of the change also depends on spatial proximity. When the displays are presented side by side, increasing the spatial disparity, large changes do not benefit relative to small changes (Smilek, Eastwood, & Merikle, 2002). Again, this finding poses a challenge to interpretations based on small changes being missed more often than large changes because it is not clear why such mechanisms would apply only to the flicker paradigm and not to the side-by-side paradigm.

In another flicker paradigm, participants were asked to report a change in a complex scene (Hollingworth & Henderson, 2002). A 1° rotational change was incrementally introduced every flicker, until the change was reported. It took more than 30 flickers (i.e., 30° of rotation) for half of the observers to report the change. This difficulty in detecting the change was not due to a failure to engage in a comparison of successive flickers: Control experiments showed that changes of a smaller size (15°) were readily detected if they occurred in a single step. A follow-up experiment also ruled out the possibility that the initial frame was encoded but was not updated until the change was explicitly detected. In this follow-up experiment, gradual increases were introduced up to a rotation of 20°, the point at which the image reversed to the initial scene in a single step. Most observers readily noticed this single-step reversal. This last result suggests that the rotational changes were

incorporated into visual memory despite the fact that observers failed to detect those incremental changes. In other words, there was an implicit update in the representation of the scene as displayed in the 20° rotational image.

One final example of recent findings for implicit representation of change comes from a perceptual-motor synchronization task (Repp, 2001). In this task, participants tapped a key in synchrony with a sequence of auditory tones. In some trials, a change in tempo was introduced at a certain point of the sequence. Participants continued tapping after the end of the sequence, following which they reported whether they had noticed a change. For sequences in which participants reported being aware of the change, they made both period corrections and phase corrections. Most important, phase corrections occurred even in trials in which the participants failed to report the change, suggesting that awareness of change is not needed for such correction.

These behavioral studies and their findings give credibility to the claim for implicit representation of change. This argument is given further strength by converging evidence from other methodologies. Already, there are a number of findings using eye movements (Hayhoe et al., 1998; Henderson & Hollingworth, 1999; Ryan et al., 2000) and brain imaging (Beck et al., 2001; Fernandez-Duque et al., 2003) that support the existence of implicit representation of change (Thornton & Fernandez-Duque, 2002).

Mitroff et al.'s (2002) critique has raised a number of important issues relating to the implicit change detection debate. At the methodological level, the authors raised specific concerns about some of the original empirical findings. In this comment, we have addressed those concerns and, in so doing, have provided further evidence for the existence of implicit change detection. At a more general level, Mitroff et al. have correctly brought attention to some of the lingering difficulties that are faced whenever implicit processes are studied. Chief among these concerns is the reliance on subjective measures of awareness (e.g., "I don't see a change").

Such subjective measures of awareness may reflect true absence of explicit knowledge, but they may also indicate biases introduced by the experimental instructions, the response criteria, or the observers' beliefs. In other words, subjective measures of awareness fail to satisfy the *exhaustiveness* criterion—the assumption that a measure can account for each and every piece of consciously perceived information available to an observer (Merikle & Reingold, 1992). In failing to satisfy this assumption, measures of implicit performance by a subjective criterion cannot rule out

---

[2] In support of their claim, Mitroff et al. (2002) reported that the search slope benefit for large changes occurs even when the location of the change is constantly varied from cycle to cycle (e.g., original scene → scene with change in bottom left → original scene → scene with change in top left). These results are difficult to interpret because the experimental design used by Mitroff et al. required a full stop after every cycle for observers to report whether they had noticed a change. It seems quite likely that implicit accumulation of change information would be disrupted by such an interruption. As a reply to this concern, Mitroff et al. presented evidence that the slope differences occur even in a full-stop condition in which the change remains in the same location across cycles. However, that finding is relevant only if one is willing to make the unlikely assumption that the implicit accumulation of information would survive such a strong disruption.

contamination by trials in which the observer was actually aware of the stimulus but failed to report it (i.e., the contamination hypothesis). Put in other words, it may be that subjective report is not as sensitive as other methods of assessing awareness.

One solution to this problem is to adopt an *objective* criterion of awareness, according to which observers are said to be unaware of something if they perform at chance levels in a forced-choice task (C. W. Eriksen, 1960; Holender, 1986). For example, while trying to assess the possibility that subliminally presented words could influence perception, Cheesman and Merikle (1984) first engaged observers in a forced-choice naming task in which display duration was varied until observers were at chance in naming written words. After this adjustment phase, the words were presented using the same display durations in a second task designed to tap noncon-scious processing. Previous studies using a subjective measure of awareness (i.e., "Did you see a word?") had shown implicit effects using such a task (Marcel, 1983). However, when paired with Cheesman and Merikle's objective measure, no such implicit effects were obtained.

