Spatial Vision, Vol. 14, No. 1, pp. 21–44 (2000)

An implicit measure of undetected change

IAN M. THORNTON[1,*] and DIEGO FERNANDEZ-DUQUE[2,*]

[1] *Cambridge Basic Research, Nissan Technical Center North America, Inc., Cambridge, MA, USA*
[2] *Psychology Department, University of Oregon, Eugene, OR, USA*

Received 3 September 1999; revised 1 February 2000; accepted 12 May 2000

Abstract—Several paradigms (e.g. change blindness, inattentional blindness, transsaccadic integration) indicate that observers are often very poor at reporting changes to their visual environment. Such evidence has been used to suggest that the spatio-temporal coherence needed to represent change can only occur in the presence of focused attention. However, those studies almost always rely on explicit reports. It remains a possibility that the visual system can implicitly detect change, but that in the absence of focused attention, the change does not reach awareness and consequently is not reported. To test this possibility, we used a simple change detection paradigm coupled with a speeded orientation discrimination task. Even when observers reported being unaware of a change in an item's orientation, its final orientation effectively biased their response in the orientation discrimination task. Both in aware and unaware trials, errors were most frequent when the changed item and the probe had incongruent orientations. These results demonstrate that the *nature* of the change can be represented in the absence of awareness.

1. INTRODUCTION

Visual search has proven to be a very useful technique for exploring human perception (see Wolfe, 1998, for a review). In a typical visual search task, observers are asked to detect a target within a field of distracting items. The target can differ from the distractors along a single feature dimension, such as color or orientation (e.g. find the red circle among blue circles) or in the way basic features are combined (e.g. find the red circle among red squares and blue circles). Speed and accuracy

---

[*] Correspondence may be addressed to either author at the following current addresses: Diego Fernandez-Duque, Rotman Research Institute, Baycrest Centre for Geriatric Care, 3560 Bathurst Street, Toronto, ON, M6A 1E6, Canada (diego@rotman-baycrest.on.ca), and Ian M. Thornton, Max Planck Institute for Biological Cybernetics, Spemannstrasse 38, 72076 Tübingen, Germany (ian.thornton@tuebingen.mpg.de).

of search can be measured as a function of set size (total number of items) and/or display duration.

Findings from visual search tasks have been very influential, producing ideas about the basic building blocks of vision, the overall organization of visual information processing and, in particular, the relationship between vision and attention (e.g. Treisman and Gelade 1980; Nakayama and Silverman, 1986; Duncan and Humphreys, 1989; Enns and Rensink, 1990; Wolfe, 1994, 1998). While some researchers have recently urged caution regarding the interpretation of findings from traditional search experiments (e.g. Nakayama and Joseph, 1998), it seems likely that the technique itself will remain an important tool for helping scientists understand human vision.

A recent variant of this basic technique involves visual search for *change*. Here, targets are defined not in terms of static features, but rather in terms of dynamic spatio-temporal patterns. For example, instead of searching for a blue rectangle, observers might be required to find the rectangle that is changing from red to blue among a field of static red and blue rectangles. Search for change is theoretically very interesting as it involves the dimensions of both space and time. As a number of researchers have pointed out, while we live in a world where time is a fundamental dimension, most research aimed at understanding mental representation has focused predominantly on spatial, and moreover, static patterns (e.g. Jones, 1976; Freyd, 1987). This is certainly true for the vast majority of visual search studies.

However, Rensink (2000a) has recently used simple stimuli, such as those shown in Fig. 1, to demonstrate that search for change is formally very similar to traditional static search paradigms. Examination of reaction times during search for change can yield estimates of search speed and search selectivity (i.e. the ability to filter out irrelevant dimensions), just as with static visual search. More importantly, search for change can shed new light on dynamic aspects of human vision: that is, it can provide estimates of attentional capacity for representations that are maintained over time, indicate how this capacity might vary with search dimension (e.g. changes in color versus changes in orientation), and provide information about the role of memory systems during dynamic visual search (cf. Horowitz and Wolfe, 1998). Visual search for change has also proved very useful for probing individual differences in the dynamic allocation of attention (Harp and Rensink, 1999).

While visual search for change has great promise as a general tool for exploring dynamic aspects of human vision, there is one finding that has generated the most interest. That finding is the extreme difficulty observers often have in detecting that anything is changing at all (Rensink *et al.*, 1997; Simons and Levin, 1997). When display durations are brief, detection rates are generally very low. Given ample time, observers often take many seconds or even minutes to locate the changing item. While such ‘change blindness’ could have been predicted from earlier work on visual short term memory (Phillips, 1974; Pashler, 1988), saccadic integration (Irwin, 1991) or even frustrating games such as ‘spot the difference’, it

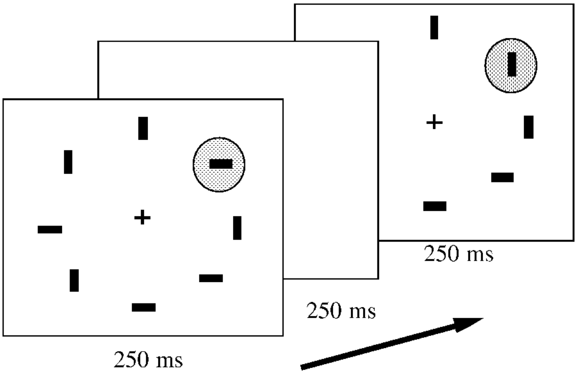

Figure 1. Schematic representation of the layout and timeline for stimuli used in studies of visual search for change (Rensink, 2000a). For our experiments, we used stimuli such as these in the initial portion of each trial (for more details, see General Methods and Fig. 3).

was the development of the so-called ‘flicker paradigm’ that firmly established the phenomenon as an important area for further research (Rensink *et al.*, 1997).

The flicker paradigm, which is illustrated in Fig. 2, is essentially a version of the search for change technique mentioned above. Here, however, complex scenes are used instead of simple stimuli. In the flicker paradigm, two views of a complex scene are separated by a blank masking field and are alternated in the sequence — scene 1, mask, scene 2, mask, scene 1, mask, scene 2, and so on. The two scenes differ from one another only with respect to a single changing item or scene location. Once the changing item has been detected it is clearly visible and often appears very ‘obvious’. The crucial factor in making the change hard to detect is the masking field. This global transient event disrupts the local transients that usually accompany change. In addition to blank fields, other types of distracting separators have been used, including blinks (e.g. O’Regan *et al.*, 2000), saccades (e.g. Bridgeman *et al.*, 1975; Grimes, 1996), movie cuts (Levin and Simons, 1997) and multiple small masking elements called ‘mud splashes’ (O’Regan *et al.*, 1999).

Subsequent work has shown that detection of change is also difficult during virtual reality simulations (e.g. Wallis and Bülthoff, 1999), dynamic animation sequences (e.g. Scholl and Pylyshyn, 1999) and even real world, face to face, interactions (Simons and Levin, 1998). Flicker techniques using arrays of simple stimuli (Rensink, 2000a) or common objects (Zelinsky, 1997), as well as side-by-side presentations of images (e.g. Pomplun *et al.*, 1999) also make the detection of change very difficult. The wide range of situations in which change blindness can be observed — apparently any circumstances that remove or reduce change-related transients — suggests that these studies may be tapping into something quite fundamental to human vision.

