

Picture Changes During Blinks: Looking Without Seeing and Seeing Without Looking

J. Kevin O'Regan

Laboratoire de Psychologie Expérimentale, CNRS, Paris, France

Heiner Deubel

Institute of Psychology, Ludwig-Maximilians University, Munich, Germany

James J. Clark

Electrical Engineering Department, McGill University, Montreal, Canada

Ronald A. Rensink

*Cambridge Basic Research, Nissan Research & Development Inc.,
Cambridge, MA, USA*

Observers inspected normal, high quality colour displays of everyday visual scenes while their eye movements were recorded. A large display change occurred each time an eye blink occurred. Display changes could either involve “Central Interest” or “Marginal Interest” locations, as determined from descriptions obtained from independent judges in a prior pilot experiment. Visual salience, as determined by luminance, colour, and position of the Central and Marginal Interest changes were equalized.

The results obtained were very similar to those obtained in prior experiments showing failure to detect changes occurring simultaneously with saccades, flicker, or “mudsplashes” in the visual scene: Many changes were very hard to detect, and Marginal Interest changes were harder to detect than Central Interest changes.

Analysis of eye movements showed, as expected, that the probability of detecting a change depended on the eye's distance from the change location. However a surprising finding was that both for Central and Marginal Interest changes, even when observers were directly fixating the change locations (within 1 degree), more than 40% of the time they still failed to see the changes. It seems that looking at something does not guarantee you “see” it.

Please address all correspondence to J. Kevin O'Regan, Laboratoire de Psychologie, Expérimentale, 28 rue Serpente 75006 Paris, France. Email: oregan@ext.jussieu.fr

The results are interpreted in terms of the idea that what the observer sees during scene exploration is not determined primarily by the location in the scene being fixated by the eyes, but by the scene *aspects* being attended to. A theory is summarized which applies to the other experiments in the literature on change blindness, and some new predictions are put forward.

INTRODUCTION

Recent studies have shown that surprisingly large changes in a visual scene can go undetected if the changes coincide temporally with the moment an eye saccade occurs (Grimes, 1996; McConkie, 1991; McConkie & Currie, 1996). The phenomenon seems not to be related to some kind of saccade-specific suppression mechanism, but simply to the visual transient that the saccade produces: Indeed, the effect can also be obtained *without* saccades if a large visual transient such as a screen flicker (Rensink, O'Regan, & Clark, 1995, 1997) or a film cut (in a motion picture sequence, cf. Levin & Simons, 1997) occurs simultaneously with the change, or if the change occurs at the same time as some small local transients, such as mudsplashes on a car windshield, are briefly spattered over the scene, (O'Regan, Rensink, & Clark, 1996, 1999).

If the explanation of the effect truly lies in the distracting influence of the visual transients that occur simultaneously with the change, then another, ecologically rather important form of transient is likely also to be effective in preventing changes from being detected, namely blinks. Like saccades, blinks create a global disruption of the retinal image, and furthermore, at around 100–200 msec, their durations are actually several times longer than saccades (20–70 msec). Blinks occur at a relatively high rate (20–30 times per minute, cf. Sun et al., 1997) and so in tasks like driving, which requires continual monitoring of visual conditions, it would be important to know to what extent a change in a traffic light or brake light, or a dangerous road event, may be missed if it occurs during a blink.

It is clear that in any scene there will be aspects that observers will preferentially attend to, and others that they will tend not to attend to. Because this attentional component of scene recognition would be expected to interact strongly with the likelihood of detecting a change in the scene, we wished to find a way of quantifying it. A variety of studies have been done in the attention and picture memory literature to attempt to understand what causes a picture aspect to be noticed or remembered, and factors such as visual salience (contrast, size, location, colour, onset characteristics), semantic (coherence, novelty), contextual, and task-related factors, as well as subjects' individual preferences, are certainly all involved. For the purposes of our study however, a simple operational measure was required, and we chose to obtain it in the following way.

We asked a group of five judges to give a short description of each picture. The judges' descriptions were very coherent, enabling us to define, for each picture, several aspects, which we termed "Central Interest" aspects, defined as those aspects that were mentioned by at least three of the judges. Central Interest aspects tended to concern what one would be tempted to call the main theme of the picture. Similarly, we noted several aspects of each picture, which we termed "Marginal Interest" aspects, defined as those aspects that were mentioned by none of the judges. In the main experiment, then, by manipulating whether changes were caused to occur to Central or Marginal Interest aspects, we had a way of controlling the degree of attention that our subjects would be expected to pay to the changes.

An additional point that will be addressed in this paper concerns the question of eye position: Up until now, little research has been done in natural visual scenes to investigate how close to a display change the eye has to be in order for the change to be seen (but see Currie, McConkie & Irwin, 1997; Hayhoe, Bensinger, & Ballard, 1998; McConkie & Currie, 1996; Zelinsky, 1997). In the present experiment we were able to monitor eye movements while observers inspected the images. We expected that a change would be easier to detect if the eye was looking directly at the change location than if the eye was some ways away from it. We indeed found that this was the case. However a curious additional fact emerged, namely that directly fixating the change was still no guarantee that the change would be seen. In fact, in more than 40% of the cases, the eye looked directly at the change and yet did not notice it. This phenomenon seemed independent of whether the location was a Central or Marginal Interest location. We called it "looking but not seeing".

METHOD

Materials

The experiment involved 48 digitized colour photographs depicting a variety of natural indoor and outdoor scenes, derived from a commercial picture database (Corel). With each picture was associated a modified picture, in which some region had been changed. There were three kinds of changes: Appearance or disappearance of a picture element (this could be an object, part of object, surface, or region such as a shadow or the sky); a shift in the position of an object; a change in the colour of an object, surface or region. The changes were made by editing the images with image manipulation programs (Adobe Photoshop, Deneba Canvas). Pictures and associated changes were the same as used in our earlier experiments (cf. Rensink et al., 1997, this issue).