This approach could be adapted to studies of implicit change processing. For example, a researcher could choose a direct per-ceptual task in which an observer would be asked to report whether a change was displayed or not. An inability to discriminate be-tween catch trials and change trials, as measured by d', would provide objective evidence that observers were unaware of the change. If under those conditions, observers were capable of selecting the location of change at better-than-chance, it would constitute evidence for implicit location of change. However, a problem with such an approach is the inability to rule out the possibility that the so-called implicit effect was an artifact of the localization task being a more sensitive measure of aware pro-cesses than is the detection task. In other words, there is no a priori way to decide which is the most sensitive measure of aware performance. Another weakness of such an approach is that it assumes that implicit processes cannot contribute to correct per-formance. In other words, it assumes that a correct answer must always be accompanied by awareness, an assumption that is clearly questionable.

Mitroff et al.'s (2002) proposed that objective measures should be adopted in the study of change blindness (see also Simons & Silverman, in press). They correctly emphasized that using an objective criterion for awareness has the benefit of providing indisputable protection against the contamination hypothesis. In other words, that approach satisfies the exhaustiveness criterion referred to above. However, Mitroff et al. failed to acknowledge that adopting an objective measure of awareness denies the pos-sibility that implicit processes could ever influence performance. Essentially, in ruling out any contribution from consciousness, such measures may also exclude the very thing that one is trying to measure. Put in more formal terms, adopting an objective measure of awareness violates the *exclusiveness* criterion, which states that a method must only measure aware performance, never running the risk of misclassifying unaware trials as products of awareness (Reingold & Merikle, 1990).

Satisfying both the exhaustiveness and the exclusiveness crite-rion requires a measure that is both a very sensitive (exhaustive-ness criterion) and a very selective (exclusiveness criterion) mea-sure of awareness, targeting all aware trials but only aware trials. These criteria are so strict that the pursuit of a method that would

satisfy them both would not be fruitful. This difficulty led Rein-gold and Merikle (1990) to suggest a more tractable solution, the establishment of *qualitative* differences between aware and un-aware performance. Unlike quantitative differences, qualitative differences are not easily explained by a contamination between aware and unaware trials. Such qualitative differences can be expressed in a variety of ways. For example, qualitative differ-ences may be expressed as a dissociation in the way that two dependent variables are affected by the level of awareness. One such possibility would be for aware trials to influence RTs, whereas unaware trials would influence accuracy.

A qualitative difference based on level of awareness may also be revealed as different effects on two cognitive operations. For example, it is conceivable that unaware trials would facilitate localization of a change but not the cuing of attention to such a location, whereas aware trials would facilitate both. Another dis-sociation may be revealed in the way that unaware processes affect two different populations. For example, unaware trials may have an effect on normal observers but not on patients, whereas aware trials affect both groups. Finally, dissociations may also be re-vealed by the qualitatively different effect that a third variable has on aware and unaware trials. For example, increased duration of interval between changes may affect only unaware trials, not aware ones. The strength of this dissociation approach is that it permits the adoption of criteria that are conservative enough to minimize false discoveries while at the same time ensuring that genuinely new phenomena are not impossible to establish.

How does the current evidence for implicit processing of change fare against this background? In our past research, we have dem-onstrated that accuracy and RT measures can be differentially affected by implicit and explicit detection of change. Specifically, in the experiments of Thornton & Fernandez-Duque (2000), we found that conflicts between an explicitly detected change and a subsequent probe led to an increase in error rates and a slowing of RT. When similar changes were present but not reported, only accuracy levels were affected. Using a similar design, Experiments 3 and 4 of the current study showed the same pattern of dissoci-ation. In Experiment 2, however, the congruency effect for un-aware trials was evident in both RTs and accuracy rates. None-theless, a more detailed analysis showed that the same dissociation was present, though only for trials in which the probe occurred close to the change.

In another study, we explored whether unaware localization was mediated by a cuing of attention (Fernandez-Duque & Thornton, 2000). Contrary to our initial hypothesis, we found that attention was cued to the location of change only when observers were aware of the change, not when the change was unaware. Thus, awareness of change had qualitatively different effects on local-ization and attentional cuing. This result is difficult to reconcile with a contamination hypothesis unless the further assumption is made that localization measures (based on accuracy rates) were more sensitive than attentional cuing measures (based on RTs). Further evidence for a dissociation comes from an experiment in which the response criterion for awareness of change was manip-ulated via instructions (Fernandez-Duque & Thornton, 2000, Ex-periment 1). Adopting a more conservative criterion for awareness significantly increased the accuracy with which aware changes were localized, but it did not modify the accuracy of unaware changes.