According to many researchers, that ‘something’ is attention. More specifically, the generally accepted explanation for why we are so bad at detecting change in the aforementioned situations is that, in the absence of motion transients, attention is not

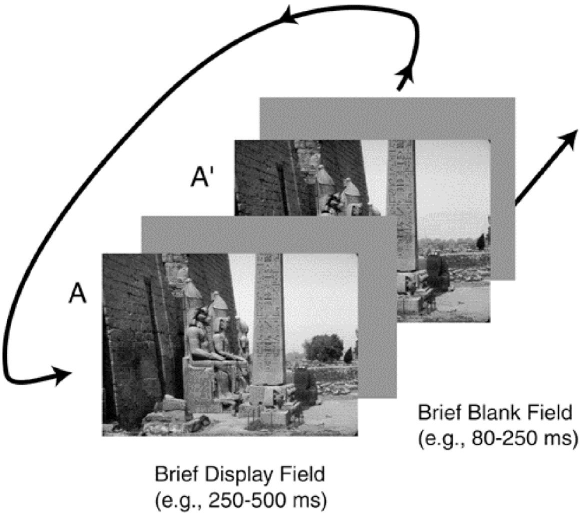

Figure 2. Schematic representation of the layout and timeline for the stimuli used in the flicker paradigm (Rensink *et al.*, 1997). Two views of a complex scene are separated by a blank masking field and are alternated in the sequence — scene 1, mask, scene 2, mask, scene 1, mask, scene 2, and so on. These two scenes differ from one another only with respect to a single changing item or scene location. In this example, taken from the Cambridge Basic Research database, a large tree on the right of the screen disappears in the second image.

being drawn to the location of change (Simons, 1996; Rensink *et al.*, 1997; Rensink, 2000b). Focused attention is seen as the 'glue' that integrates representations across space and time (Treisman and Gelade, 1980), giving rise to coherent 'object files' that can be updated as a represented stimulus changes (Kahneman *et al.*, 1992). Therefore, in the absence of attention, successive views of an object will replace each other rather than being integrated into a single coherent representation. According to this view, change can only be represented — and therefore detected — in the presence of focused attention.

Indeed, if attention is directed to the location of change, either by motion transients, object saliency (Rensink *et al.*, 1997; Shore and Klein, 2000) semantic cues (Rensink e*t al.*, 1997), or exogenous cues (Scholl, 2000), detection is fast and efficient. In the absence of such cues, search is slow, error prone and effortful. In situations lacking external cues, search is guided by the volitional control of the observer. That is, items in the display successively become the focus of attention, with the search moving to a new item after the observer decides that the currently held item is not changing over time (and therefore is not the target). Thus, locating an item that is changing is much like locating a conjunction target in traditional visual search tasks, the difference being that in the latter case the target is defined only by its spatial properties, whereas in the former, the target is defined by its spatio-temporal properties.

The idea that spatio-temporal coherence can only exist in the presence of focused attention has important implications for our general view of perception, because it suggests that, at a given moment in time, there will be large portions of the visible world in which changes cannot be represented. Similar claims have previously been made for a more general lack of object structure in the absence of attention (e.g. Treisman and Gelade, 1980; Wolfe and Bennett, 1997). This lack of spatial and temporal integration has led a number of researchers to suggest that our subjective impression of a detailed, stable representation of the physical world is little more than a 'Grand Illusion' (O'Regan, 1992). It is argued that we only have a stable detailed representation of the visual area we are attending to at any given time. We fail to notice that most of our visual world lacks detail and coherence because as soon as we 'look' at a new region of space, we bring that region into the focus of attention (Rensink, 2000b).

Studies of change blindness are considered to be important evidence in support of perception as a 'Grand Illusion'. However, these studies of change blindness, as with earlier studies that invoke the Grand Illusion, typically require observers to make explicit reports. Thus, while they may provide direct evidence about perceptual *awareness*, such findings are less informative about perceptual *representation.* Indeed, there is a great deal of work on perception without awareness suggesting that conscious report is not always a good indication of visual representation (e.g. Marcel, 1983a, b; Graves and Jones, 1992; Kolb and Braun, 1995; Luck *et al.*, 1996; McCormick, 1997; Moore and Egeth, 1997; Bar and Biederman, 1998; Chen, 1998; Dehaene *et al.*, 1998; Mack and Rock, 1998).

Recently, we used a simplified change blindness paradigm coupled with a two-alternative forced choice (2AFC) selection task to demonstrate that typical studies of change blindness underestimate the representation of change in the visual system (Fernandez-Duque and Thornton, 2000). Even when observers report seeing no change, they are better than chance at selecting the changed item if forced to make a choice. These results suggest that the visual system is capable of representing the location of a change even in the absence of awareness. Interestingly, this implicit localization of change does not appear to be mediated by a reallocation of attention. That is, in control experiments, we demonstrated that undetected changes were ineffective at reorienting attention to the location of change.

The purpose of the current studies was to further explore the representation of change in the absence of awareness. More specifically, our goal was to extend our findings on implicit localization of change (Fernandez-Duque and Thornton, 2000), by exploring whether the *nature* of a change can also be implicitly represented. To achieve this, we combined a simplified change blindness paradigm with a speeded orientation discrimination task. Using the logic of priming studies, this combination allows us to explore the influence that a changing item might have on a subsequent horizontal/vertical decision. For instance, following an undetected change from vertical to horizontal, will responses to a vertical probe be facilitated, inhibited, or unaffected? Modulation during unaware trials would suggest that the specific nature of the change can be represented in the absence of awareness.

2. GENERAL METHOD

Both experiments reported in this paper use the same basic design illustrated in Fig. 3. These displays combined a simplified change detection paradigm with a speeded orientation discrimination task. On each trial a simple change detection display began with the presentation of a ring of 8 rectangles arranged around a central fixation cross. The ring arrangement, first introduced by Eriksen and Collins (1969), ensures that all items are equidistant from fixation. This initial display remained visible for 250 ms, was replaced by a blank 250 ms inter-stimulus interval (ISI) and then the whole ring reappeared. Observers were instructed to fixate at the center of the ring, but to remain alert for a possible change of orientation in one of the rectangles. The complete ring subtended 4.6 deg visual angle from the fixation point. Each rectangle measured 10 × 30 pixels, which subtended 0.46 × 1.38 deg visual angle. Rectangles were drawn in black on a medium gray background which was also the color of the blank ISI.

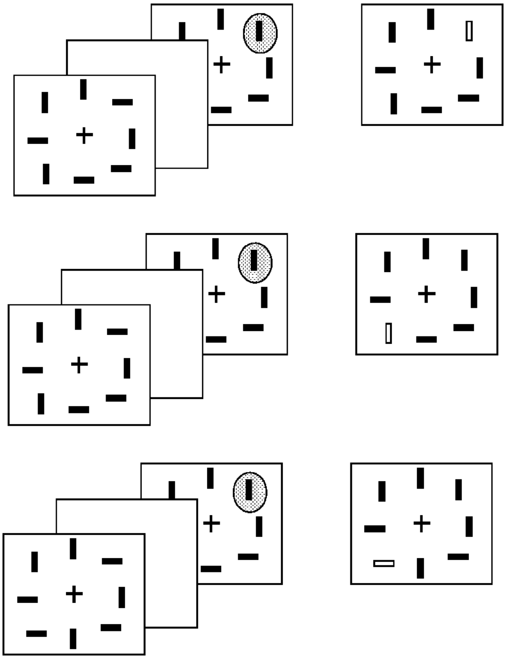

Figure 3. The four experiments in the current study combine a simplified change detection paradigm with a speeded orientation discrimination task. In the initial portion of each trial, the two rectangle frames were identical except for a single object which changes orientation during the blank interstimulus interval (ISI). Change was equiprobable at any location and here involves the rectangle located between 12 and 3 o'clock. Following this initial portion of the trial, after a zero-second ISI, a response display with a target item appeared. Observers made a speeded response using one of two marked keys. The position of the target could overlap either with the change location (valid trial, top) or the distractor location (invalid trials, middle and bottom). In invalid trials, the target could have the same orientation as the changed item (congruent invalid trial, middle) or the opposite orientation (incongruent invalid trial, bottom).

On ‘change’ trials (66%), one of the items in the ring changed orientation between the first and second display by rotating 90 deg around its center point. Each item in the ring had an equal probability of being the change target on a given trial. The remaining 33% of trials were catch trials in which no object changed orientation between the first and second display. These trials were included to assess the accuracy with which observers were subjectively reporting the presence of change.

Immediately following the change detection display on each trial, one item in the ring changed color from black to light gray. This color change, which occurred 250 ms after the second display of the ring, indicated the target for the orientation task. Observers were instructed to respond as rapidly as possible to the orientation of this highlighted item by pressing one of two indicated keys. Feedback on the accuracy and speed of this response was always provided in the form of audible beeps.