In half of the pictures the locations of the picture that were changed were chosen to be Central Interest (CI) aspects, and half were chosen to involve Marginal Interest (MI) aspects, defined as explained previously. The change

locations were the same as those used in our earlier experiments. Even though Marginal Interest changes tended to involve aspects of the picture that did not constitute the "main theme" of the picture, we were careful to choose them so that they physically occupied as much space and were approximately as centrally located in the picture as Central Interest changes. Thus, the mean centroid of the CI locations were at (x,y) pixel coordinates ($-8 \pm 136, -5 \pm 86$, i.e. $0.4 \pm 6, 0.2 \pm 4$ degrees) relative to the centre of the 640×480 pixel screen, and at coordinates ($92 \pm 195, -11 \pm 139$, i.e. $4 \pm 8.7, 0.5 \pm 6$ degrees) for the MI locations (the \pm values are standard deviations). The nature of the change between the original and modified pictures was also carefully controlled so as to ensure that there was as little difference as possible in the visual conspicuity of the change for the CI and MI changes. Thus, the mean proportion of modified pixels in the picture was $4.2 \pm 4\%$ and $5.1 \pm 4\%$, for CI and MI changes respectively. The mean Euclidean distance in the RGB values (each coordinate potentially ranging from 0 to 255) across all the changed pixels was 104 ± 24 and 101 ± 51 respectively for CI and MI changes. Finally, the mean intensity change (potentially ranging from 0, i.e. black to 255, i.e. white) of the changed pixels was 127 ± 63 and 126 ± 110 for CI and MI changes. These values show that on purely visual measures, the CI and MI changes were very similar. Any difference in the probability of noticing the changes was therefore probably attributable to differences in the semantic relevance of the changes.

Procedure

Biting the bite plate made with dental compound, the subjects sat with their eyes at a distance of 80 cm from a 100 Hz colour video monitor in a darkened room. The area of the screen in which the picture appeared subtended 21° vertically and 28° horizontally at the eye. An SRI dual Purkinje eye tracker registered the subjects' eye movements and was interfaced to the computer that controlled stimulus presentation.

After calibration of the eye tracker using an array of nine dots on the screen, the subjects were shown 10 practice pictures, followed by the 48 test pictures randomly mixed anew for each subject. A short break was made half way through the experiment.

Subjects were told that at some moment while they were looking at each picture, the picture would change, and that their task was to press a button immediately this occurred. While the button was maintained pressed, the computer displayed a 20×20 pixel ($0.8^\circ \times 0.8^\circ$) white outline square on the screen which moved over the picture and continuously indicated where the eye was looking. As this was occurring, subjects were asked to look at the location where they had seen the change, thereby indicating its location with the eye-contingent moving square, and release the button. After this the screen went dark and a

fixation cross appeared in the screen's middle. The subject looked at the fixation cross, and pressed the button again, causing the next picture to appear.

Picture changes occurred every time the subject blinked: The original picture was replaced by the changed picture, or vice versa. If after 48 seconds or 15 blinks the change had not yet been detected by the subject, the computer waited for a random time between 0.5 and 1 sec, and displayed the change without waiting for a blink. This ensured that subjects would always ultimately see the change.

The subjects were not told that their blinks caused the display changes, nor were they told how many times a change would occur. Since they were actively searching for changes, some subjects adopted the strategy of staring at the screen intently, which unfortunately had the consequence of preventing blinks. Such subjects were told that the eye tracker worked better when the eye was moist and that they should therefore blink every now and then. After the experiment the experimenter interviewed the subjects to determine whether they had detected a correlation with blinks. Of the 10 subjects, 2 said they had noticed some kind of correlation.

Eye movements were sampled continuously at 400 Hz during the trials. Blinks were detected by monitoring an output channel of the SRI tracker which indicates the presence of a "valid" pupil in the infrared image. The display change was made by switching video display buffers, and could thus be effected during the retrace period between two successive video frames. Because display changes occurred entirely within the blink period, generally no visual transient at all was visible to subjects and the image appeared perfectly steady. Exceptions to this were cases when the subjects blinked without completely closing their eyes, or when reflections, eyelashes, or other artefacts caused the computer to register a blink when in fact none had occurred (see later).

Subjects

Ten paid naïve subjects, students at the Ludwig Maximilians Universität, participated in the experiment.

RESULTS

Data Analysis

Detection of blinks by the pupil signal of the SRI tracker is not completely reliable. The eyelid may only partially close, or reflections or eyelashes may partially interfere with the image. Whereas during the experiment the display changes were triggered by the pupil-disappearance signal, after the experiment it was possible to additionally examine the eye-movement record in relation to

the pupil signal. A normal blink can be recognized because it is accompanied by a distinctive "signature" in the output of the eye tracker, involving a typical rapid vertical movement and return of the eye (cf. Collewijn, van der Steen, & Steinman, 1985). By carefully examining the eye traces by hand we eliminated records where we thought the change had been incorrectly triggered by some artefact. This data selection was done by the experimenters, but blind to the nature and outcome of each trial.

A further question in the data analysis concerns what we mean by the location of the eye at the moment of the blink. Analysis of the records showed that in the majority of cases, the eye blink occurred at the beginning of a fixation and ended before the end of the fixation. In some cases, however, the blink straddled an eye movement. However, in these cases the eye very rarely moved further than 1 degree from the location that it had occupied before the blink. In total, 95% of the time, the eye's position after the blink was within 1 degree of its position before the blink. The other 5% of the cases were eliminated from the data analysis.

Overall, through the two causes of blink detection artefacts and through the eye moving further than 1 degree during the blink, 64 records from the total of 480 had to be eliminated. This is a significant loss in the quantity of available data and forced us to pool the data for the three different types of display changes instead of analysing them separately.

Number of Blinks Before Detection

Figure 1 shows the number of cases in which observers detected the change on the first, second, third, etc. blink, for the Central Interest and the Marginal Interest type of changes. The "W" bars on the right-hand side of the histograms correspond to "wrong" cases, that is, cases where the subject indicated a change had occurred somewhere where no change had occurred. The "N" bars correspond to cases where no change was detected until the computer caused the change to occur without a blink taking place.

We see that for the Central Interest changes (Fig. 1a), the change was most often detected on the first blink. Some changes were only detected on the second and later blinks. A few changes were wrongly detected or never detected at all before the computer displayed the change without the blink taking place. On the contrary, in the Marginal Interest case (Fig. 1b), the change was generally not detected the first time it occurred, but had to wait till the second or later occurrences. Most often, the change was never detected before the computer displayed the change without the blink taking place.