In the current experiments, we sought further evidence for qualitative differences between aware and unaware performance. In Experiment 1, we asked whether performance for unaware trials would be worse than chance, under the assumption that item selection would still be biased toward the change even though task instructions required selection of the nonchanged item. Although the pattern of results was in this predicted direction, it failed to reach statistical significance. Future studies will have to explore whether this failure was due to a lack of statistical power or a real absence of dissociation. In Experiment 2, we explored the qualitative differences between aware and unaware trials in the way that awareness of change interacted with the probe proximity. Probes that were close to the changed item led to an increase in RT and a subsequent reduction in error, but only when observers were aware of the change, a pattern of results that cannot be explained by simple contamination.

Dissociations between aware and unaware performance have also been reported by other researchers. For example, performance in a perceptual-motor synchronization task has been shown to be qualitatively different depending on awareness of change, with aware changes leading to corrections in period and phase but unaware changes only triggering phase corrections (Repp, 2001). Eye movement studies have revealed a qualitative difference in the influence of unaware changes on normal and amnesic observers (Ryan et al., 2000). When seeing a complex scene for a second time, normal observers reveal increased viewing times of the changed region even when failing to report the change. In contrast, the implicit effect is absent in amnesic patients.

Concluding Remarks

Change detection is a powerful and a flexible tool for exploring visual processing. In our previous work, we have suggested that explicit reports underestimate the ability of the visual system to represent change (Fernandez-Duque & Thornton, 2000; Thornton & Fernandez-Duque, 2000). Consistent with this claim, our studies have demonstrated that undetected changes can still influence behavior. These findings, as well as data from several other laboratories (Hayhoe et al., 1998; Hollingworth, 2003; Hollingworth & Henderson, 2002; Ryan et al., 2000), provide clear evidence that visual representations are richer and more stable than initial studies of change blindness may have led us to believe (cf. O'Regan, 1992).

Mitroff et al. (2002) have suggested that those findings could be explained by explicit mechanisms without invoking an implicit process. In the four experiments reported in this article, we have shown that such explicit mechanisms do not provide a viable alternative to the notion of implicit detection of change. Rather, the current findings, together with a growing body of converging evidence from other methodologies (e.g., eye movement analysis, motor control and brain imaging studies), suggest that implicit processing of change is more than an alluring theoretical possibility, it is a solid, empirical fact. These are not the songs of sirens, these are the sounds of science.

References

Beck, D. M., Rees, G., Frith, C. D., & Lavie, N. (2001). Neural correlates of change detection and change blindness. *Nature Neuroscience, 4,* 645–650.

Castiello, U., Paulignan, Y., & Jeannerod, M. (1991). Temporal dissociation of motor responses and subjective awareness. *Brain, 114,* 2639–2655.

Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics, 36,* 387–395.

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

Cowey, A., & Stoerig, P. (1991). Reflections on blindsight. In A. D. Milner & M. D. Rugg (Eds.), *The neuropsychology of consciousness* (pp. 11–39). London: Academic Press.

DeJong, R., Liang, C.-C., & Lauber, E. (1994). Conditional and unconditional automaticity: A dual-process model of effects of spatial stimulus–response correspondence. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 731–750.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics, 16,* 143–149.

Eriksen, C. W. (1960). Discrimination and learning without awareness: A methodological survey and evaluation. *Psychological Review, 67,* 279–300.

Fernandez-Duque, D., Grossi, G., Thornton, I. M., & Neville, H. J. (2003). Representation of change: Separate electrophysiological markers of attention, awareness, and implicit processing. *Journal of Cognitive Neuroscience, 15,* 491–507.

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,* 324–344.

Goodale, M. A. (1996). One visual experience, many visual systems. In T. Inui and J. L. McClelland (Eds.), *Attention and performance XVI: Information integration in perception and communication* (pp. 369–393). Cambridge, MA: MIT Press.

Hayhoe, M., Bensinger, D., & Ballard, D. H. (1998). Task constraints in visual working memory. *Vision Research, 38,* 125–137.

Henderson, J. M., & Hollingworth, A. (1999). The role of fixation position in detecting scene changes across saccades. *Psychological Science, 10,* 438–443.

Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behavioral and Brain Sciences, 9,* 1–23.

Hollingworth, A. (2003). Failures of retrieval and comparison constrain change detection in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance, 29,* 388–403.

Hollingworth, A., & Henderson, J. M. (2002, May). *Sustained insensitivity to incremental scene rotation: A dissociation between explicit change detection and visual memory.* Talk presented at the Second Annual Meeting of the Visual Sciences Society, Sarasota, FL.

Lambert, A., Naikar, N., McLachlan, K., & Aitken, V. (1999). A new component of visual orienting: Implicit effects on peripheral information and subthreshold cues on covert attention. *Journal of Experimental Psychology: Human Perception & Performance, 25,* 321–340.

Lu, C.-H., & Proctor, R. W. (1994). Processing of an irrelevant location as a function of the relevant stimulus dimension. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 286–298.

Lu, C.-H., & Proctor, R. W. (1995). The influence of irrelevant location information on performance: A review of the Simon and spatial Stroop effects. *Psychonomic Bulletin & Review, 2,* 174–207.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin, 109,* 163–203.

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

Merikle, P. M., & Reingold, E. M. (1992). Measuring unconscious perceptual processes. In R. F. Bornstein & T. S. Pitman (Eds.), *Perception*

without awareness: Cognitive, clinical, and social perspectives (pp. 55–80). New York: Guilford Press.

Milner, B., Corkin, S., & Teuber, H. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H. M. *Neuropsychologia, 6,* 215–234.

Mitroff, S. R., & Simons, D. J. (2002). Changes are not localized before they are explicitly detected. *Visual Cognition, 9,* 937–968.

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

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. (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology, 46,* 461–488.

Pelli, D. G., & Zhang, L. (1991). Accurate control of contrast on microcomputer displays. *Vision Research, 31,* 1337–1360.

Reingold, E. M., & Merikle, P. M. (1990). On the inter-relatedness of theory and measurement on the study of unconscious processes. *Mind and Language, 5,* 9–28.

Rensink, R. A. (1990). Toolbox-based routines for Macintosh timing and display. *Behavior Research Methods, Instruments, & Computers, 22,* 105–117.

Rensink, R. A. (2000). Visual search for change: A probe into the nature of attentional processing. *Visual Cognition, 7,* 345–376.

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

Rensink, R. A. (in press). Visual sensing without seeing. *Psychological Science.*

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.

Repp, B. H. (2001). Phase correction, phase resetting, and phase shifts after subliminal timing perturbations in sensorimotor synchronization. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 600–621.

Richards, E., Jolicœur, P., Stolz, J. A., & Vogel-Sprott, M. (2001, November). *The transition from feature-based to object-based processing.*

Poster session presented at the 42nd Annual Meeting of the Psychonomic Society, Orlando, FL.

Ryan, J. D., Althoff, R. R., Whitlow, S., & Cohen, N. J. (2000). Amnesia is a deficit in relational memory. *Psychological Science, 11,* 454–461.

Schacter, D. L. (1995). Implicit memory: A new frontier for cognitive neuroscience. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 815–824). Cambridge, MA: MIT Press.

Simon, J. R., & Small, A. M., Jr. (1969). Processing auditory information: Interference from an auditory cue. *Journal of Applied Psychology, 53,* 433–435.

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

Simons, D. J., & Silverman, M. (in press). Neural and behavioral measures of change detection. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences.* Cambridge, MA: MIT Press.

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

Smilek, D., Eastwood, J. D., & Merikle, P. M. (2002). *Further evidence that unattended changes guide attention.* Unpublished manuscript.

Steinman, S. B., & Nawrot, M. (1992). Real-time color-frame animation for visual psychophysics on the Macintosh computer. *Behavior Research Methods, Instruments, & Computers, 24,* 439–452.

Stroop, J. R. (1992). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology: General, 121,* 15–23. (Original work published 1935)

Thornton, I. M., & Fernandez-Duque, D. (2000). An implicit measure of undetected change. *Spatial Vision, 14,* 21–44.

Thornton, I. M., & Fernandez-Duque, D. (2002). Converging evidence for the detection of change without awareness. In J. Hyönä, D. P. Munoz, W. Heide, and R. Radach (Eds.), *The brain's eye: Neurobiological and clinical aspects of oculomotor research: Progress in brain research* (Vol. 140, pp. 99–118). Amsterdam: Elsevier Science.

Watanabe, K. (2003). Differential effect of distractor timing on localizing versus identifying visual changes. *Cognition, 88,* 243–257.