Following the speeded orientation response, observers indicated whether they had been aware of any change during the flicker part of the display. In Experiment 1 this awareness response was collected using a go/no-go protocol. That is, observers were required to make a keypress if they were aware of a change (the ‘go’ response) and did nothing if they were unaware (the ‘no-go’ response). In Experiment 2, a two-alternative forced choice response (2AFC) was used to indicate awareness. No feedback was provided on the accuracy of awareness responses in either experiment.

The most important manipulation in these experiments concerned the relationship between the changed item and the subsequent target for the orientation discrimination task. This relationship is illustrated in Fig. 3. On one half of the change trials the target item was identical to the item that changed: that is, it occurred in the same location and always had the same orientation. These will be referred to as *valid* trials (Fig. 3, top). For the remaining half of the change trials the target item was located diametrically opposite to the changed item. These will be referred to as *invalid* trials (Fig. 3, middle and bottom). Invalid trials were further subdivided depending on the relation between the orientation of the probe and the final orientation of the changed item. In half of the invalid trials the target item of the orientation discrimination task was the same orientation as the changed item. These will be referred to as *congruent* invalid trials (Fig. 3, middle). The remaining half of the invalid trials contained a target item with the opposite orientation to the changed item. These will be referred to as *incongruent* invalid trials (Fig. 3, bottom). (Note that there are no incongruent valid trials, as this would have involved a second change of orientation.)

3. EXPERIMENT 1

The purpose of Experiment 1 was to explore the impact of a changing item on a subsequent orientation discrimination response as a function of whether the change was consciously perceived or not. This design allowed us to explore both *validity* effects and *congruency* effects.

Validity effects are thought to arise due to a shift in the allocation of attention. If attention is ‘cued’ to a particular location, subsequent processing at that location will be facilitated. This facilitation typically manifests itself as a reaction time difference between cued and uncued objects. When a change is consciously detected, it should act as an attentional cue. Therefore, a target occurring at a valid location should be responded to more quickly than a target at an invalid location.

Congruency effects are thought to reflect some form of match/mismatch between a primed representation and a subsequent target. For example, when a horizontal item changes orientation to become vertical, such ‘verticality’ is primed so that a ‘vertical’ response to a subsequent target will proceed faster and with fewer errors than a ‘horizontal’ response. This prediction is based on the assumption that the change, at least when consciously detected, will make the properties of the changed item, such as its orientation, more salient.

While we can predict the response patterns following conscious detection of change with relative ease, our primary interest is in what happens after changes that are not consciously perceived. If change blindness reflects a failure to represent change outside the current focus of attention, then speed and accuracy of orientation responses should be unaffected by the presence of an undetected change. However, if undetected changes are represented at some level in the visual system, then validity and congruency effects may occur even when observers report being unaware.

3.1. Method

3.1.1. Observers. Eighteen University of Oregon students took part in this study for partial course credit. All observers had normal or corrected to normal vision, were right handed and were naive as to the purpose of the study.

3.1.2. Observer selection. Based on our previous work, we imposed a number of observer selection criteria. The most important of these involved the overall level of awareness for changing items. In order to calculate accurate estimates of behavior, each observer needed to provide a minimum number of aware and unaware responses. Therefore, any observer who responded aware or unaware on more than 80% of trials was excluded from the analysis. We also excluded any observer whose overall accuracy in the orientation discrimination task fell more than 2 standard deviations below the sample mean, as it seems likely that such a pattern reflects either miscommunication of instructions or uncooperative behavior. In Experiment 1, no observers were excluded.

3.1.3. Equipment. Stimuli display and response collection were carried out on a Power Macintosh 7200 attached to a standard 15″ (27 × 20 cm) RGB monitor with a frame rate of 75 Hz and a screen resolution of 832 × 624 pixels. Software was custom written using Think C version 7.0. Many of the routines were based on work by Steinman and Narwot (1992), Pelli and Zhang (1991) and Rensink (1990).

3.1.4. Stimuli. Stimuli were identical to those described in the general method (see Fig. 3). The probe remained on the screen until a response was made.

3.1.5. Procedure. Observers were asked to fixate on a central cross during the first two rectangle displays, but were allowed to move their eyes at all other times. They were informed that on some trials, the orientation of one of the rectangles would change and that they should try to notice such a change.

The basic trial structure was illustrated and the nature of the two responses was explained. Orientation responses were collected using a 2AFC response and awareness was indicated via a go/no-go decision. The go/no-go technique was used to reduce the response mapping load on observers, as they were already required to make a speeded 2AFC response to the orientation task. Observers completed an initial training block of 10 trials to ensure they understood this basic structure.

Observers were then instructed that the horizontal/vertical decision should be a speeded response and that audible feedback would be given each time they were too slow or made an incorrect orientation response. Observers then completed a minimum of 10 practice trials that included this feedback. Practice was terminated when the observer felt comfortable making the speeded responses.

No feedback was given on the awareness response. Observers were instructed to adopt a liberal criterion for awareness. That is, they were instructed to report awareness of change if they either ‘saw’ the change or even if they simply thought or felt that something had changed. The instruction to adopt a liberal criterion for awareness aimed to minimize the contamination of ‘unaware’ responses by mislabeled aware responses.

After the practice phase, observers completed 384 experimental trials. The trials were divided in 4 blocks of 96 trials each. Observers were encouraged to take short breaks between each block.

3.2. Results

In accordance with previous studies of change blindness, observers were generally quite poor at detecting the changes present in these displays. In trials with a change, observers indicated being aware of it only 54% of the time. Importantly, changes were reported on only 8% of catch trials. The discrimination index, d_{L}, revealed that, despite the fact that observers missed many changes, they were able to discriminate better than chance between the change trials and the catch trials, $d_{\mathrm{L}} = 2.798$, $t(17) = 11.2$, $p < 0.0001$, $SEM = 0.25$, CI (2.27, 3.32).

In the current work, a hit was a trial in which there was a change and the subject reported being aware of the change. A false alarm (FA) was a catch trial (i.e. trial without a change) in which the subject reported seeing a change. We used d_{L} to index sensitivity. d_{L} is the logistic analogue of d' (Snodgrass and Corwin, 1988), sometimes referred to as $2\log(\alpha)$ (Macmillan and Creelman, 1996). Snodgrass and Corwin (1988) have shown that sensitivity indices based on signal detection theory with logistic distributions (i.e. d_{L}) yields equivalent results as indices based on signal detection theory with normal distributions (i.e. d'). We preferred d_L because it is easier to compute, and therefore errors in its calculation are minimized. d_L is computed as follows:

$$d_L = \ln\{[\mathrm{Hit}(1 - \mathrm{FA})]/[(1 - \mathrm{Hit})\mathrm{FA}]\}. \quad (1)$$

We used reaction times and error rates to explore the validity and the congruency effect, as well as possible modulations by awareness. We first describe the reaction time data, after which we describe the error data.