These results are very similar to our earlier results showing that changes that coincide with visual transients created by flicker or by "mudsplashes" are hard to detect, particularly when the changes are "Marginal Interest" changes.

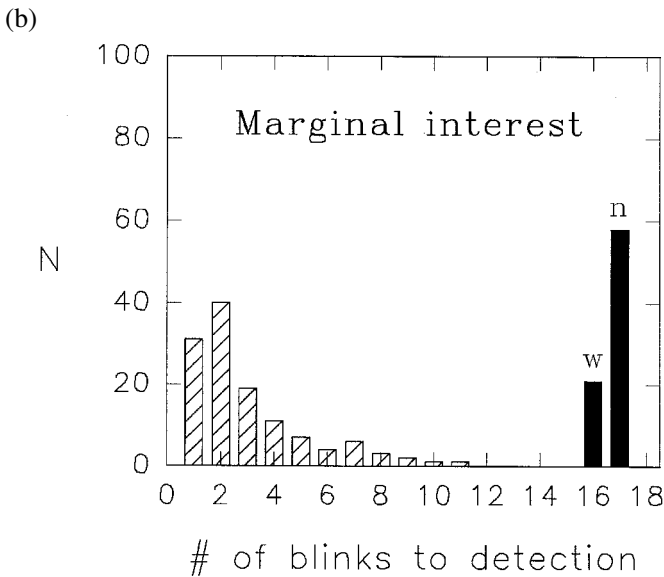
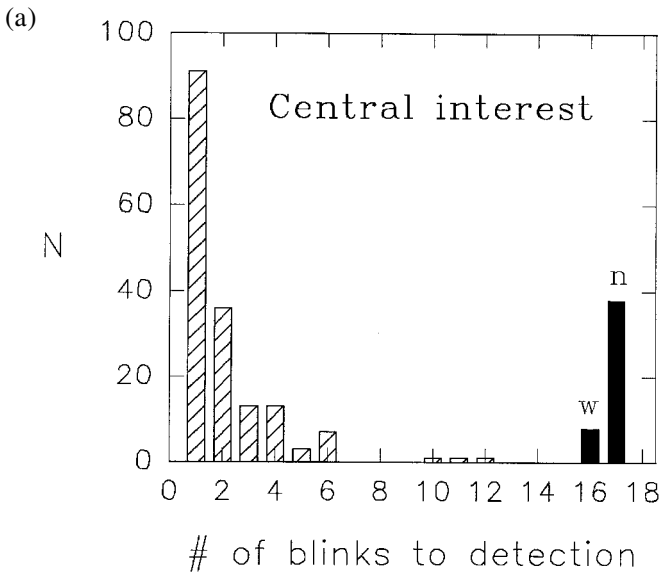


FIG. 1. The number of cases, summed over all subjects, where a change was detected after 1, 2, 3, ... 15 blinks, for (a) Central Interest, and (b) Marginal Interest changes. The columns marked W (wrong) correspond to cases where a subject indicated a change at a location where there had been no change. The N (no change) columns correspond to cases where the subject had still detected no change after 15 blinks or 48 sec, and the computer triggered the change outside a blink.

Analysis of the eye movement trace may allow us to understand why changes are so hard to detect and what is the origin of the difference between the Central and Marginal Interest cases.

Eye Movement Measures

Figure 2 shows a typical eye movement scanning pattern for a picture. It is seen that even though the observer was looking at the picture for 48 sec, and searching actively for possible changes that might occur anywhere in the picture, the eye continued to follow a surprisingly stereotyped, repetitive, scanpath in which large areas of the picture are never directly fixated. Similar observations were made by Yarbus (1967) and other authors, who observed that many portions of a picture are never directly fixated, and that the particular scanpath that is used depends on what the observer is looking for in the picture.

Could this be the reason why some changes are not noticed? Could it be that those cases when the change is missed correspond to cases where the scanpath happens not to include the change location? This hypothesis might explain the difference between the MI and CI changes: Thus, it might be that MI locations, being less "interesting" to observers, tend to be less likely to be included in the scanpath than CI locations.

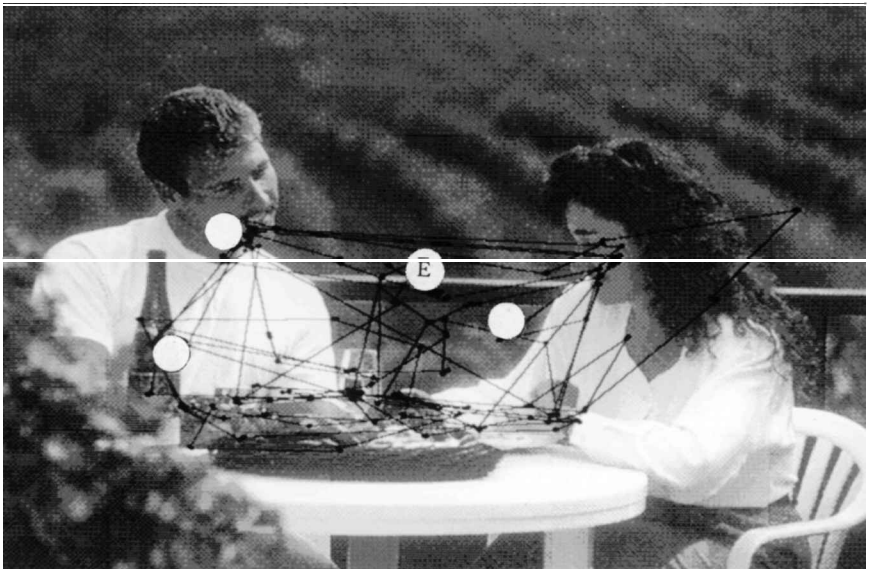


FIG. 2. Typical scanpath while a subject searched for changes. The original picture was in colour. The change that occurred in this picture was a vertical displacement of the railing in the background to the level of the man's eyes. In this record, the change was detected at the moment that the observer blinked for the fourth time. The positions of the eye when the blinks occurred are shown as white circles. The last, "effective" blink, marked "E", occurred when the eye was in the region of the bar.