Reaction time data included trials in which the observer made a correct response to the orientation task. Trials on which an error was made were excluded from the RT analysis because they were insufficient in number. To examine the *validity* effect, we ran a 2 × 2 repeated measure ANOVA in which type of trial (valid, invalid congruent) was crossed with level of awareness of change (aware, unaware). There was a main effect of awareness, $F(1, 17) = 11.8$, $p < 0.003$, $MSE = 3957$. Reaction time was generally slower when observers were aware of the change than when the change was undetected (see Fig. 4). More importantly, there was a main effect of validity, $F(1, 17) = 67$, $p < 0.0001$, $MSE = 701$. Figure 4 shows that valid trials were faster than invalid trials. However, this validity effect interacted with level of awareness, $F(1, 17) = 9.8$, $p < 0.006$, $MSE = 1608$. Further analysis of this interaction revealed a strong 81 ms validity effect when observers were aware of the change, $t(17) = 5.8$, $p < 0.0001$, $SEM = 14$, CI (51, 110), and a smaller, though still reliable, 21 ms validity effect when observers were unaware of the change, $t(17) = 2.7$, $p < 0.014$, $SEM = 7.9$, CI (5, 38). The presence of a validity effect for unaware trials suggests that, even in the absence of awareness,

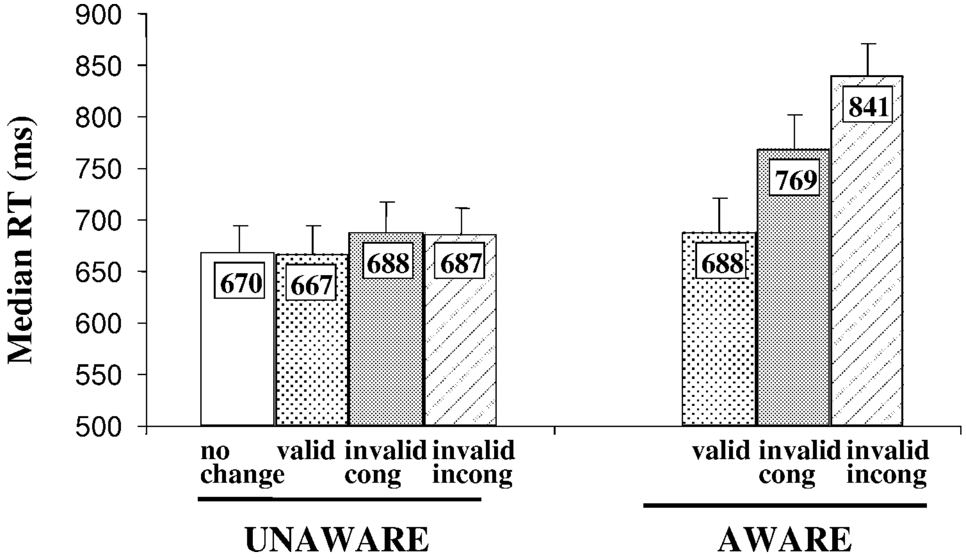

Figure 4. Median Reaction Times (in ms) for Experiment 1. Error bars indicate the Standard Error of the Mean.

some information about the location of the change was being represented and served as an orienting cue, drawing attention to the location of change.

To assess the *congruency* effect, reaction times to invalid congruent and invalid incongruent trials were compared using a 2 (congruent, incongruent) × 2 (aware, unaware) ANOVA. As with the validity effect, this analysis revealed a main effect of awareness, $F(1, 17) = 61$, $p < 0.0001$, $MSE = 4033$, aware trials were slower than unaware trials. More importantly, there was a congruency effect, $F(1, 17) = 8.4$, $p < 0.01$, $MSE = 2679$, invalid incongruent trials were slower than invalid congruent trials. However, this congruency effect interacted with the level of awareness, $F(1, 17) = 9.6$, $p < 0.006$, $MSE = 2519$. Further analysis of this interaction revealed that there was a strong 72 ms congruency effect when observers were aware of the change, $t(17) = 4.3$, $p < 0.0001$, $SEM = 16.7$, CI (37, 107), but this effect completely disappeared when observers were unaware of the change. Thus, while conflict between the final orientation of the changed item and the orientation of the probe slows down responses when observers are aware of the change, there is no such effect on reaction times when observers are unaware of the change.

Error data were analyzed using the same design as for the RTs. For this analysis, errors were those trials in which observers responded horizontal when they should have responded vertical, or vertical when they should responded horizontal. Trials in which observers pressed the wrong key or pressed the aware key before the orientation response, were coded as errors in the program, but were not included in the error analysis (these type of errors were very infrequent, occurring in less than 0.3% of trials). The validity analysis revealed a main effect of awareness, $F(1, 17) = 18$, $p < 0.001$, $MSE = 7.4$. There were more errors when observers were unaware of the change than when they reported seeing a change (see Fig. 5).

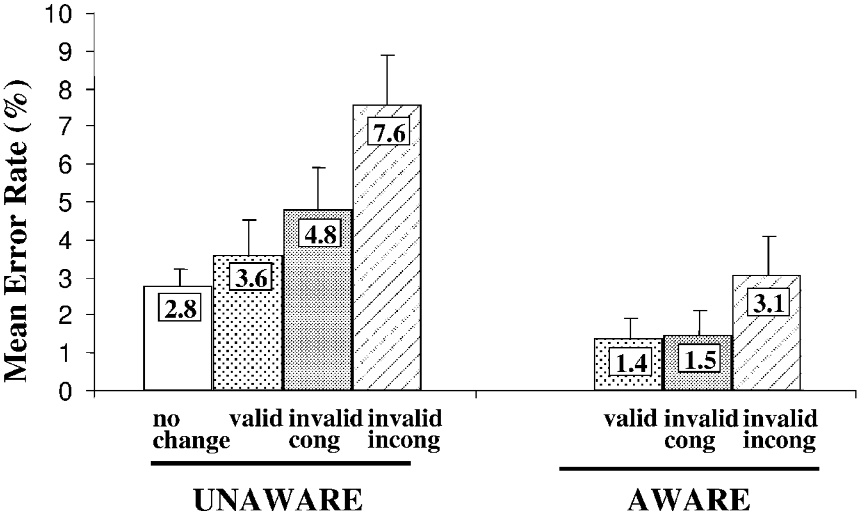

Figure 5. Error Rates for Experiment 1.

There were no effects of validity nor were there any interactions between validity and awareness.

The congruency analysis also revealed a main effect of awareness $F(1, 17) = 13.9$, $p < 0.002$, $MSE = 19$, with more errors being made in the unaware condition. More importantly, there was a main effect of congruency, $F(1, 17) = 6.9$, $p < 0.02$, $MSE = 13$. Figure 5 shows that errors were more frequent when the final orientation of the changed item was incongruent to the probe orientation than when these orientations were congruent. Notably, this congruency effect was present for both aware and unaware responses, as there was no interaction between congruency and awareness. Thus, even when observers reported being unaware of the change, the nature of the change affected their response.

3.3. Discussion

The results of Experiment 1 demonstrate that undetected change can influence performance on a subsequent task. The presence of a reaction time validity effect for unaware responses is consistent with a reallocation of attention to the location of change even in the absence of awareness. In our previous work we had shown a sensitivity to the location of change (Fernandez-Duque and Thornton, 2000) but found no evidence that such sensitivity was mediated by a reallocation of attention. One important difference between the current paradigm and our previous work is the nature of the probe itself. In our previous work, the dimension of change (i.e. orientation) was irrelevant for the subsequent discrimination task, which required in some experiments a color judgment on the rectangle, and in other experiments a discrimination between two objects. In the current experiment, the dimension of change (i.e. orientation) is also the relevant dimension for the probe response, and as such it might be a more effective cue. The error data failed to reveal a validity effect. However, this lack of validity effect in error rates is not very surprising as invalid congruent trials provide correct information for the line orientation task.

The most important finding of the current experiment is that, even in the absence of awareness, the *nature* of the change had an impact on the accuracy of responses. More errors were triggered by incongruent trials than by congruent ones, both in aware and unaware conditions. Given the very low overall error, a significant congruency effect is quite impressive. On the other hand, the size of the effect is undeniably small and low error rates are often accompanied by reduced sample variability — due to ceiling performance — that can sometimes artificially increase statistical power. We specifically address this issue in Experiment 2.

Unlike the error rates, reaction time data only revealed a congruency effect for aware trials. The lack of a reaction time congruency effect in the unaware trials indicates that the unaware congruency effect in the error rate is not due to contamination from aware trials. That is, if trials in which the observer was aware were incorrectly reported as unaware, then such contamination should also have been expressed in the reaction time data.

Aware trials were generally slower than unaware trials. Besides some form of strategic slowing, this reaction time effect may be a reflection of some form of attentional blink (e.g. Raymond *et al.*, 1992) or psychological refractory period (e.g. Welford, 1952). That is, additional processing or response preparation associated with consciously detecting the change is likely to interfere with the orientation response.

In general, error rates were higher in unaware trials than in aware trials. One possible explanation for this effect is that a conscious detection of change allows observers to inhibit the prepotent response, by allocating more time to press the key. In the absence of such a voluntary inhibitory effect, unaware changes in the incongruent trials would lead to fast RTs but lower accuracy. Alternatively, the larger number of errors in the unaware trials might be an artifact of the way we assessed awareness of change. More specifically, in this study we used a go/no-go response method to measure subjective awareness. After making an orientation response, observers pressed a key if they were aware of a change (the 'go' response) and did nothing if they were unaware (the 'no-go' response). To ensure that observers performed the discrimination task properly, errors in the discrimination task were immediately followed by a feedback tone. It is conceivable that observers' reaction to such feedback (e.g. concern, frustration, confusion, etc.) made them fail to respond to the go/no-go portion of the trial, thus artificially increasing the unaware error rates. In Experiment 2, we attempt to replicate the error congruency effect found in Experiment 1, while controlling for this possible confound.

4. EXPERIMENT 2

Experiment 1 provides initial evidence that the nature of an undetected change can influence subsequent performance. Even when unaware of the change, observers were still more prone to make an error on the orientation task when the orientation decision followed an incongruent change.

The size of this congruency effect, however, and the error rates in general, were modest at best. In Experiment 2, we attempt to boost error rates by making the orientation discrimination task more difficult. This was achieved by reducing the duration of the orientation probe. In Experiment 1, the probe and final display had remained visible until the observer responded. Here, we present the probe for only 20 ms. While such a briefly flashed item would almost certainly be available for substantially longer than 20 ms, due to visual or informational persistence, (Di Lollo and Wilson, 1978; Coltheart, 1980) our goal was simply to make detection of the probe relatively more difficult. We were interested in whether the previously observed unaware congruency effect would be maintained under these conditions.

Experiment 2 also tries to rule out the possibility that the error congruency effect in unaware trials was artificially inflated by the withholding of the awareness response following an error. To achieve this, we replaced the go/no-go response method with a 2AFC decision. If a congruency effect is still observed for unaware trials in Experiment 2, this will rule out the possibility that the pattern observed in Experiment 1 was due to withheld aware responses following an error.

4.1. Method

4.1.1. Observers. Sixteen University of Oregon students took part in this study for partial course credit. All observers had normal or corrected to normal vision, were right handed and were naive as to the purpose of the study. In order to calculate accurate estimates of behavior, each observer needed to provide a minimum number of aware and unaware responses. Four observers were dropped from the analysis due to failure to discriminate between catch trials and change trials.

4.1.2. Stimuli and equipment. All aspects of this experiment were identical to Experiment 1, except that the final probe item and display only remained visible for 20 ms.

4.1.3. Procedure. The procedure was the same as in Experiment 1, except that additional training was given to familiarize observers with the brief probe. Specifically, observers were first shown the structure of a trial with a probe that remained visible for 200 ms. They then completed 20 training trials with this probe duration, followed by a further 20 trials with a 40 ms probe, before finally being trained on 20 trials with the 20 ms probe. Also, as both responses now involved 2AFC decisions, we attempted to reduce key mapping problems by assigning the speeded responses (horizontal or vertical) to the dominant hand and the unspeeded awareness responses (aware or unaware) to the opposite hand. We also aligned the two response keys for the discrimination task vertically and the two response keys for the awareness task horizontally.

4.2. Results

Observers reported being aware in 64% of the trials with a change, but also falsely reported a change in 28% of catch trials. These levels of reported awareness were higher than in Experiment 1, particularly for the catch trials, revealing a shift in response criterion toward the more liberal end. It seems possible that the brief probe might sometimes have been mistaken for the change in orientation. This would explain the increase in false alarms. In general, however, observers were able to discriminate better than chance between the change trials and the catch trials. This was true for trials in which observers' responses in the orientation task were correct, $d_{\mathrm{L}} = 1.943$, $t(11) = 6.33$, $p < 0.0001$, $SEM = 0.307$, $CI\ (1.27, 2.62)$, and also for error trials $d_{\mathrm{L}} = 1.3945$, $t(11) = 5.6$, $p < 0.0001$, $SEM = 0.25$, $CI\ (0.85, 1.94)$, although the discriminations in error trials was lower than in the correct trials, $t(11) = 2.26$, $p < 0.05$, $SEM = 0.25$, $CI\ (0.014, 1.08)$.

Reaction time and error rates were again analyzed using 2 × 2 ANOVAs for both validity and congruency effects. Figure 6 shows the reaction time data for

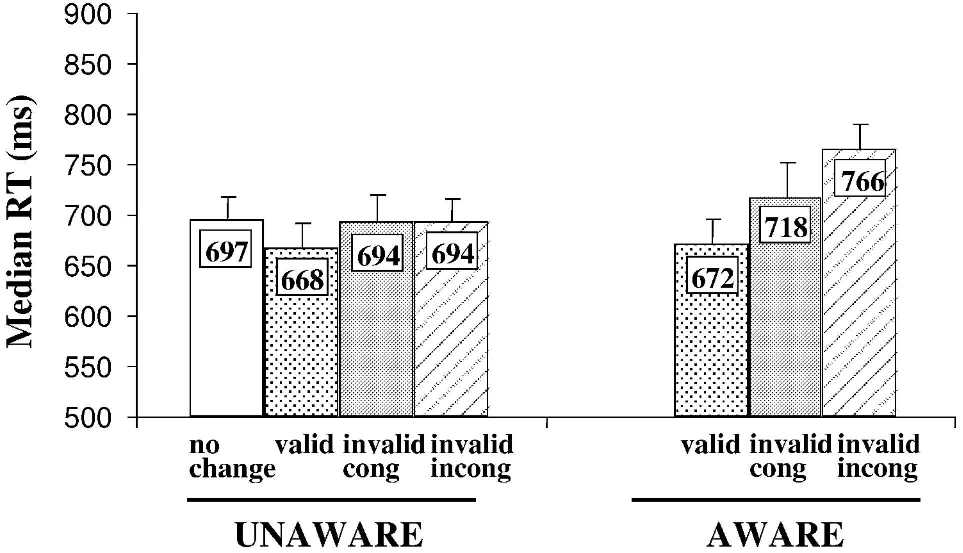

Figure 6. Median Reaction Times (in ms) for Experiment 2. Error bars indicate the Standard Error of the Mean.

all conditions. As in Experiment 1, there was a clear validity effect, $F(1, 11) = 7$, $p < 0.05$, $MSE = 2198$, with valid trials being faster than invalid trials. This validity effect did not interact with awareness.

There was a fairly strong congruency effect for aware trials, $t(11) = 2.7$, $p < 0.05$, $SEM = 17.6$, CI (9, 87), which was completely absent for unaware trials, a pattern of results that again replicates the findings of Experiment 1. In the analysis of variance there was a trend toward a main effect of congruency, $F(1, 11) = 2.8$, $p = 0.12$, $MSE = 2420$, and a trend towards a congruency × awareness interaction, $F(1, 11) = 2.4$, $p = 0.15$, $MSE = 2955$.

There was a general slowing of aware responses relative to unaware responses, although this effect was slightly weaker than in Experiment 1. Specifically, the difference was absent in the validity analysis, but was present in the congruency analysis, $F(1, 11) = 5.7$, $p < 0.04$, $MSE = 4935$.

Figure 7 reveals a large increase in the overall error rate for this experiment relative to Experiment 1 (note change of scale). Obviously, our attempt to increase error rates was successful.

Unlike in Experiments 1, there was a validity effect in the pattern of error data. This effect was present in both aware and unaware trials, $F(1, 11) = 7.6$, $p < 0.02$, $MSE = 24.6$, and did not interact with awareness. This increased error rate in invalid trials might represent the cost of being cued to the location of change, and therefore missing the 20 ms probe.