The data plotted in Figs. 3–5 will help elucidate this question. First of all, we plotted general histograms of the eye movement characteristics in the CI and MI cases. Figure 3 plots fixation durations and saccade amplitude for the CI and MI cases: The curves are identical, and typical of what is known about eye movements in scanning pictures in general. More interesting, to get an idea of the scanpaths and how far they tended to be from CI and MI type change locations, we plotted in Fig. 4 the distance of the eye's fixation from the change location in the CI and MI cases, cumulated over the whole period of viewing of each picture (in the case where the change involved a change in position, the effective change location was taken to be in the middle between the centres of gravity of the initial and final positions). Overall, the CI locations receive more direct fixations than the MI locations, as might be expected from the fact that CI locations were chosen to be more interesting or relevant for interpretation of the picture than MI locations. It is clear therefore that when a blink occurs, the eye is likely to be closer to a CI change location than to a MI change location. This is confirmed in Fig. 5(a), which now restricts the preceding histogram to the moments when the blinks occur. That is, it plots the distance of the eye from the change location at the moment of occurrence of each blink. The data here include all blinks, whether the observer detected the change or not. Again, we see that, at the moment of a blink, the eye's location tended to be closer to the change location when it was a CI location than when it was a MI location.

The probability of the eye being near the change at the moment of its occurrence would thus be distinctly smaller in the MI than in the CI case. However, we shall see that this is not the whole explanation of the effects observed in this experiment. It is true that the probability of detecting MI changes is smaller because the scanpath intersects them less often, but an additional, more interesting factor is also active.

Looking Without Seeing

Figure 5(b) plots a histogram of the eye's location on the “effective” blinks, that is, those blinks that were immediately followed by detection of the change. Dividing the values in Fig. 5(b) by the values in Fig. 5(a), we obtain in Fig. 5(c) the proportion of “effective” blinks, that is, the probability that a blink will give rise to a change detection, as a function of the eye's distance to the change location.

It is difficult to do statistics on data that has been obtained by dividing two values in this way. Furthermore, as can be seen from Fig. 5(b), beyond about 6 or 8 degrees from the change position, only 10 or so effective blinks are contributing to the data in each curve—these cases will correspond to only a small sub-set of the 24 pictures and the 10 subjects for each curve. The data must therefore be considered with caution. Nevertheless, the main aspects of the curves in Fig. 5(c), at least up to eccentricity of 8 degrees, are probably reliable,

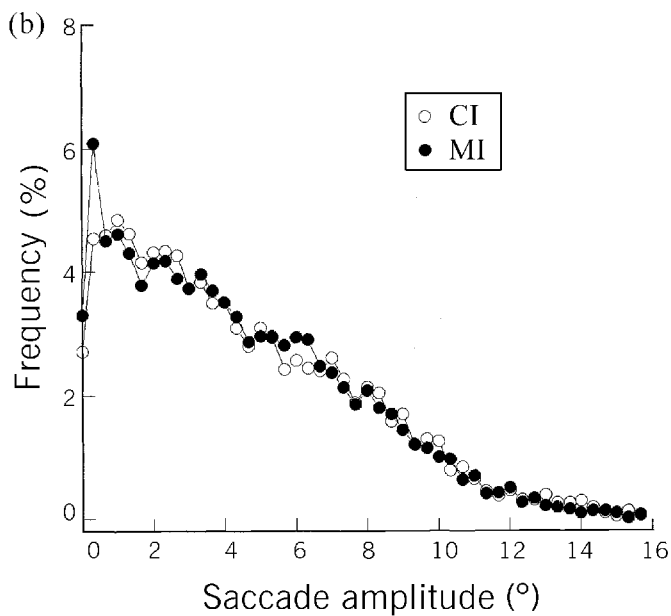
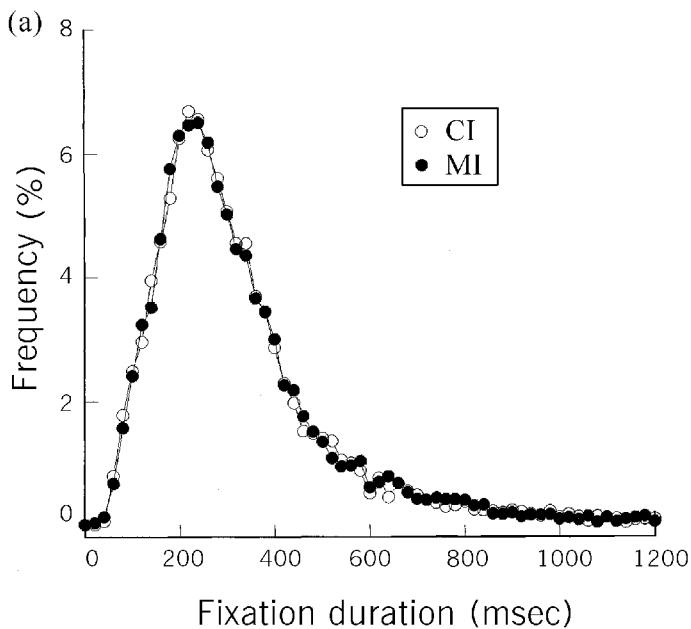


FIG. 3. Histograms of (a) fixation duration and (b) saccade amplitude over all subjects, pictures, and for the whole observation period prior to change detection for Central Interest (CI) and Marginal Interest (MI) cases.

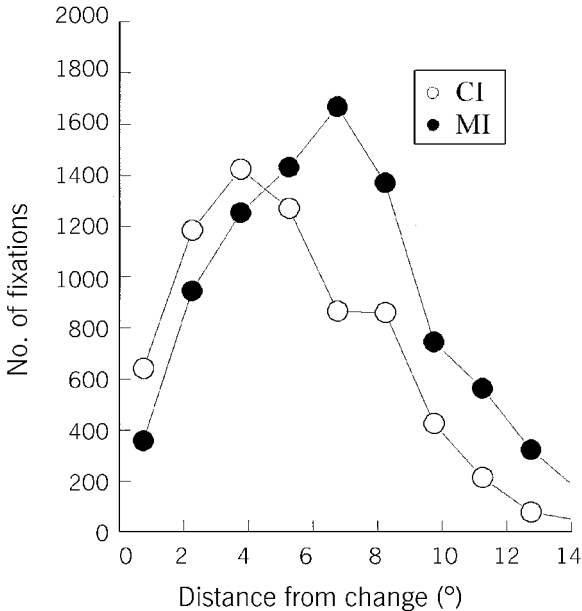


FIG. 4. Distance in degrees of the eye's landing positions from the change location, cumulated over all subjects, pictures, and for the whole observation period prior to change detection for Central Interest (CI) and Marginal Interest (MI) cases.

since they replicate in a reasonable trend for each of the independent data points in the figure. We therefore feel justified in deducing the following interesting points from the figure.