Analysis of the congruency effect for error rates revealed both a main effect of congruency, $F(1, 11) = 38$, $p < 0.0001$, $MSE = 131$, and an interaction with awareness, $F(1, 11) = 18$, $p < 0.001$, $MSE = 60$. Although the congruency

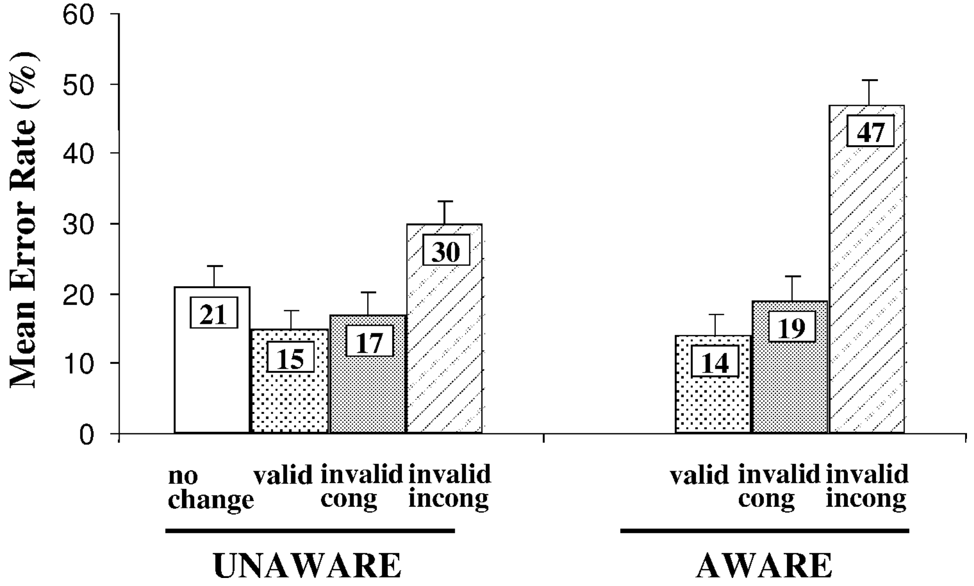

Figure 7. Error Rates for Experiment 2. Notice that the scale has changed between Experiments 1 and 2.

effect was larger for aware trials (a 28% difference in error rate) than for unaware (a 13% difference), in both conditions it was highly significant (aware: $t(11) = 6.5$, $p < 0.0001$, $SEM = 4.2$, CI (19, 37); and unaware: $t(11) = 2.8$, $p < 0.02$, $SEM = 4.5$, CI (3, 23)). It seems plausible that this interaction arises because becoming aware of the change increases the likelihood of missing the brief probe. If the probe is missed, observers may tend to respond based on the final orientation of the change. This would result in an error for invalid incongruent trials, in particular for trials in which the change is consciously perceived. Consistent with this notion, the congruency analysis indicated that there were generally more errors in aware trials than unaware trials, $F(1, 11) = 18$, $p < 0.001$, $MSE = 60$.

4.3. Discussion

The main goal of Experiment 2 was to boost the overall error rate and explore whether under conditions of lower accuracy, the error congruency effect in unaware trials was maintained. The use of a brief probe was effective in increasing overall error rates. More importantly, the congruency effect for unaware trials was still present, now representing a 13% difference in errors between congruent and incongruent trials.

A new finding in the current experiment was the appearance of validity effects in the error data. For both aware and unaware trials, making an orientation response away from the location of change led to more errors than making the decision at the location of change. This validity effect in the error data probably reflects the increased difficulty of detecting the probe. In Experiment 1 the probe remained visible until response, making it quite unlikely that observers would fail to detect it. In contrast, in Experiment 2 the probe only remained visible for 20 ms. Therefore, the cost of being cued away from the probe location was greatly increased in Experiment 2, and this could account for the error validity effect.

Compared to Experiment 1, the level of awareness was somewhat higher in Experiment 2 and the discrimination of change trials from catch trials was also poorer. It is likely that these effects were also due to the brief probe presentation. Observers were instructed to use a liberal criterion for awareness, that is, to respond ‘aware’ even if they only had a ‘feeling’ that something had changed during the flicker part of the task. In some trials, observers might have completely missed the change but partially detected the flash. In those trials, it is plausible that observers would report being aware of a change, although in fact they were not.

In Experiments 1 and 2, the final frame of the change display contained an unequal number of horizontal and vertical items. Such an imbalance in the final frame could have influenced patterns of responding. However, a control study in which equal numbers of horizontal and vertical items were presented in the final frame, strongly argues against this explanation. Specifically, when this confound was eliminated, a healthy congruency effect for unaware trials was still observed.

5. GENERAL DISCUSSION

Two experiments demonstrated that the *nature* of an undetected change can influence responses on a subsequent task. Specifically, even when observers reported being unaware that a change had taken place, that change still biased the pattern of errors in a subsequent orientation discrimination task. For example, when an item changed from horizontal to vertical, observers were more likely to respond ‘vertical’ to a horizontal probe. This effect is not dependent on the spatial overlap between the changed item and the probe as on the critical trials they always appear on opposite sides of the screen.

The current findings are an important extension of our previous work, in which we demonstrated that the *location* of a change can be represented without awareness (Fernandez-Duque and Thornton, 2000). These two effects, along with recent findings using a variety of different paradigms (e.g. Chun and Jiang, 1998; Hayhoe *et al.*, 1998; Cohen *et al.*, 1999; Henderson and Hollingworth, 1999; Klein and MacInnes, 1999; Smilek *et al.*, 2000) converge on the idea that the visual system can represent more about the dynamics of a visual scene than previous research on change blindness would lead us to believe (e.g. Rensink *et al.*, 1997; Horowitz and Wolfe, 1998).

As the evidence continues to accumulate, an important focus for future research will be on trying to identify the underlying *mechanisms* that allow dynamic events to be processed in the absence of awareness. A fundamental question in this regard concerns the role of attention. Specifically, does attention, which is thought to provide much of the spatial and temporal integration needed for *conscious* perception, perform a similar function during *unconscious* perception?

A fairly close link is generally assumed to exist between attention and awareness (Treisman and Kanwisher, 1998). That is, attending to an object typically entails being aware of at least some of its properties (Levin and Simons, 1997). Proposing an attentional modulation of implicit representations might thus seem somewhat counterintuitive. Nevertheless, it remains possible that integrating representations across time always involves attention, but at levels that do not necessarily involve conscious perception.

Attention might affect performance in implicit tasks by biasing the response. It is possible that, when forced to guess the location of a change not consciously perceived, observers might select the most attended location. Thus, if a changing item can effectively cue attention toward it, observers could successfully guess its location, even if they do not consciously perceive the change (Graves and Jones, 1992; McCormick, 1997). In previous studies, we directly tested the hypothesis that an undetected change could effectively cue attention to its location (Fernandez-Duque and Thornton, 2000). We reasoned that such a cueing would facilitate a discrimination response at the location of change while impairing it at the opposite location (i.e. validity effect).

In these original studies we found no evidence for re-deployment of attention as a mechanism in implicit localization of change. In the current studies we again tested for possible cueing effects, but this time modified our paradigm to include a type of change (i.e. change in orientation) that was relevant for the discrimination task (i.e. an orientation discrimination). With this modification, we did find reliable validity effects for aware and unaware trials in both experiments reported here. Thus, it is possible that attention was cued to the location of change, but at levels below the threshold for conscious awareness.

While the findings of the current study are consistent with spatial cueing as a mechanism for implicit localization of change, there are several reasons to urge caution in accepting this explanation. In addition to the lack of a validity effect in our previous work (Fernandez-Duque and Thornton, 2000), other studies in our laboratory using the current orientation discrimination task, have also failed to produce unaware validity effects, despite the presence of reliable unaware congruency effects (Thornton and Fernandez-Duque, 1999). The validity effect might also be particularly susceptible to contamination from aware trials. As discussed below, while we are confident that the unaware congruency effect is not due to observers mistakenly reporting lack of awareness, we cannot rule out this possibility for the validity effect. Finally, even if future studies establish subthreshold spatial cueing as an important mechanism for implicit *localization* of change it would not provide a parsimonious explanation for the *congruency* effect, in which a changing item affects the response to a probe diametrically opposite to it.

In addition to spatial cueing, there are other ways in which attention could facilitate implicit representations. For example, a *distributed* allocation of attentional resources across the visual field, rather than a shifting of *focused* attention, could also provide spatio-temporal integration in the absence of awareness. In favor of this hypothesis, preliminary studies reveal that forcing observers to focus attention at the center of the screen via a secondary task effectively eliminates the ability to implicitly localize change (Fernandez-Duque and Thornton, 1999).

It is clearly not possible to rule out attentional modulation as a contributing factor to the current findings of change detection without awareness. There are, however, alternative, non-attentional mechanisms, that can also be considered. For example, a recent theory of scene perception has proposed a 'triadic' architecture for the visual mechanisms responsible for scene and object processing (Rensink, 2000b). In addition to a low-level feature system and an attentional system, Rensink proposes a 'setting' system that provides information about the scene's layout and gist. The 'setting' system is theorized to be fast, automatic and to work in parallel to (and outside of) the attentional system, providing quite detailed information about the scene. As such, this 'setting' system might implicitly represent the nature of the change (i.e. gist), and its location (i.e. 'layout'). If the location of change could become 'marked' in a layout system, then such a non-attentional layout mechanism might allow the 'selection without awareness' observed in our localization paradigm (Fernandez-Duque and Thornton, 2000).

The 'triadic' model represents a major advance in our understanding of scene perception due to its explicitness about the structure of the components, its computational sophistication, and its testable predictions. Milner and Goodale's (1995) model, which emphasizes the distinction between perception and action systems, provides another framework within which to explore our implicit effects of localization and congruency. This model has the additional advantage of making strong predictions about the underlying neuroanatomy, predictions that we are currently exploring using brain imaging techniques (Fernandez-Duque *et al.*, 2000).

As the above discussion should make clear, at the moment we can do little more than speculate on the mechanisms underlying our findings. However our current finding about implicit representation of the *nature* of the change provides certain general constraints on the possible mechanisms for representing change. For example, in our original study, the *location* of change might have been distinguished from all other positions in the array by some trace of the first frame that persisted across the blank interstimulus interval and became integrated with the second frame (Di Lollo and Wilson, 1978; Coltheart, 1980; Grossberg and Mingolla, 1985). This possibility is admittedly unlikely, given the relatively long interstimulus interval (250 ms) used in those studies, but it cannot be completely ruled out. As the only place where this integration would add information to the display is at the location of change, this could drive object selection when observers were asked to guess the location of change. However, visual persistence cannot account for the congruency effect. If there was some persistence-based merging or interference between frames 1 and frames 2, this should weaken or confuse the representation of the final change orientation. Thus, persistence would decrease the likelihood of observing an unaware congruency effect, rather than increasing the likelihood. If persistence is operating in our current displays, then we might actually be underestimating the level of implicit detection!

The current finding about implicit representation of the *nature* of the change also protects us against certain artifacts that might have explained the above chance performance in the original studies. In those studies, it was possible that above chance selection of the changing item was based on knowledge about what did not change rather than about what did change. When asked to choose which of two items had changed, an observer who was confident of seeing that one did *not* change would correctly select the opposite item (i.e. the changing one) even in the absence of any implicit detection of change. Although this exclusion strategy was unlikely (see Fernandez-Duque and Thornton, 2000) it could not be completely ruled out. In contrast, the congruency effect is insensitive to the exclusion strategy. Whether the observer knows that the item diametrically opposite to the probe remained unchanged is irrelevant to the task of reporting the orientation of the probe. Thus, even when the exclusion strategy is not at play, there is evidence that the change is being represented.

Even so, to further explore whether an exclusion strategy could account for the better-than-chance performance in the original localization task (Fernandez-Duque and Thornton, 2000), we ran a control study in which the display duration of the first frame was increased from 250 ms to 750 ms. A longer display should allow the observer to 'hold' more items in mind, and therefore increase the percentage of trials in which the change is consciously detected. More importantly, in trials in which the change is not detected, an increase in the number of 'hold' items should increase the probability of noticing that the other item in the 2AFC did not change. Using the exclusion strategy of selecting the item opposite to the unchanged one should then improve performance. Although increasing the duration of the first frame was effective in increasing the percentage of aware trials, performance in unaware trials was not improved. This result strongly argues against the exclusion strategy (Thornton and Fernandez-Duque, 1999).

An alternative explanation for the current findings is that the congruency effect in unaware trials is an artifact, a contamination from trials in which observers were somewhat aware of the change, but for one reason or another they pressed the 'unaware' key. It is notoriously difficult to rule out contamination using only behavioral methods, but in the current set of findings, there are reasons to believe that this is not a major factor. Specifically, in both experiments, there were clear indications of a congruency effect in the error rates for both aware and unaware responses. However, the congruency effect only appeared in the RTs for aware trials. If a substantial number of aware trials were incorrectly reported as unaware, then such contamination should also have been expressed in the reaction time data.

Instead, in both Experiments 1 and 2, reaction times for unaware congruent and incongruent trials were almost identical.

Other lines of evidence also suggest contamination cannot be a general explanation for the detection of change without awareness. For instance, in a recent control study we replaced the aware/unaware decision with a six-point confidence scale. Even when reporting being very sure of not seeing any change, observers were able to localize the change better than chance (Fernandez-Duque and Thornton, 1999). Another approach to the contamination problem is to move away from purely behavioral measures. A technique particularly suitable for this approach would be Event Related Potentials (ERPs). Using ERPs to establish a neural marker of change detection in the absence of awareness is appealing as the technique allows post-test comparisons based on the type of response such as correct/error and aware/unaware (Fernandez-Duque *et al.*, 2000). Examination of eye-movement records may also be very useful (e.g. Hayhoe *et al.*, 1998; Cohen *et al.*, 1999; Henderson and Hollingworth, 1999).

In conclusion, we believe the current results clearly demonstrate that the *nature* as well as the *location* of changing items can be represented in the absence of awareness. The use of implicit measures of performance, in the context of change-over-time tasks, promises to shed new light on the visual processes involved in the perception of dynamic events (Rensink, 2000b; Fernandez-Duque and Thornton, 2000; Smilek *et al.*, 2000). As future studies attempt to identify possible mechanisms for these unaware processes, the neural substrates of unaware perception are likely to emerge. Those advances will almost certainly help us understand how much detail can be represented without attention, as well as the functional significance of the representational systems that operate beyond the realm of conscious awareness.

Acknowledgements

The authors contributed equally to this research, and authorship order was determined arbitrarily. Thanks are extended to Jodie Baird, Mike Posner, Ron Rensink and two anonymous reviewers for commenting on an earlier draft of this manuscript. Thanks are also extended to Carrie Harp and Kristen Macuga for helping with data collection.

REFERENCES

Bar, M. and Biederman, I. (1998). Subliminal visual priming, *Psychol. Sci.* **9** (6), 464–469.

Bridgeman, B., Hendry, D. and Stark, L. (1975). Failure to detect displacements of the visual world during saccadic eye movements, *Vision Research* **15**, 719–722.

Chen, Z. (1998). Inattentional amnesia: Evidence of perception without awareness, *Investigative Ophthalmology and Visual Science* **39** (4), S629.

Chun, M. M. and Jiang, H. (1998). Contextual cueing: implicit learning and memory of visual context guides spatial attention, *Cognitive Psychology* **36**, 28–71.

Cohen, N. J., Ryan, J. and Althoff, R. (1999). The eye movement-based relational manipulation effect: Evidence for detailed representations of the world. Paper presented at the *European Conference on Eye Movements*, Utrecht University, The Netherlands.

Coltheart, M. (1980). The persistence of vision, *Phil. Trans. Roy. Soc. London, B* **290**, 57–69.

Dehaene, S., Naccache, L., Le Clec'H, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., van de Moortele, P. F. and Le Bihan, D. (1998). Imaging unconscious semantic priming, *Nature* **395**, 597–600.

Di Lollo, V. and Wilson, A. E. (1978). Iconic persistence and perceptual moment as determinants of temporal integration in vision, *Vision Research* **18** (12), 1607–1610.