First of all, as would seem reasonable on acuity grounds, the probability of detecting a change drops off rapidly as a function of the eye's distance from the change location. But what is particularly interesting is the fact that at the left-most point of the probability plots, corresponding to the situation when the eye is directly fixating the change at the moment it occurs (within 1 degree), there is less than a 60% chance of detecting it. In other words, in more than 40% of the cases, when the eye is directly fixating the change at the moment it occurs, the change goes undetected. It seems that looking at something is not a guarantee of seeing it: It seems you do not always see where you look.

A second point concerns the visibility of changes when the eye is not directly fixating the change location. We see that at about 2 degrees from the change location, there is a sharp drop in the probability of detecting the change: For an eccentricity of between 2 and 10 degrees, change detection is almost constant at a value of about 0.1. This is the probability of "seeing without looking".

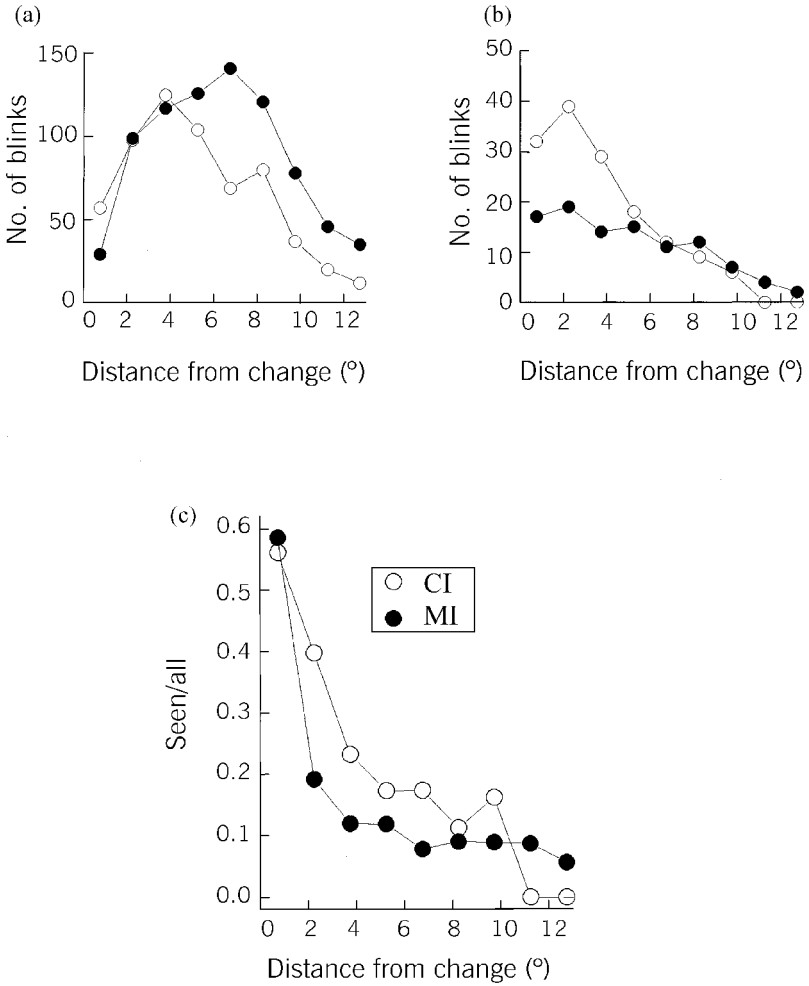


FIG. 5. (a) Distance of the eye's landing positions from the change location, at the moment of the change (i.e. at the blink), cumulated over all subjects, pictures, for Central Interest (CI), and Marginal Interest (MI) cases. (b) The same as (a), but only for those "effective" blinks that were immediately followed by detection of the change. (c) Ratio of effective to all blinks, i.e. proportion of changes that were detected, as a function of the distance of the eye from the change at the moment of the change.

A third point concerns the difference between the probability of recognizing CI and MI changes: For eye-to-change position distances between 2 and 9 degrees, there appears to be a small but consistent difference in the curves in Fig. 5(c): At a given distance between the eye and the change location, CI changes are slightly more likely to be detected than MI changes.

This shows that the scanpath explanation is not the whole story: Even when eccentricity relative to the eye position is controlled for, MI changes are *still* less likely to be detected than CI changes. (Note that since the visual conspicuity of the changes was carefully equalized for the two types of changes, any difference in their detectability can only be due to the meaningfulness of the change within the context of the picture.) It seems that at equal eccentricities, an MI change will overall be less likely to be detected than a CI change, presumably because, by definition, it is not what the observer considers the picture to be about.

A final point concerns MI and CI changes when the eye is fixating the change location at the moment of the blink. Contrary to what happens at other eccentricities, here there appears to be no difference in the probability of detection between MI and CI. Possibly an explanation of this would be as follows. When the eye moves to a MI location, then this location has become, temporarily, the centre of interest. Thus, when the eye is in the 0–1 deg region of an MI location, it effectively constitutes a CI location. A change occurring there should therefore be just as noticeable, independently of whether it is a CI or an MI change.

Eccentricity or Attention?

Perhaps a parsimonious explanation of all the observed effects would dispense entirely with the idea that eccentricity with respect to the eye's fixation point is an important factor in determining the probability of change detection. After all, the picture elements that changed were perhaps sufficiently large to be perfectly distinguishable from a fixation point completely on the opposite side of the screen. Perhaps what looks like an eccentricity effect in our data is actually just another manifestation of an attentional effect.