Duncan, J. and Humphreys, G. W. (1989). Visual search and stimulus similarity, *Psychol. Rev.* **96**, 433–458.

Enns, J. T. and Rensink, R. A. (1990). Sensitivity to three-dimensional orientation in visual search, *Psychol. Sci.* **1**, 323–326.

Eriksen, C. W. and Collins, J. F. (1969). Temporal course of selective attention, *J. Exper. Psychol.* **80** (2), 254–261.

Fernandez-Duque, D. and Thornton, I. M. (1999). Detecting change in the absence of awareness, *Investigative Ophthalmology and Visual Sciences* **40** (4), S4212.

Fernandez-Duque, D. and Thornton, I. M. (2000). Change detection without awareness: Do explicit reports underestimate the representation of change in the visual system? *Visual Cognition* **7** (1/2/3), 323–344.

Fernandez-Duque, D., Grossi, G., Thornton, I. M. and Neville, H. (2000). Neural correlates of change detection: An ERP study. Poster presented at the *Annual Meeting of the Cognitive Neuroscience Society*, San Francisco.

Freyd, J. J. (1987). Dynamic mental representations, *Psychol. Rev.* **94**, 427–438.

Graves, R. E. and Jones, B. S. (1992). Conscious visual perceptual awareness vs. non-conscious visual spatial localisation examined with normal subjects using possible analogues of blindsight and neglect, *Cognitive Neuropsychology* **9** (6), 487–508.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades, in: *Perception, Vancouver Studies in Cognitive Science*, Akins K. (Ed.), Vol. 5, pp. 89–109. Oxford University Press, New York.

Grossberg, S. and Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations, *Perception and Psychophysics* **38**, 141–171.

Harp, C. J. and Rensink, R. A. (1999). A comparison of attentional processing in younger and older observers, *Investigative Ophthalmology and Visual Science* **40** (4), S270.

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

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

Horowitz, T. S. and Wolfe, J. M. (1998). Visual search has no memory, *Nature* **394**, 575–577.

Irwin, D. E. (1991). Information integration across saccadic eye movements, *Cognitive Psychology* **23**, 420–456.

Jones, M. R. (1976). Time, our lost dimension: Toward a new theory of perception, attention, and memory, *Psychol. Rev.* **83**, 323–355.

Kahneman, D., Treisman, A. M. and Gibbs, B. (1992). The reviewing of object files: object-specific integration of information, *Cognitive Psychology* **24**, 175–219.

Klein, R. M. and MacInnes, W. J. (1999). Inhibition of return is a foraging facilitator in visual search, *Psychol. Sci.* **10** (4), 346–352.

Kolb, F. C. and Braun, J. (1995). Blindsight in normal observers, *Nature* **377**, 336–338.

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

Luck, S. J., Vogel, E. K. and Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink, *Nature* **382**, 616–618.

Mack, A. and Rock, I. (1998). *Inattentional Blindness*. MIT Press, Cambridge, MA.

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

Marcel, A. J. (1983b). Conscious and unconscious perception: an approach to the relations between phenomenal experience and perceptual processes, *Cognitive Psychology* **15**, 238–300.

McCormick, P. A. (1997). Orienting attention without awareness, *Journal of Experimental Psychology: Human Perception and Performance* **23** (1), 168–180.

Macmillan, N. A. and Creelman, C. D. (1996). Triangles in ROC space: History and theory of 'nonparametric' measures of sensitivity and response bias, *Psychonomic Bulletin and Review* **3**, 164–170.

Milner, A. D. and Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford University Press, Oxford.

Moore, C. M. and Egeth, H. (1997). Perception without attention: Evidence of grouping under conditions of inattention, *Journal of Experimental Psychology: Human Perception and Performance* **23** (2), 339–352.

Nakayama, K. and Joseph, J. S. (1998). Attention, pattern recognition, and pop-out in visual search, in: *The Attentive Brain*, Parasuraman, R. (Ed.), pp. 711–733. MIT Press, Cambridge, MA.

Nakayama, K. and Silverman, G. H. (1986). Serial and parallel processing of visual feature conjunctions, *Nature* **320**, 264–265.

O'Regan, J. K. (1992). Solving the 'real' mysteries of visual perception: The world as an outside memory, *Canad. J. Psychol.* **46** (3), 461–488.

O'Regan, J. K., Rensink, R. A. and Clark, J. J. (1999). 'Mudsplashes' causes blindness to large scene changes, *Nature* **398** (6722), 34.

O'Regan, J. K., Deubel, H., Clark, J. J. and Rensink, R. A. (2000). Picture changes during blinks: Looking without seeing and seeing without looking, *Visual Cognition* **7** (1/2/3), 191–212.

Pashler, H. (1988). Familiarity and visual change detection, *Perception and Psychophysics* **44**, 369–378.

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

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

Pomplun, M., Shen, J. and Reingold E. M. (1999). The visual span in comparative visual search. Paper presented at the *Annual Meeting of the Cognitive Science Association for Interdisciplinary Learning (CSAIL)*, Hood River, Oregon.

Raymond, J. E., Shapiro, K. L. and Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance* **18** (3), 849–860.

Rensink, R. A. (1990). Toolbox-based routines for Macintosh timing and display, *Behavior Research Methods, Instruments and Computers* **22** (2), 105–117.

Rensink, R. A. (2000a). Visual search for change: A probe into the nature of attentional processing, *Visual Cognition* **7** (1/2/3), 345–376.

Rensink, R. A. (2000b). The dynamic representation of scenes, *Visual Cognition* **7** (1/2/3), 17–42.

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

Scholl, B. J. and Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood, *Cognitive Psychology* **38** (2), 259–290.

Scholl, B. J. (2000). Attenuated change blindness for exogenously attended items in a flicker paradigm, *Visual Cognition* **7** (1/2/3), 377–396.

Shore, D. I. and Klein, R. M. (2000). The effects of scene inversion on change blindness, *J. General Psychol.* **127** (1), 27–43.

Simons, D. J. (1996). In sight, out of mind: When object representations fail, *Psychol. Sci.* **7**, 301–305.

Simons, D. J. and Levin, D. T. (1997). Change blindness, *Trends in Cognitive Sciences* **1** (7), 261–267.

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

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

Snodgrass, J. G. and Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia, *Journal of Experimental Psychology: General* **117** (1), 34–50.

Steinman, S. B. and Nawrot, M. (1992). Real-time color-frame animation for visual psychophysics on the Macintosh computer, *Behavior Research Methods, Instruments and Computers* **24** (3), 439–452.

Thornton, I. M. and Fernandez-Duque, D. (1999). Representing the 'what' and the 'where' of undetected change. Poster presented at the *40th Annual Meeting of the Psychonomic Society*, Los Angeles, CA.

Treisman, A. M. and Gelade, G. (1980). A feature-integration theory of attention, *Cognitive Psychology* **12**, 97–136.

Treisman, A. M. and Kanwisher, N. G. (1998). Perceiving visually presented objects: recognition, awareness, and modularity, *Current Opinion in Neurobiology* **8**, 218–226.

Wallis, G. and Bülthoff, H. H. (2000). What's scene and not seen: Influences of movement and task upon what we see, *Visual Cognition* **7** (1/2/3), 175–190.

Welford, A. T. (1952). The 'psychological refractory period' and the timing of high-speed performance — A review and a theory, *Brit. J. Psychol.* **43**, 2–19.

Wolfe, J. M. (1994). Guided Search 2.0 — A revised model of visual-search, *Psychonomic Bulletin and Review* **1** (2), 202–238.

Wolfe, J. M. (1998). Visual search, in: *Attention*, Pashler, H. (Ed.), pp. 13–73. University College London Press, London, UK.

Wolfe, J. M. and Bennett, S. C. (1997). Preattentive object files: Shapeless bundles of basic features, *Vision Research* **37** (1), 25–43.

Zelinsky, G. (1997). Eye movements during change detection search task, *Investigative Ophthalmology and Visual Science* **38**, S373.