To be more precise, let us assume, with Posner (1980; Posner, Snyder, & Davidson, 1980), that observers are able to attend covertly to regions that they are not fixating directly. If the changes are visible from any screen location, the change detection task could still have been done, even if the eye had not moved at all. We would then have found no effect at all of the eye's eccentricity from the change location, and the only effect found would have been the "interest" effect, since this will have determined the likelihood of processing a particular picture location and being able to report a change there. However, it is known that eye movements are strongly linked to attention (cf. e.g. Deubel & Schneider, 1996), with attention "pulling" the eye to the attended location. Statistically, therefore, we will observe a tendency for the eye position to be close to locations where changes are detected.

In other words, it is quite possible that in the present experiment the finding of an effect of the eye's position is a *consequence* and not a *cause*: It is a consequence of the fact that a location was being attended. Further work must be

done to disentangle the respective roles of visibility and attention in the present phenomena.

Other Results on Looking Without Seeing

Let us turn now to the finding that the eye was often looking directly at a location and yet the change was not noticed. Though surprising, this finding is not entirely unexpected, given a number of earlier observations. For example, analogously to the ability of listeners to attend to one of two simultaneously audible conversations, Neisser and Becklen (1975) have shown subjects video sequences in which two different game sequences were transparently superimposed. The observers had no trouble attending to one sequence, despite the fact that the other sequence was occupying the same spatial location on the screen. Neisser (1976) interpreted this as evidence that the encoding of visual events involves actively making sense of them: This has the consequence that only those parts of a scene which are being integrated or constructed into a sequence will be noticed. A similar, striking, demonstration was given by Haines (1991), who showed that air pilots landing an aircraft in a flight simulator can be totally oblivious of the presence of a perfectly visible aircraft parked in the middle of the runway. Presumably such an occurrence is essentially inconceivable in a real situation and cannot be integrated into the pilots' momentary conceptual framework. Thus, instead of aborting their landing, the pilots blithely run through the obstructing aircraft.

A wealth of related findings have been demonstrated more recently by Mack and Rock (1998), who refer to what they call "inattention blindness": When observers are intently engaged in one visual task, they will often be unaware of what would normally be a perfectly visible simultaneously occurring event. Though many of Mack and Rock's findings are restricted to cases where the visual event that is missed is *not* at the location being directly fixated, several experiments showed that inattention blindness was surprisingly strong at the very location where the eyes were looking. Mack and Rock concluded that what is important in seeing an event is not the position where the eye is fixating, but what aspects of the scene are being attended to.

Some recent studies by Ballard, Hayhoe, and Pelz (1995) and Hayhoe et al. (1998) also provide evidence, in a block copying task on a computer screen, that looking directly at blocks in a region does not guarantee that observers will notice changes in those blocks. The authors interpret their results in terms of a "deictic" theory in which the eye is not only being used to acquire information, but also as a spatial marker in the scene, analogous to how the finger is used as a pointer in counting targets.

Simons (1996) and Zelinsky (1997, 1998) have also found that in scenes consisting of small collections of objects, changes made to directly attended objects can be missed. Finally, particularly striking demonstrations have been

given by Levin and Simons (1997), who show that, in film sequences, the identity of an actor who is the centre of attention can be changed during a film cut without this being noticed. Simons and Levin (1997), furthermore, refer to real-life situations where the person to whom a subject is talking can be replaced during a brief interruption in the conversation without the subject noticing it. They interpret all these results in terms of the view that not all aspects of even attended parts of a scene are generally encoded. In order to preserve the impression of visual continuity across eye fixations (or, in cinema, across film cuts) only the global layout and the gist of a scene are encoded, and the details are left to be filled in at will by reference to the information available in the visual field. Though giving a special role to layout, this view is very similar to the notion suggested in O'Regan (1992), according to which there is no need to "re"-present the outside world inside the brain, since the outside world can be used as an external memory to be probed when details need to be obtained.

To summarize: In order to explain why in more than 40% of the cases the changes at the location being directly fixated were not detected, it may be useful to consider that when an observer is attending to a particular location in a picture, not all the aspects, or objects, or parts of the picture that share that location are being simultaneously attended to. Some aspects of the picture in the same location may be perfectly visible, but still not processed: These aspects of the picture would then be missed, even though the eye is fixating them directly. This idea relates to the notion that attention might be linked, not to *locations* in the visual field, but to *aspects* of the visual field that have been perceptually grouped, for example into objects (Driver & Baylis, 1989; Duncan, 1984). Furthermore, these ideas are coherent with the notion, inherent in the literature on covert attentional orienting, that the attended location is not necessarily the location being currently fixated (Posner, 1980; Posner et al., 1980; cf. Klein, Kingstone, & Pontefract, 1992).

DISCUSSION

A Theory and Some New Predictions for Change Blindness

In the following paragraphs we will summarize the theory of change blindness which has evolved out of this experiment, as well as from earlier experiments using other paradigms. However, instead of considering conditions in which a scene change coincides with some kind of visual disturbance such as a flicker, blink, saccade, or "mudsplash", it will be useful to begin by considering what happens under *normal* conditions when a change occurs in a scene, that is, *without* imposed, extraneous visual disturbances.

Detecting Changes Under Normal Conditions. When a person looks at a picture, they come to it with their own predispositions and interests that will determine which aspects of the picture they will preferentially process. However, at the same time, low-level visual factors such as contrast, contour complexity, colour distribution, and global layout in the picture may modulate these prior dispositions. Let us suppose that the combined effect of these two factors is such that for a particular picture, over a number of observers, aspect A is most likely to be processed, and that aspects B, C, D, E, F, G, etc. are next most likely to be processed, in decreasing order of likelihood. Aspects at the top of this list correspond to what we have called Central Interest aspects, whereas the other aspects correspond to Marginal Interest aspects of the picture.

Now consider a particular observer on a particular occasion. As the process of recognizing and encoding the picture unfolds, and as the observer's comprehension of the content of the picture builds up, processing develops and shifts from one aspect of the picture to another. Suppose that in this particular case, the observer processes the aspects of the picture in the order: B, A, B, A, B, E, A, C, B, ... and inserts them into their internal representation, which therefore now contains the elements A, B, C, and E.

If a change occurs in the picture, the transient that the change inevitably causes will be picked up by transient detector mechanisms in the observer's low-level visual system, and their attention will be exogenously directed to the location of the transient. If the change concerns a Central Interest aspect such as A, for example, with A changing to A', then the transient will attract attention to A', and, comparison with the internal representation, which contains A, B, C, and E, will be possible. By virtue of the transient, the observer will be able to report *that* a change has occurred, and by virtue of the comparison with the contents of the internal representation, the observer will be able to say *what* the change consisted of, namely a change from A to A'. On the other hand, if the change concerns a Marginal Interest aspect such as G, say, with G changing to G', then low-level transients will direct attention to G', and again, by virtue of the transient, the observer will be able to say *that* a change has occurred. However, since G' is not in the internal representation, the observer will, a priori, not be able to say precisely *what* the change was.

There are now two possibilities. One possibility is that the information contained in the transient itself—we might call this the “flavour” of the transient—will nevertheless suffice to give the observer an idea of what the change was. If, for example, the change comprised a picture element changing position, it is quite likely that the currently processed picture element (G'), plus the information provided by the motion-detectors detecting the transient indicating that there has been a movement, will suffice for the observer to have the impression of seeing the picture element move.

On the other hand, it may happen that the flavour of the transient does not provide sufficient information for the observer to be able to deduce what was at the change location before the change. For example, if element G simply disappears, then the transient may let the observer know that an element has disappeared, but he or she will not know exactly *what* disappeared, since the observer had no internal record of G (we assume retinal or screen remanence provides insufficient information).

In conclusion, for changes occurring as they would normally occur in a scene, that is, without additional, extraneous imposed disturbances, then in the case of Central Interest aspects observers should be able to see the changes without any problem, but in the case of Marginal Interest aspects, in some circumstances observers should have difficulty knowing in detail what was at the change location before the change.

Interestingly, this prediction about seeing changes under normal conditions has not been studied empirically. It may correspond to the "what was that?" situation which one encounters sometimes when one has the impression that something has happened in the visual field, without knowing exactly what it was. Certainly further work would be useful to determine to what extent people miss Marginal Interest changes that take place in full view under normal conditions. Analogously to the results for the present blink experiments, under the present theory it should be possible to make such changes even though the eye is fixating them directly.

Detecting Changes with Concomitant Visual Disturbances. Now consider the situation when a blink, flicker, eye saccade, film cut, or "mudsplash", occurs simultaneously with the scene change. These visual disturbances create visual transients that are spread over the visual field, and which compete attentionally with the transient corresponding to the location where the scene changed. Attention, instead of being directed to the picture aspect that really changed, is likely to be directed to an irrelevant location. The observer must serially scan through each location where a transient is detected, and check whether the element he or she is observing corresponds to what has been stored in the internal representation. The necessity of doing this serial scan is thus a first cause for difficulty in change detection. A second, more important difficulty derives from the fact that change detection must rely now on a comparison process with what has been encoded. If the change corresponds to a Central Interest aspect, comparison will generally be possible because the Central Interest aspect will most likely have been encoded into the internal representation. But if the change concerns a Marginal Interest aspect, that is, an aspect which is likely not to have been encoded into the internal representation, then the observer will not be able to report the change.

This is what is found in all our experiments: Marginal Interest changes are hard to detect. Note that this explanation of the change blindness easily accommodates the results of the “mudsplash” experiments (O'Regan et al., 1999), in which the changing picture item occurs in full view and is not at all masked by a global disturbance of any sort. In these experiments, the local transients corresponding to mudsplashes compete with the local transient corresponding to the true change, and the only way an observer can know that a change has occurred is by comparing what he or she is attending to now, with what has been previously encoded into the internal representation. Since Marginal Interest aspects tend not to be encoded, there is nothing to be compared to, and such changes will tend not to be detected. This is true even though the changes occur in full view.

This explanation of the change blindness effects also explains the “masking rectangle” experiment, where a rectangle both cues the observer to the change location, but simultaneously suppresses the distinctive flavour of the transient (O'Regan et al., 1999). Again, only changes to Central Interest aspects of the picture can be readily detected.

The Moment of Change and Very Slow Changes. There are two interesting further predictions that can be made on the basis of this theory of change blindness. The first prediction concerns the moment at which the change occurs. Suppose again that an observer has inspected the aspects of a picture in the order B, A, B, A, B, E, A, C, B, ... , and has thus encoded elements A, B, C, and E into his or her internal representation. Now encoding may continue and concentrate more on elements C and D, say. Consider the situation where we use a flicker, blink, or saccade paradigm, where global transients prevent any local transient from standing out. If a change were to occur in aspect B of the picture while the observer were concentrating on aspects C and D, even though B is the most Central Interest item of the picture, since it has already been inserted into the internal representation, the observer may simply not return to it to check whether it has changed.

We therefore have the interesting prediction that under some circumstances we may expect that even a Central Interest aspect of the picture may change without being noticed, if this change occurs a sufficiently long time into the period of exploration, so that the observer has already encoded this picture element into his or her internal representation, and is unlikely to return to it.

A second prediction concerns what we expect to happen if conditions are arranged so that changes occur so slowly that they do not provoke visual transients at all. Just as in the classic change blindness studies where the transients are masked by extraneous transients, we expect that under circumstances where changes are so slow that no transient detector signals the location of the change, we should also be able to make large changes in Marginal Interest

locations in full view of observers without these being noticed. Simons (personal communication) has preliminary results confirming this claim.

CONCLUSION

The present experiment has provided further evidence that the cause for the inability to detect changes during eye saccades, observed by a number of researchers over the last years, is most probably linked, not to some specifically saccade-dependent suppression mechanism, but simply to the fact that the visual transient caused by the saccade swamps the transient corresponding to the sought-for change.

Thus, blinks, which produce visual transients similar to those of saccades, though of longer duration, also prevent simultaneously occurring changes from being noticed. This finding is of practical significance, given the high frequency of blinking in everyday life, and given their fairly long duration (> 100 msec).

The paper has also elucidated the role of eye movements in the failure to detect changes. One reason changes are missed may partly be that the eye moves in a scanpath that does not always encompass the change location. The probability of detection may be lowered by the fact that the detection must take place in peripheral vision, where sensitivity may be less. However, this explanation is not the whole story, since even when the eye is directly fixating the change location, an appreciable proportion (more than 40%) of changes are *still* not detected.

Our interpretation of this is that what an observer “sees” at any moment in a scene is not the *location* he or she is directly fixating with the eyes, but the *aspect* of the scene he or she is currently attending to, that is, presumably, what he or she is processing with a view to encoding for storage into memory. In particular, at a given moment, a scene *aspect* may constitute only a sub-set of the elements that are currently being directly fixated, namely a particular sub-set that, when grouped together, corresponds to some meaningful scene entity. Furthermore, global regions or elements of the picture well outside the region of central fixation may also be part of the sub-set of elements that constitute a scene *aspect*, in that for example they may form the setting or backdrop within which more centrally fixated items are inserted. The eye's fixation location will thus in general only be an unreliable indicator of what is being processed.

The discussion section of the paper summarized the theory of change blindness that arises from the present result and previous results in the literature using eye saccades, flicker, “mudsplashes”, “masking rectangles” or film cuts. Two main ideas enable the results to be explained: The idea that the internal representation of the visual field essentially contains only those few aspects of the scene that have been attended to, plus the idea that visual transients caused by picture changes may compete with extraneous transients imposed by

experimental manipulations in attracting attention to possible change locations. Several new predictions came forth from the application of these ideas. One prediction is that even under normal conditions, when no artificially imposed transients occur simultaneously with the display change, we expect that under some circumstances, the exact nature of Marginal Interest changes occurring in full view should be difficult for observers to specify, even though they are aware *that* a change has occurred. A second prediction concerns the moment at which a change occurs: We expect that if a change occurs late in the period of exploration of a scene, it may be missed even if it is a Central Interest change. Finally, we predicted that results similar to the classic change blindness effects should be obtainable by using very slow changes that do not produce visual transients.

REFERENCES

- Ballard, D.H., Hayhoe, M.M., & Pelz, J.B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, *7*, 66–80.
- Collewijn, H., van der Steen, J., & Steinman, R.M. (1985). Human eye movements associated with blinks and prolonged eyelid closure. *Journal of Neurophysiology*, *54*(1), 11–27.
- Currie, C.B., McConkie, G.W., & Irwin, D.E. (1997). Visual stability: Retinal and extra-retinal factors. *Investigative Ophthalmology and Visual Science*, *38*, B307.
- Deubel, H., & Schneider, W.X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Research*, *36*, 1827–1837.
- Driver, J., & Baylis, G.C. (1989). Movement and visual attention: The spotlight metaphor breaks down. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 448–456.
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, *113*, 501–517.
- Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.) *Vancouver studies in cognitive science: Vol 2: Perception*, (pp. 89–110). New York: Oxford University Press.
- Haines, R. (1991). A breakdown in simultaneous information processing. In L. Stark & G. Obrecht (Eds), *IVth international symposium on Presbyopia*. (pp. 171–176). New York: Plenum.
- Hayhoe, M.M., Bensinger, D., & Ballard, D.H. (1998). Task constraints in visual working memory. *Vision Research*, *38*, 125–137.
- Klein, R., Kingstone, A., & Pontefract, A. (1992). Orienting of visual attention. In K. Rayner (Ed.), *Eye movements and visual cognition: Scene perception and reading*. (pp. 46–65). New York: Springer.
- Levin, D.T., & Simons, D.J. (1997). Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin and Review*, *4*, 501–506.
- Mack, A., & Rock, I. (1998). *Inattention blindness*. Cambridge, MA: MIT Press.
- McConkie, G.W. (1991). Perceiving a stable visual world. In J. Van Rensbergen, M. Devijver, & G. d'Ydewalle (Eds), *Proceedings of the sixth european conference on eye movements*. (pp. 5–7). Leuven, Belgium: University of Leuven, Laboratory of Experimental Psychology.
- McConkie, G.W., & Currie, C.B. (1996). Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 563–581.

- Neisser, U. (1976). *Cognition and reality*. San Francisco: Freeman.
- Neisser, U., & Becklen R. (1975). Selective looking: Attending to visually-specified events. *Cognitive Psychology*, 7, 480–494.
- O'Regan, J.K. (1992). Solving the “real” mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*. 46, 461–488.
- O'Regan, J.K., Rensink, R.A., & Clark, J.J. (1996). “Mudsplashes” render picture changes invisible. *Investigative Ophthalmology and Visual Science*. 37, S213.
- O'Regan, J.K., Rensink, R.A., & Clark, J.J. (1999). Change-blindness as a result of “mudsplashes”. *Nature*, 398, 34.
- Posner, M.I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–26.
- Posner, M.I., Snyder, C.R.R., & Davidson, B.J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, 109, 160–174.
- Rensink, R., O'Regan, J.K., & Clark, J.J. (1995). Image flicker is as good as saccades in making large scene changes invisible. *Perception*, 24 (Suppl.), 26–27.
- 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(5), 368–373.
- Rensink, R.A., O'Regan, J.K., & Clark, J.J. (this issue). On the failure to detect changes in scenes across brief interruptions. *Visual Cognition*, 7, 127–145.
- Simons, D.J. (1996). In sight, out of mind: When object representations fail. *Psychological Science*, 7, 301–305.
- Simons, D.J., & Levin, D.T. (1997). Change blindness. *Trends in Cognitive Science*, 1(7), 261–267.
- Sun, W.S., Baker, R.S., Chuke, J.S., Rouholiman, B.R., Hasan, S.A., Gaza, W., Stava, M.W., & Porter, J.D. (1997). Age related changes in human blinks. *Investigative Ophthalmology and Visual Science*, 38(1), 92–99.
- Yarbus, A.L. (1967). *Eye movements and vision*. New York: Plenum Press.
- Zelinsky, G.J. (1997). Eye movements during a change detection search task. *Investigative Ophthalmology and Visual Science*, 38, S373.
- Zelinsky, G.J. (1998). *Detecting changes between scenes: A similarity-based theory using iconic representations*. (CNS Tech. Rep. No. CNS-98-01) Beckman Institute for Advanced Science and Technology, Champaign, IL.